# The Twin Measure for Queueing System Predictability 

David Raz<br>School of Computer Science<br>Tel-Aviv University, Tel-Aviv, Israel<br>davidraz@post.tau.ac.il

Hanoch Levy<br>School of Computer Science<br>Tel-Aviv University, Tel-Aviv, Israel hanoch@cs.tau.ac.il

Benjamin Avi-Itzhak<br>RUTCOR, Rutgers University<br>New Brunswick, NJ, USA<br>aviitzha@rutcor.rutgers.edu


#### Abstract

Two identical jobs with deterministically identical processing times arrive at a Web server simultaneously (Twins), but leave the system thirty seconds apart. Is the service predictable? Is their sojourn time predictable? This issue arises in modern day networking systems such as call centers and Web servers as well as in other queueing systems. We propose a novel measure based on the principle that in a predictable system, "twin" jobs should not depart the system very far apart. We analyze this measure for a number of common scheduling policies and compare the results. We compare the results to those of other predictability approaches proposed recently and discuss its usefulness.


## Categories and Subject Descriptors

C. 4 [Performance of Systems]: Performance AttributesPredictability; F.2.2 [Nonnumerical Algorithms and Problems]: Sequencing and Scheduling; G. 3 [Probability and Statistics]: Queuing Theory

## General Terms

Performance, Measurement

## Keywords

predictability, FCFS, LCFS, SJF, LJF, LAS, SRPT, LRPT, round robin, job scheduling, processor sharing, PS, queue disciplines, twin measure

## 1. INTRODUCTION

How does one measure the predictability of a queueing system?
The importance of the issues of predictability is widely recognized in many works and applications. For example the issue of predictability was recently discussed in [13], which provides references to related work discussing its importance.

[^0]The issue of predictability is also strongly related to the issue of Fairness. If the system isn't predictable, it is most likely since some jobs are treated unfairly. See for example [1] that surveys several studies discussing the subject of fairness, and some recent works quantifying job fairness.
In modern day systems such as call centers and computer systems, where the actual workings of the system are hidden from the customer, one practical way to measure and check for system predictability is to launch a pair of identical jobs with deterministically identical service times (Twins). For example, it is very natural that a Web user who submits two concurrent identical, or close to identical, requests to the same Web site, will appreciate the site predictability based on their relative response times. We propose a measure that is based on this notion. Obviously, a low "twin measure" is not a sufficient requirement for customer satisfaction. In fact, we do not even claim that it is sufficient for guaranteeing system predictability, as this issue is too complicated to be captured by this simple notion. However, we do claim that it is a required feature, and a simple one to measure and analyze, as we will show.
Specifically, we propose a measure based on the expected value, and possibly higher moments, of the difference between the waiting times ,or sojourn times, of identical, simultaneously arriving jobs, conditioned on the jobs' service requirement. By conditioning on the service requirement this measure also captures the notion of a policy which is not equally predictable to all job sizes.
The model and the measure we propose are described in Section 2. Analysis of the twin measure in single server systems, for common scheduling policies, is given in Section 3. We then (Section 4) discuss the measures obtained: in Section 4.1 we classify the scheduling policies into four classes; in Section 4.2 we compare the classification results to those obtained by another recently proposed predictability criterion; in Section 4.3 we discuss what optimality under the twin measure means, and we propose a policy which is optimal with respect to both the twin measure and the sojourn times.
In Section 5 we analyze the twin measure for some common multiple server systems.
Finally, in Section 6 we discuss several ways to extend the twin measure and we conclude with some concluding remarks (Section 7).

## 2. MODEL AND NOTATION

We use an $M / G I / s$ model, namely, a queueing system with $s$ servers, each with one unit of service rate. Arrivals
are Poisson at rate $\lambda$, and the service requirements are sampled independently with probability density function (pdf) $b(x)$ and cumulative distribution function (cdf) $B(x) . B_{x}$ is the probability of a service requirement of exactly $x$, and of course for continuous $b(x), \forall x, B_{x}=0$. However, for many practical distributions there is at least one value for which there is an accumulation of probability and $B_{x} \neq 0$. e.g. the maximum value.
The service requirement has expected value $\bar{x}$, and a second moment $\overline{x^{2}}$. The load (utilization) of the server is defined as

$$
\rho \stackrel{\text { def }}{=} \frac{\lambda \bar{x}}{s}=\frac{\lambda}{s} \int_{0}^{\infty} t b(t) d t,
$$

and for stability we require $\rho<1$.
One useful quantity is the load made up by the jobs of size less than or equal to $x$, denoted $\rho(x)$, which is

$$
\rho(x) \stackrel{\text { def }}{=} \frac{\lambda}{s} \int_{0}^{x} t b(t) d t .
$$

We also define he load made up by the jobs of size strictly less than $x$, denoted $\rho\left(x^{-}\right)$, which is

$$
\rho\left(x^{-}\right) \stackrel{\text { def }}{=} \frac{\lambda}{s} \int_{0}^{x^{-}} t b(t) d t .
$$

Note that for a continuous pdf $b(t), \rho(x)=\rho\left(x^{-}\right)$.
We use the notation $X \sim B P(x)$ to denote that $X$ is distributed as a busy period starting with a job of size $x$. $X \sim B P_{y}(x)$ and $X \sim B P_{y^{-}}(x)$ denote the same, except the busy period is only composed of jobs of size not larger than $y$ and smaller than $y$, respectively.
We use the notation $X \preceq F(t)$ to denote that a random variable $X$ stochastically dominates $F(t)$, i.e. $\mathbb{P}\{X \leq t\} \leq$ $F(t) \quad \forall t \geq 0$. We use the notation $\prec$ in a similar way.
We assume that the server is work conserving, i.e. a preempted job retains the service it was already given.
We use Typrwriter-Style to denote scheduling policies.

### 2.1 The Twin Measure

Let $C_{1}$ and $C_{2}$ be identical jobs, with service requirements $x$ and $x+\delta$ respectively. We call such jobs twins. Let the arrival, departure, and first service epochs of twin $i, i=1,2$ be $a_{i}, d_{i}$, and $s_{i}$ respectively. Assume that $C_{1}$ arrives when $\left(a_{1}\right)$ the system is in steady state. Assume that the twins arrive $\epsilon$ time units apart, that is, $a_{2}-a_{1}=\epsilon, \epsilon>0$.

Definition 2.1 (Twin Measure).
Define the random variable $Z(x, \epsilon, \delta)=\left|d_{2}-d_{1}\right|$, given $x, \epsilon$, and $\delta$. Let $z^{n}(x, \epsilon, \delta)$ be the $n$-th moment of $Z(x, \epsilon, \delta)$, i.e. $z^{n}(x, \epsilon, \delta)=\mathbb{E}\left\{Z(x, \epsilon, \delta)^{n}\right\}$.
For scheduling policy $\phi$ and job size $x$, the $n$-th twin measure $\mathcal{T}_{n}^{\phi}(x)$ is defined as the limit, when $\epsilon$ and $\delta$ tend to zero, of $z^{n}(x, \epsilon, \delta)$, assuming a single limit exists. Namely $\mathcal{T}_{n}^{\phi}(x)=\lim _{\epsilon \rightarrow 0, \delta \rightarrow 0} z^{n}(x, \epsilon, \delta)$.

The shortened term twin measure, denoted $\mathcal{T}^{\phi}(x)$, is used to describe the first twin measure, namely $\mathcal{T}^{\phi}(x)=\mathcal{T}_{1}^{\phi}(x)=$ $\lim _{\epsilon \rightarrow 0, \delta \rightarrow 0} \mathbb{E}\{Z(x, \epsilon, \delta)\}$. While we will focus in this paper on the first twin measure, we find it useful to define it in the scope of higher moments, for future research.

Remark 2.1. The $n$-th twin measure is only defined when a single limit exists for both $\delta \searrow 0$ and $\delta \nearrow 0$, while $\epsilon \rightarrow 0$.

One can devise a policy for which there is no such single limit. For example, consider a policy $\phi$ that serves jobs in a First-Come-First-Served (FCFS) manner, unless the second job in the queue has a smaller size than the first one, in which case it servers the first two jobs in a Processor Sharing (PS) manner. In this case for $\delta \searrow 0$ the twins are served in a PS manner and $\mathcal{T}^{\phi}(x)=0$, while for $\delta \nearrow 0$ the twins are served in a FCFS manner, and $\mathcal{T}^{\phi}(x)=x$. The twin measure in such a case can be chosen to be the maximum of all the limits, the mean value, or it can remain undefined, as appropriate for the application.

Definition 2.1 has the benefit that it applies to service distributions for which same size arrivals are impossible. It also avoids a pitfall that a policy $\phi$ can artificially serve equal sized jobs in a different manner than non-equal sized ones. As in some size distributions equal sized jobs are extremely rare, this will not hinder the expected performance of $\phi$, yet allow it to have an artificially low twin measure. However, this definition makes the analysis tedious so for the sake of analysis we use a simpler definition:

Definition 2.2 (Simplified Twin Measure).
Let $\delta=0$, i.e. both jobs have the same size. Define the random variable $Z(x, \epsilon)=\left|d_{2}-d_{1}\right|$, given $x$ and $\epsilon$. Let $z^{n}(x, \epsilon)$ be the $n$-th moment of $Z(x, \epsilon)$, i.e. $z^{n}(x, \epsilon)=\mathbb{E}\left\{Z(x, \epsilon)^{n}\right\}$.
For scheduling policy $\phi$ and job size $x$, the $n$-th twin measure $\mathcal{T}_{n}{ }^{\phi}(x)$ is defined as the limit, when $\epsilon$ tend to zero, of $z^{n}(x, \epsilon)$, assuming a limit exists. Namely $\mathcal{T}_{n}^{\phi}(x)=\lim _{\epsilon \rightarrow 0}$ $z^{n}(x, \epsilon)$.

The shortened term twin measure, is again used to describe the first twin measure, namely $\mathcal{T}^{\phi}(x)=\lim _{\epsilon \rightarrow 0} \mathbb{E}\{Z(x, \epsilon)\}$.
Whether we use the original or the simplified definition will be clear from the context.
One can also choose to normalize the twin measure, the obvious normalization factor being $x$. This does not change the results in any significant manner.

## 3. ANALYSING COMMON SCHEDULING POLICIES FOR SINGLE SERVER SYSTEMS

### 3.1 Processor Sharing (Ps)

When $C_{1}$ departs, $C_{2}$ can have at most $\epsilon+\delta$ remaining service time. Therefore $z(x, \epsilon, \delta) \leq(\epsilon+\delta) \bar{N}$ where $\bar{N}$ is the mean number of jobs in the system. Thus

$$
\mathcal{T}^{P S}(x)=\lim _{\epsilon \rightarrow 0, \delta \rightarrow 0} z(x, \epsilon, \delta) \leq \lim _{\epsilon \rightarrow 0, \delta \rightarrow 0}((\epsilon+\delta) \bar{N})=0
$$

Intuitively, both jobs are identical and arrive simultaneously, so they will receive exactly the same service, and leave the system simultaneously.

### 3.2 First Come First Served (fffs)

Let $W(x, \epsilon, \delta)$ be a random variable denoting the time elapsing between $d_{1}$ and $s_{2}$. For any non-preemptive scheduling policy we have $Z(x, \epsilon, \delta)=W(x, \epsilon, \delta)+x+\delta$. For FCFS $W(x, \epsilon, \delta)$ is the amount of work arriving in the interval between the arrival epochs of the twins and it has an expected value of $\rho \epsilon$, i.e. $z(x, \epsilon, \delta)=\mathbb{E}\{Z(x, \epsilon, \delta)\}=\rho \epsilon+x+\delta$ and

$$
\mathcal{T}^{F C F S}(x)=\lim _{\epsilon \rightarrow 0, \delta \rightarrow 0}(\rho \epsilon+x+\delta)=x
$$

### 3.3 Last Come First Served (LCFs)

We start with the non-preemptive case. There are two possible orders of service, either (i) $C_{1}$ is served before $C_{2}$, i.e. either the server was idle on $a_{1}$, or the server finished serving the job that was served on $a_{1}$ at some epoch in the interval $\left[a_{1}, a_{2}\right)$, and no other job arrived between $a_{1}$ and that epoch, or (ii) $C_{2}$ is served first.
(i) We can ignore all jobs served before $C_{1}$. The service order from then onward is as follows. First $C_{1}$ is served for $x$ units of time. $C_{2}$ then waits for a busy period created by all jobs arriving while $C_{1}$ was served, and is served at the completion of this busy period, for $x+\delta$ units of time. Let $V(x, \epsilon, \delta)$ be a random variable denoting the time elapsing between $s_{1}$ and $s_{2}$. Thus $Z(x, \epsilon, \delta)=V(x, \epsilon, \delta)+\delta$. Clearly $V(x, \epsilon, \delta) \sim B P(x)$, thus

$$
z(x, \epsilon, \delta)=\frac{x}{1-\rho}+\delta
$$

For example, this can be derived from the transform of the distribution of a busy period starting with a job of size $x$, $G^{*}(s, x)=e^{-x\left[s+\lambda-\lambda G^{*}(s)\right]}$, where $G^{*}(s)$ is the transform of the distribution of the busy period length (see e.g. [5, p. 212]).

Taking the limit we have

$$
\mathcal{T}^{L C F S}(x)=\lim _{\epsilon \rightarrow 0, \delta \rightarrow 0}\left(\frac{x}{1-\rho}+\delta\right)=\frac{x}{1-\rho}
$$

(ii) We can ignore all jobs served before $C_{2}$. The service order from then onward is as follows. First $C_{2}$ is served for $x+\delta$ units of time. Then a busy period created by all jobs arriving while $C_{2}$ was served is being served. Then a busy period created by all jobs arriving between $a_{1}$ and $a_{2}$ is served, followed by $C_{1}$ being served for $x$ units of time. Thus $Z(x, \epsilon, \delta)=V(x, \epsilon, \delta)-\delta$, where $V(x, \epsilon, \delta) \sim B P(x+\delta+\epsilon)$. Therefore

$$
\begin{aligned}
z(x, \epsilon, \delta) & =\frac{x+\delta+\epsilon}{1-\rho}-\delta \\
\mathcal{T}^{L C F S}(x) & =\lim _{\epsilon \rightarrow 0, \delta \rightarrow 0}\left(\frac{x+\delta+\epsilon}{1-\rho}-\delta\right)=\frac{x}{1-\rho}
\end{aligned}
$$

For the preemptive case, all jobs arriving in the interval $\left[a_{1}, a_{2}\right)$, including $C_{1}$, are served in total for a period of length $\epsilon$ until they are preempted by $C_{2} . C_{2}$ and any preempting jobs are then served until $d_{2}$, and can be ignored. The period $\left(d_{2}, d_{1}\right)$ is composed of $C_{1}$ and jobs arriving in the interval $\left(a_{1}, a_{2}\right)$, and jobs preempting them, minus a service of $\epsilon$ units of time already done. Therefore $Z(x, \epsilon, \delta) \sim B P(x+\epsilon)-\epsilon$, so

$$
\mathcal{T}^{L C F S}(x)=\lim _{\epsilon \rightarrow 0, \delta \rightarrow 0}\left(\frac{x+\epsilon}{1-\rho}-\epsilon\right)=\frac{x}{1-\rho}
$$

To summarize, for both the non-preemptive and the preemptive case

$$
\mathcal{T}^{L C F S}(x)=\frac{x}{1-\rho}
$$

### 3.4 Shortest Job First (SJF)

The analysis for SJF is quite tedious if one uses Definition 2.1. Therefore, the analysis provided here uses Definition 2.2 , i.e. we assume the twin jobs are of equal size. One can verify that this simplification does not alter the result as well. We will use this definition from this point onwards.

We start with the non-preemptive case.
Note that in general, SJF does not determine the order of service between equally sized jobs. Obvious choices are either FCFS or LCFS. We call these policies SJF-FCFS and SJF-LCFS respectively.

Starting with SJF-FCFS, $C_{1}$ is always served first. We can ignore all jobs served before $C_{1}$. The service order from then onwards is as follows. First $C_{1}$ is served for $x$ units of time. This is followed by a busy period composed of jobs of sizes smaller than $x$ arriving while $C_{1}$ was served. Following this, jobs of size $x$ arriving in the interval $\left(a_{1}, a_{2}\right)$ are served, followed by a busy period composed of jobs of size smaller than $x$ arriving while they were served. Lastly, $C_{2}$ is served for $x$ units of time. If we let $V(x, \epsilon)$ be a random variable denoting the time elapsing between $s_{1}$ and $s_{2}$ it is easy to see that $Z(x, \epsilon)=V(x, \epsilon)$. Considering the service order we have $V(x, \epsilon) \sim B P_{x^{-}}(x)+B P_{x^{-}}\left(\epsilon \lambda x B_{x}\right)$, and therefore

$$
\begin{align*}
\mathcal{T}^{S J F-F C F S}(x)=\lim _{\epsilon \rightarrow 0}\left(\frac{x}{1-\rho\left(x^{-}\right)}+\right. & \left.\frac{\epsilon \lambda x B_{x}}{1-\rho\left(x^{-}\right)}\right) \\
& =\frac{x}{1-\rho\left(x^{-}\right)} \tag{1}
\end{align*}
$$

For SJF-LCFS, note that $C_{1}$ can still be served first, e.g. if the server is idle on $a_{1}$. Therefore, let $C_{f}$ and $C_{s}$ be the first and second twins to be served, respectively. The order of service starts with $C_{f}$, followed by a busy period composed of jobs of sizes not larger than $x$ arriving while $C_{f}$ was served. Following this, jobs of size $x$ arriving in the interval $\left(a_{1}, a_{2}\right)$ are served, followed by a busy period composed of jobs of size not larger than $x$ arriving while they were served. Lastly, $C_{s}$ is served for $x$ units of time. Using the same notation $Z(x, \epsilon)=V(x, \epsilon)$ and $V(x, \epsilon) \sim$ $B P_{x}(x)+B P_{x}\left(\epsilon \lambda x B_{x}\right)$, leading to

$$
\begin{array}{r}
\mathcal{T}^{S J F-L C F S}(x)=\lim _{\epsilon \rightarrow 0}\left(\frac{x}{1-\rho(x)}+\frac{\epsilon \lambda x B_{x}}{1-\rho(x)}\right) \\
=\frac{x}{1-\rho(x)} \tag{2}
\end{array}
$$

Note that for continuous service distributions SJF-FCFS and SJF-LCFS have the same twin measure.
We now move on to the preemptive case. Again we consider SJF-FCFS and SJF-LCFS.

For SJF-FCFS $C_{1}$ is served first. When $C_{1}$ finishes service there are no jobs of size smaller than $x$ in the queue. The next to be served are jobs of size $x$ arriving in $\left(a_{1}, a_{2}\right)$, followed by $C_{2}$. Each of those can be interrupted, but only by jobs of size smaller than $x$. Therefore $Z(x, \epsilon) \sim$ $B P_{x^{-}}\left(\epsilon \lambda x B_{x}\right)+B P_{x^{-}}(x)$ and (1) holds.

For SJF-LCFS $C_{1}$, and other jobs of size $x$ arriving in $\left(a_{1}, a_{2}\right)$, can be served for a total service no loner than $\epsilon$ before being preempted by $C_{2}$. Let $\beta(\epsilon)$ be a random variable denoting this amount of service. When $C_{2}$ finishes service the next to be served are jobs of size $x$ arriving in $\left(a_{1}, a_{2}\right)$, followed by $C_{1}$. Each of those can be interrupted by jobs of size not larger than $x$. Therefore $Z(x, \epsilon) \sim B P_{x}\left(\epsilon \lambda x B_{x}\right)+B P_{x}(x)-\beta(\epsilon)$. As $0 \leq \beta(\epsilon) \leq \epsilon$ we have $\lim _{\epsilon \rightarrow 0} \beta(\epsilon)=0$ and (2) holds.

### 3.5 Longest Job First (LJf)

We define LJF-FCFS and LJF-LCFS in a similar manner to the ones defined for SJF.

Using the same arguments as in Section 3.4, and the fact
that

$$
\begin{aligned}
& \int_{x^{+}}^{\infty} t f(t) d t=\rho-\rho(x) \\
& \int_{x}^{\infty} t f(t) d t=\rho-\rho\left(x^{-}\right)
\end{aligned}
$$

we have

$$
\begin{aligned}
\mathcal{T}^{L J F-F C F S}(x) & =\frac{x}{1-(\rho-\rho(x))} \\
\mathcal{T}^{L J F-L C F S}(x) & =\frac{x}{1-\left(\rho-\rho\left(x^{-}\right)\right)}
\end{aligned}
$$

This applies to both the non-preemptive and the preemptive case.

### 3.6 Least Attained Service (Las)

In the LAS scheduling policy, also called FB, service is given to the jobs which received the least service so far. See [8] for a survey of results regarding this policy.

Note that in the LAS scheduling policy jobs with equal attained service share the processor. Upon arrival, $C_{1}$ will be served for at most $\epsilon$, then $C_{2}$ will be served for an equal amount, and from then on they will have equal attained service, and keep sharing the processor, until they leave the system together. Therfore

$$
\mathcal{T}^{L A S}(x)=0
$$

### 3.7 Shortest Remaining Processing Time (SRPT)

Note that since both jobs start with equal jobs sizes, from the first epoch in which one of them is served, the other will not be served until the first one leaves the system. It can easy to observe that the twin measure is the same no matter which of the jobs is first served. We will therefore assume w.l.g. that $C_{1}$ is served first. Following [11], $Z(x, \epsilon)$ can be decomposed into the sum

$$
Z(x, \epsilon)=W(x)^{S R P T}+R(x)^{S R P T}
$$

where $W(x)^{S R P T}$ is a random variable denoting the waiting time for $C_{2}$ of size $x$, i.e. the time from $d_{1}$ to $s_{2}$, and $R(x)^{S R P T}$ is a random variable denoting the residence time for $C_{2}$ of size $x$, i.e. the time from $s_{2}$ to $d_{2}$. Both $W(x)^{S R P T}$ and $R(x)^{S R P T}$ do not depend on $\epsilon$ and therefore

$$
\begin{align*}
\mathcal{T}^{S R P T}(x)=\lim _{\epsilon \rightarrow 0} & \left(W(x)^{S R P T}+R(x)^{S R P T}\right) \\
& =\mathbb{E}\{W(x)\}^{S R P T}+\mathbb{E}\{R(x)\}^{S R P T} \tag{3}
\end{align*}
$$

Starting with $\mathbb{E}\{W(x)\}^{S R P T}$, note that once $C_{1}$ enters service, we can divide the arriving jobs into three categories: 1) jobs with service requirement over or equal to $x$. These jobs are served after $d_{2}$ and therefore can be ignored. 2) jobs with service requirement below the remaining service requirement of $C_{1}$ on the epoch of their arrival. These jobs will preempt $C_{1}$ and be served before $d_{1}$. 3) jobs with service requirement lower than $x$, but above the remaining service requirement of $C_{1}$ on the epoch of their arrival. These jobs are served in the interval $\left(d_{1}, s_{2}\right)$.

To carry out this analysis, observe jobs arriving in an infinitesimal interval of size $d t$ when $C_{1}$ has remaining service requirement $t$. Category 2) jobs preempt $C_{1}$ and create a sub busy period of size $d t /\left(1-\rho\left(t^{-}\right)\right)$(including the initial
$d t$ interval). Category 3) jobs arriving in this sub busy period are to be served after $d_{1}$, and the work load created by these jobs is

$$
\frac{d t}{1-\rho\left(t^{-}\right)} \lambda \int_{t^{+}}^{x^{-}} y b(y) d y=\frac{\left(\rho\left(x^{-}\right)-\rho(t)\right) d t}{1-\rho\left(t^{-}\right)}, 0 \leq t<x
$$

and zero for $t=x$.
Note that any job arriving while these jobs are served, with service requirement below $x$, will also be served before $s_{2}$. Therefore we are facing a busy period of size

$$
\frac{\left(\rho\left(x^{-}\right)-\rho(t)\right) d t}{1-\rho\left(t^{-}\right)} \frac{1}{1-\rho\left(x^{-}\right)}
$$

Integrating this yields the mean size of the waiting interval, namely

$$
\mathbb{E}\{W(x)\}^{S R P T}=\int_{0}^{x^{-}} \frac{\rho\left(x^{-}\right)-\rho(t)}{\left(1-\rho\left(t^{-}\right)\right)\left(1-\rho\left(x^{-}\right)\right)} d t
$$

As for $\mathbb{E}\{R(x)\}^{S R P T}$, this is simply the residence time of a job with service requirement $x$ under SRPT. This is true since like a regular job, jobs already in the system once $C_{2}$ begins service are guaranteed to have remaining processing time over $x$, and therefore will not affect the residence time. Thus,

$$
\mathbb{E}\{R(x)\}^{S R P T}=\int_{0}^{x^{-}} \frac{d t}{1-\rho\left(t^{-}\right)}
$$

Using (3) we get

$$
\begin{aligned}
& \mathcal{T}^{S R P T}(x)= \\
& \qquad \begin{aligned}
\left.\int_{0}^{x^{-}} \frac{\rho\left(x^{-}\right)-\rho(t)}{\left(1-\rho\left(t^{-}\right)\right)( } 1-\rho\left(x^{-}\right)\right)
\end{aligned} t+\int_{0}^{x^{-}} \frac{d t}{1-\rho\left(t^{-}\right)} \\
& =\int_{0}^{x^{-}} \frac{1-\rho(t)}{\left(1-\rho\left(t^{-}\right)\right)\left(1-\rho\left(x^{-}\right)\right)} d t
\end{aligned}
$$

which for a continuous pdf $b(t)$ is simply

$$
\mathcal{T}^{S R P T}(x)=\frac{x}{1-\rho(x)}
$$

### 3.8 Longest Remaining Processing Time (LRPT)

Under LRPT all jobs leave the system at the end of the busy period in which they arrive. Specifically, $C_{1}$ and $C_{2}$ leave the system simultaneously, leading to

$$
\mathcal{T}^{L R P T}(x)=0
$$

### 3.9 Round Robin (RR)

We analyze the RR policy with service quantum $\Delta$ where $\Delta \ll x$, and for simplicity we assume that service times are multiples of $\Delta$. We use a model quite similar to the one described in $[6, \operatorname{Sec} 4.4]$, except that newly arriving jobs join the queue after the last arriving job.

As $\Delta \ll x$ we can be certain that even if $C_{1}$ begins service before $a_{2}$, it cannot finish a single quantum before $a_{2}$. Therefore the only jobs between $C_{1}$ and $C_{2}$ in the queue are jobs arriving in the interval $\left(a_{1}, a_{2}\right)$ and $C_{1}$ will never be more than one service cycle ahead of $C_{2}$. Let $N(\epsilon)$ be a random variable denoting the number of jobs arriving or in the interval $\left(a_{1}, a_{2}\right)$. Some of these jobs will have shorter service requirement than $x$, and therefore $\Delta \leq Z(x, \epsilon) \leq(N(\epsilon)+1) \Delta$.

However, note that $N(\epsilon)$ has an expected value is $\lambda \epsilon$, and therefore $\Delta \leq y(x, \epsilon) \leq \Delta+\lambda \epsilon$ leading to

$$
\begin{equation*}
\mathcal{T}^{R R}(x)=\lim _{\epsilon \rightarrow 0} y(x, \epsilon)=\Delta \tag{4}
\end{equation*}
$$

Note that (4) holds also for other models of RR, e.g. when arriving jobs join the queue in other positions, though the analysis in some cases is somewhat more complicated.

## 4. DISCUSSION ON THE TWIN MEASURE OF SCHEDULING POLICIES

In this section we discuss the twin measures obtained in the previous sections. We start with proposing a classification and comparing the policies in each class. We then compare this classification to the classification provided by [13, 12] which we call the Conditional Response Time Criterion. We finalize with a discussion of the meaning of optimality under the twin measure.

For simplicity, we use the measures obtained for continuous service distributions. This is also convenient since [13, 12] deals mainly with continuous service distributions (although the results can probably be extended to the noncontinuous case).

### 4.1 Classifying the Scheduling Policies

Definition 4.1 (Twin Measure Classification). A scheduling policy $\phi$ will be called Absolutely Twin Predictable if $\mathcal{T}^{\phi}(x)=0$ for every $x$.

A scheduling policy $\phi$ will be called Strongly Twin Predictable if $\mathcal{T}^{\phi}(x) \leq x$ for every $x$.

A scheduling policy $\phi$ will be called Weakly Twin Predictable if $\mathcal{T}^{\phi}(x)<x /(1-\rho)$ for every $x$.

A scheduling policy $\phi$ will be called Not Twin Predictable if $\mathcal{T}^{\phi}(x) \geq x /(1-\rho)$ for at least one value of $x$.

The reason we use this classification will be made clear once we discuss the scheduling policies within each class.

### 4.1.1 Absolutely Twin Predictable

This class includes all policies for which twin jobs will leave the system simultaneously, namely PS, LAS and LRPT.

This clearly demonstrates that the twin measure has very little to do with the efficiency aspects of performance, as LAS is obviously much more efficient than PS.

It also demonstrates that the twin measure is not sufficient to guarantee sojourn time predictability. For example LRPT has notoriously unpredictable sojourn times.

### 4.1.2 Strongly Twin Predictable

This class includes policies for which the twin measure for a job of size $x$ isn't larger than $x$. Observe that the minimum twin measure one can expect from any non-preemptive scheduling policy is $x$. One can see that FCFS achieves this optimal value for non-preemptive scheduling policies.

The other policy in this class is RR which in the limit $\Delta \rightarrow 0$ becomes PS.

### 4.1.3 Not Twin Predictable

LCFS has the largest measure analyzed, $x /(1-\rho)$. We conjecture that LCFS had the largest twin measures amongst non-preemptive scheduling policies. Note that the difference between the conjectured best and worst non-preemptive
scheduling policies, FCFS and LCFS, can be extreme in cases where $\rho \rightarrow 1$.

### 4.1.4 Weakly Twin Predictable

Policies in this class include SJF, LJF and SRPT.
For small jobs, such that $\rho(x)<\rho-\rho(x)$, SJF has the lower measure, and the opposite for long jobs.

Interestingly, SJF and SRPT have the same measure. This can be explained in the following way. We call a job intervening if the job gets served in the interval $\left(d_{1}, d_{2}\right)$, or $\left(d_{2}, d_{1}\right)$ if $C_{2}$ is served first, and is not $C_{1}$ or $C_{2}$. Observe that under SJF a job of size $x$ is only intervened by jobs of size $y<x$ arriving in a period of size $x$. In the nonpreemptive case this period is the period in which the first job is served, while in the preemptive case this period is the period in which the second job is served.

Now consider SRPT. Observe a period of time of length $d t$ in which the first job to be served was already served for $t$ units of time. Jobs intervening in this period are jobs with size $t \leq y<x$. Observe a period of time of length $d t$ in which the second job to be served was already served for $t$ units of time. Jobs intervening in this period are jobs with size $y<t$. So in total, intervening jobs for every such interval $d t$ are of size $y<t$. Now consider that the total length of such intervals of length $d t$ is $x$, so the jobs intervening a job of size $x$ under SRPT are also jobs of size $y<x$ arriving in a period of size $x$.
This is an interesting observation as it shows a similarity un-observed before between SRPT and SJF.

### 4.2 Comparing Predictability Criteria

In this section we compare the results of the twin measure classification with that of the Conditional Response Time Criterion. The Conditional Response Time Criterion was proposed in [13]. Further results are provided in [12], specifically for the case $\mathbb{E}\left\{X^{3}\right\}=\infty$.


Figure 1: Classification According to the Conditional Response Time Criterion

We start by summarizing the Conditional Response Time Criterion in the settings and notation of this work.

Definition 4.2. A job of size $x$ is treated predictably under policy $\phi$, service with $p d f b(x)$, and load $\rho$ if the conditional variance in response time seen by a job of size $x$ under
policy $\phi, \operatorname{Var}\{\mathrm{T}(\mathrm{x})\}^{\phi}$, follows

$$
\frac{\operatorname{Var}\{\mathrm{T}(\mathrm{x})\}^{\phi}}{x} \leq \frac{\lambda \overline{x^{2}}}{(1-\rho)^{3}} .
$$

A scheduling policy $\phi$ is predictable if every job size is treated predictably.

A scheduling policy $\phi$ is: (i) Always Predictable if $\phi$ is predictable under all loads and service distributions; (ii) Sometimes Predictable if $\phi$ is predictable under some loads and service distributions; and unpredictable under other loads and service distributions or (iii) Always Unpredictable if $\phi$ is unpredictable under all loads and service distributions.

Figure 1 summarizes the results of classifying common scheduling policies according to the Conditional Response Time Criterion.

| Policy | Twin | Response Time |
| :--- | :---: | :---: |
| PS | Absolutely | Always |
| LAS | Absolutely | Sometimes |
| LRPT | Absolutely | Never |
| FCFS | Strongly | Sometimes |
| NP-LCFS | Not | Never |
| P-LCFS | Not | Always |
| NP-SJF | Weakly | Sometimes |
| P-SJF | Weakly | Sometimes |
| NP-LJF | Weakly | Never |
| P-LJF | Weakly | Never |
| SRPT | Weakly | Sometimes |

Table 1: Comparing Predictability Criteria. The column "Twin" lists the class according to the Twin Measure Classification. The column "Response Time" lists the class according to the Conditional Response Time Criterion.

Table 1 compares the classification of policies according to the two criteria. A scheduling policy name starting with NP denotes the non-preemptive variant of the scheduling policy, while a name starting with $P$ denoted the preemptive variant. Note that for continuous service distributions, our analysis has shown same results for preemptive and non-preemptive policies, which is not the case with the Conditional Response Time Criterion.

Only policies for which analysis was provided in both works are listed.

For some policies the two criteria agree. For example PS is both Absolutely Twin Predictable and Always Predictable. NP-LCFS is both Not Twin Predictable and Never Predictable. However, for some policies the criteria totally disagree. LRPT is Absolutely Twin Predictable, yet Never Predictable. $P-L C F S$ is Not Twin Predictable, yet Always Predictable. LAS is Absolutely Twin Predictable and Sometimes Predictable, but this predictability is only in the case $\mathbb{E}\left\{X^{3}\right\}=\infty$ (see [12]), so for most service distributions the criteria disagree.

This dissimilarity suggests that for guaranteeing predictability one might want to combine the two criteria, or require both.

### 4.3 Optimality under the Twin Measure

As several policies have zero twin measure, these policies can all be considered twin measure wise optimal. One might therefore want to find a policy that is both optimal in the twin measure sense, and has low sojourn time. Consider the following variant of SRPT, called SRPT ${ }_{\alpha}$. Assume the job with the shortest remaining processing time has remaining processing time $x$. Service is given in a processor sharing manner, to all jobs with remaining processing time not larger than $x+\alpha$. It is easy to see that if $\alpha$ can be arbitrarily small, this policy's sojourn times are arbitrarily close to SRPT. On the other hand, if $\alpha \geq \epsilon+\delta$ (which since $\epsilon \rightarrow 0, \delta \rightarrow 0$ can be done for arbitrarily small $\alpha$ ), twin jobs will be served in processor sharing manner from $a_{2}$ until $\min \left(d_{1}, d_{2}\right)$, at which point one of them will leave the system and the other one will have at most $\epsilon+\delta$ service left. As $\epsilon \rightarrow 0, \delta \rightarrow 0$, this job will receive full service from that epoch onwards, and thus $Z(x, \delta, \epsilon) \leq \epsilon+\delta$ and $\mathcal{T}^{S R P T_{\alpha}}(x)=0$. Thus $\operatorname{SRPT}_{\alpha}$ has both optimal twin measure, and optimal sojourn time.

## 5. THE TWIN MEASURE IN MULTI-SERVER SYSTEMS

In this section we analyze the twin measure in some common multi-server settings. The scheduling policy between members of the same queue is FCFS.

### 5.1 Single Queue

We start our analysis with the simple single queue system, denoted SingleQueue, where the first job in the queue is served by the first server to become idle.

For simplicity we ignore effects caused by jobs arriving in the interval $\left(a_{1}, a_{2}\right)$ in this analysis. As the scheduling policy is FCFS, it is easy to see that these effects would be negligible as $\epsilon \rightarrow 0$.

A simple observation is that the twin measure is smaller than $x+\epsilon$, or $x$ when $\epsilon \rightarrow 0$. This is so because at the worse case (i) $C_{2}$ will be served right after $C_{1}$. However, there is a probability that either (ii) both twins arrive when two or more servers are idle, in which case both twins are served immediately, and leave the system simultaneously, or (iii) some other server than the one serving $C_{1}$ will become idle while $C_{1}$ is being served, in which case $C_{2}$ will be served partially in parallel to $C_{1}$, and leave system before $d_{2}+x$. Letting $\epsilon \rightarrow 0$ and taking expectations we have

$$
\begin{equation*}
\mathcal{T}^{\text {SingleQueue }}(x)=(1-\alpha)\left(\int_{0}^{x} \beta(y) y d y+\gamma(x) x\right) \tag{5}
\end{equation*}
$$

where $\alpha$ is the probability of arriving when two or more servers are idle. Given that no more than one server was idle on arrival, $\beta(y)$ is the probability that $C_{2}$ will be delayed for $y$ units of time until another server is idle, and $\gamma(x)$ is the probability that no other server will be idle until $C_{1}$ is served, in which case $C_{2}$ is delayed for exactly $x$ units of time.

For $\alpha$ we have

$$
\alpha=\sum_{k=0}^{s-2} p_{k}
$$

where $p_{k}$ is the probability of finding $k$ jobs in the system. For example, in the case of Exponential service times ([5,

Sec. 3.5]),

$$
\begin{aligned}
& p_{k}=p_{0} \frac{(s p)^{k}}{k!}, k \leq m \\
& p_{0}=\left[\sum_{k=0}^{m-1} \frac{(s \rho)^{k}}{k!}+\frac{(s \rho)^{s}}{s!(1-\rho)}\right] .
\end{aligned}
$$

In the general distribution case one can map the distribution to a PH distribution (for example using [10]) and use matrix analytic methods ([7]) to obtain a good approximations of $p_{k}$.
For evaluating $\beta(y)$ and $\gamma(x)$ note that due to the Poisson arrivals the remaining service time on each of the servers at the epoch $s_{1}$ is the residual life of the service time, and has a pdf $\hat{b}(x)=(1-B(x)) / \bar{x}$ and a $\operatorname{cdf} \hat{B}(x)=\int_{0}^{x} f(t) d t$. For $C_{2}$ to be delayed exactly $y$ units of time due to one specific server, the other servers need to have a residual service time larger than $y$, so the probability of that event is $\hat{b}(y)(1-\hat{B}(y))^{s-2}$, and finally $\beta(y)=(s-1) \hat{b}(y)(1-\hat{B}(y))^{s-2}$. Using the same argument $\gamma(x)=(1-\hat{B}(y))^{s-1}$.

To summarize

$$
\begin{aligned}
& \mathcal{T}^{\text {SingleQueue }}(x)=\left(1-\sum_{k=0}^{s-2} p_{k}\right) \times \\
& \left(\int_{0}^{x}(s-1) \hat{b}(y)(1-\hat{B}(y))^{s-2} y d y+(1-\hat{B}(y))^{s-1} x\right)
\end{aligned}
$$

One can observe that the twin measure is decreasing with $s$, as both $\alpha$ and $\beta(y)$ increase with $s$.

### 5.2 Multiple Queues

In this section we make some observation about the multiple queue system, where each queue is assigned one server, and that server serves the jobs in that queue in FCFS manner. If a job joins the system and finds an empty queue it joins that queue and is served immediately. Other wise, the job is assigned a queue using some queue assignment policy. Once a job is assigned a queue, it cannot jockey to another queue, even if the other queue's server is idle. We denote this setting MultipleQueue.

Note that analysis is dependent on the queue assignment policy. However, for all queue assignment policies there are three possible cases: (i) Two or more queues are empty upon arrival of the twins. In this case both twins are served in parallel and leave the system simultaneously. (ii) Both twins join the same queue. In this case $C_{2}$ departs $x$ units after $C_{1}$, and (iii) the twins join different queues. In this case the twins will depart $|y|$ units of time apart, where $y$ is the difference in the remaining work in the two queues, which can be negative.
Using very similar notation to (5)

$$
\mathcal{T}^{\text {MultipeleQueue }}(x)=(1-\alpha)\left(\int_{0}^{x} \beta(y)|y| d y+\gamma(x) x\right)
$$

where $\alpha$ is the probability of arriving when two or more queues are empty. Given that no more than one queue was empty, $\beta(y)$ is the probability that the difference in the remaining work in the two queues $C_{1}$ and $C_{2}$ joined is $y . \gamma(x)$ is the probability that the twins will join the same queue.

Evaluating $\alpha, \beta(y)$ and $\gamma(x)$ is much more complicated in this case than it was for SingleQueue. We state some methods for this evaluation below. One can also resort to simulation methods, which are in fact not very complicated.

Evaluating $\alpha$, after the general distribution is mapped to a PH distribution, involves a Markovian-chain which is infinite in more than one dimension. In these cases matrix analytic methods do not work and one needs to use other methods, such as the Dimensionality Reduction method proposed in [9], and see discussion there of other methods. This provides us with full state probabilities.

Analyzing $\beta(y)$ requires knowledge of the distribution of remaining work in a queue. This is in general much more complicated than the residual life of a single job, although it is possible, using the queue length distribution, which is obtainable from the state probabilities.
$\gamma(x)$ might be simple or complicated, depending on the job assignment policy. For example if the jobs are assigned to queues in random, $\gamma(x)=1 / s$. If jobs are assigned to the shortest queue, the queue length distribution can be used.

One queue assignment policy of interest is to assign jobs to the queue with the least remaining work. Although this joining policy might not always be practical, it is possible in some computing systems where the length of each job is predetermined. Note that in this case $C_{1}$ always departs first, and $C_{2}$ always departs $x$ or less units of time after $C_{1}$. One can therefore easily observe that the twin measure is smaller than $x$. In general, this is not possible in other queue assignment policies.

## 6. EXTENDING THE TWIN MEASURE

In this section we discuss two ways to extend the twin measure: using more than two jobs, which we call trains, and not sending the jobs simultaneously.

### 6.1 Job Trains

One way to extend the results of the twin measure is to consider Job Trains, i.e. situations where more than two identical jobs are sent.

Packet Trains were proposed as means for measuring link bandwidth and available bandwidth (e.g. $[2,3,4]$ ). In these, packet trains are injected into the network. The dispersion of the probe packets at the receiver side is then used in different bandwidth estimation algorithms, using, for example, dispersion mean values or dispersion variance. However, this entire body of work assumes that packets are served using the FCFS policy, as indeed is the case with packet routers, at least for packets of the same flow. However, job schedulers may choose different policies, with quite different results.

We provide here only the simplified definition (parallel to Definition 2.2), as the non-simplified one is much more difficult to define rigorously and provides little benefit.

Definition 6.1 (Job Train Measure).
Let $C_{1}, C_{2}, \ldots, C_{m}$ be $m$ identical jobs, with equal service requirements $x$, arriving at epochs $a_{1}, a_{2}=a_{1}+\epsilon, a_{3}=$ $a_{1}+2 \epsilon, \ldots, a_{m}=a_{1}+(m-1) \epsilon$ where $\epsilon>0$, and departing at epochs $d_{1}, d_{2}, \ldots, d_{m}$. Assume that $C_{1}$ arrives when $\left(a_{1}\right)$ the system is in steady state. Define the random variable $Z(x, \epsilon, m)=\max _{i} d_{i}-\min _{i} d_{i}$, given $x$, $m$ and $\epsilon$. Let $z^{n}(x, \epsilon, m)$ be the $n$-th moment of $Z(x, \epsilon, m)$, i.e. $z^{n}(x, \epsilon, m)=\mathbb{E}\left\{Z(x, \epsilon, m)^{n}\right\}$.

For scheduling policy $\phi$ and job size $x$, the $n$-th job train measure $\mathcal{T}_{n}^{\phi}(x, m)$ is defined as the limit, when $\epsilon$ tends to zero, of $z^{n}(x, \epsilon, m)$, assuming a limit exists. Namely $\mathcal{T}_{n}^{\phi}(x, m)=\lim _{\epsilon \rightarrow 0} z^{n}(x, \epsilon, m)$.
The shortened term job train measure, denoted $\mathcal{T}^{\phi}(x, m)$, is
used to describe the first job train measure, namely $\mathcal{T}^{\phi}(x, m)=\mathcal{T}_{1}^{\phi}(x)=\lim _{\epsilon \rightarrow 0} \mathbb{E}\{Z(x, \epsilon, m)\}$.

Analysis of the job train measure for the policies analyzed in Section 3 is quite straightforward. In fact, for all the policies $\mathcal{T}^{\phi}(x, m)=(m-1) \mathcal{T}^{\phi}(x)$. However, this isn't always the case.
Consider for example a synchronous server setting. All jobs are of size $x$ and the server works in service cycles of length $2 x$ which serve either one or two customers, depending on availability of jobs. Each job receives a service rate of $1 / 2$. If $p_{\text {odd }}$ is the probability of finding an odd number of customers in the system upon arrival, then the job train measure is $p_{\text {odd }} 2 x\lceil(m-1) / 2\rceil+\left(1-p_{\text {odd }}\right) 2 x\lceil(m-2) / 2\rceil$ which is definitely not linear with $m$. For example, for $m=2$ we get $p_{\text {odd }} 2 x$, for $m=3$ we get $2 x$.
A second example is a round robin policy, where jobs join the queue at a random location. One can observe that all the jobs in the job train will be served in the same service cycle, so the job train measure is at most as large as the service cycle length. This will grow with $m$, but not linearity.

### 6.2 Non-Simultaneous Twins

A second way to extend the results of the twin measure is to consider twins which do not arrive at the system concurrently. For example, the twins can arrive exactly $x$ units apart, i.e. $a_{2}-a_{1}=x$. This might provide more insight on the predictability of the system. For example, note that only for LRPT the measure is still zero.
Another case of specific interest is the case where the second twin enters the system when the first one departs, i.e. $a_{2}=d_{1}$. For example this setting can represent a customer refreshing a Web page just as it finished loading, expecting a similar load time. It might be more interesting to measure the difference or ratio between the sojourn times, i.e. $\mid\left(d_{1}-\right.$ $\left.a_{1}\right)-\left(d_{2}-a_{2}\right) \mid$ or $\max \left(d_{1}-a_{1}, d_{2}-a_{2}\right) / \min \left(d_{1}-a_{1}, d_{2}-a_{2}\right)$.

## 7. CONCLUDING REMARKS

We proposed a novel way of measuring the predictability of a queueing system, utilizing pairs of identical customers. This measure was analyzed for several common scheduling policies. The results show that the measure in itself is unrelated to the utilization or waiting time performance of the system. The results also suggest that the measure in itself cannot guarantee service time predictability. Comparing these results with the results of the Conditional Response Time Criterion show that in several cases they contradict, leading to the proposal that to guarantee predictability maybe both of the measures should be used as criteria.
Analysis of the measure for multiple servers shows that for single queue the measure improves when the number of server increases. For multiple queues this is dependent on the queue assignment policy.

We proposed several ways in which the measure can be extended, namely by using more than two customers, or by considering jobs that do not arrive concurrently. Analysis of these is left for future research.

## 8. ACKNOWLEDGEMENTS

This work was supported in part by the Israeli Ministry of Science and Technology, grant number 380-801, and by EURO-NGI network of excellence

## 9. REFERENCES

[1] B. Avi-Itzhak, H. Levy, and D. Raz. Quantifying fairness in queueing systems: Principles, approaches and applicability. Probability in the Engineering and Informational Sciences (PEIS), 2006. To appear.
[2] R. L. Carter and M. E. Crovella. Measuring bottleneck link speed in packet-switched networks. Performance Evaluation, 27-28:297-318, October 1996.
[3] N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. IEEE Journal on Selected Areas in Communications, 21(6):879-894, August 2003.
[4] M. Jain and C. Dovrolis. End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput. IEEE/ACM Transactions on Networking, 11(4):537-549, August 2003.
[5] L. Kleinrock. Queueing Systems, Volume 1: Theory. Wiley, 1975.
[6] L. Kleinrock. Queueing Systems, Volume 2: Computer Applications. Wiley, 1976.
[7] G. Latouche and V. Ramaswami. Introduction to Matrix Analytic Methods in Stochastic Modeling. ASA-SIAM, Philadelphia, 1999.
[8] M. Nuyens and A. Wierman. The foreground-background queue: a survey. Under submission, 2007.
[9] T. Osogami. Analysis of Multiserver Systems via Dimensionality Reduction of Markov Chains. PhD thesis, School of Computer Science, Carnegie Mellon University, 2005.
[10] T. Osogami and M. Harchol-Balter. Closed-form solutions for mapping general distributions to quasi-minimal PH distributions. Performance Evaluation, 63(6):524-552, June 2006.
[11] L. E. Schrage and L. W. Miller. The queue M/G/1 with the shortest processing remaining time discipline. Operations Research, 14:670-684, 1966.
[12] A. Wierman. Scheduling for Today's Computer Systems. PhD thesis, School of Computer Science, Carnegie Mellon University, 2007.
[13] A. Wierman and M. Harchol-Balter. Classifying scheduling policies with respect to higher moments of conditional response time. In Proceedings of ACM Sigmetrics 2005 Conference on Measurement and Modeling of Computer Systems, pages 229-239, Banff, Alberta, Canada, June 2005.


[^0]:    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
    Valuetools '07, October 23-25, 2007, Nantes, France
    Copyright 2007 ICST 978-963-9799-00-4.

