Steady State Distribution for Stochastic Knapsack with Bursty Arrivals

Venkatesh Sarangan, Donna Ghosh, Natarajan Gautam, and Raj Acharya

***Abstract*— In this letter, we develop a methodology for obtaining an approximate steady state occupancy distribution for a multi class stochastic knapsack with bursty call arrivals, and exponential holding times. Preliminary results indicate that the proposed technique is effective in studying the knapsack behavior.**

***Index Terms*— Performance analysis.**

I. Introduction

THE STOCHASTIC knapsack with Poisson call arrivals is a well known model for analyzing telecommunication systems [5]. With the evolution of the Internet and the World Wide Web, researchers have shown that the actual call arrival processes in the network are bursty in nature with long range dependence. Hence more generic knapsack solutions for non-Poisson call arrivals should be developed in order to study the current networks. To the best of our knowledge, solving for the Knapsack distribution for any non-Poisson arrival has remained an open problem since 1981.

Stochastic processes that are *bursty* have *long range dependence*, and are characterized by *slowly decaying autocorrelation* function. We can also characterize such processes through *heavy tailed distributions*, although the class of processes so defined is more inclusive than the preceding characterization [7]. Especially in the context of network traffic, it is possible to generate an arrival process that is asymptotically bursty, either at the packet or connection level, by drawing the inter-arrival times that are independent and identical from a heavy tailed distribution [7], [2]. Following this, we characterize a bursty arrival process in this work as the one whose inter-arrival times are independent and identically distributed (i.i.d.) as per a heavy-tailed distribution.

II. Stochastic Knapsack with Bursty Arrivals

The stochastic knapsack consists of C resource units (e.g., bandwidth) to which objects of K mutually independent classes arrive. Arrivals of class $k, k = 1, \cdots K$, are governed by a *bursty* stochastic process with mean rate Λ_k. If an arriving $class - k$ object is admitted into the knapsack, it holds b_k resource units for an holding time drawn from an *exponential* distribution with mean $1/\mu_k$. We assume that the knapsack follows the *complete sharing* policy for admitting the class arrivals [5]. Let $\mathcal{X}_k(t)$ be the stochastic process denoting the number of $class - k$ objects in the knapsack at time t. Let $\vec{\mathcal{X}}(t) = \{\mathcal{X}_1(t), \cdots \mathcal{X}_K(t)\}$ be the stochastic process representing the state of the knapsack at time t. The objective of this letter is to find the steady state distribution of the process $\vec{\mathcal{X}}(t)$.

A. Outline of the proposed solution

We first consider a knapsack with infinite capacity. Let $\{Y_k(t)\}$, $\forall k$ be the stochastic process governing the number of $class - k$ objects in the infinite capacity knapsack, defined over the state space $\mathcal{S}' = \{\vec{n} \in \mathcal{I}^K : \vec{b}.\vec{n} \leq \infty\}$; where $\vec{b} = (b_1, b_2, \cdots b_K)$ with b_k denoting the bandwidth requested by an arrival from class k, and $\vec{n} = (n_1, n_2, \cdots n_K)$ with n_k representing the number of class k objects present in the knapsack. As the knapsack capacity is infinite, the K processes $\{Y_1(t)\}, \cdots \{Y_K(t)\}$ are all mutually independent. For each $\{Y_k(t)\}$, we construct an equivalent process $\{Y'_k(t)\}$ which is *Markovian* and *reversible*, and defined over $\mathcal{S}'$. Let $\{\vec{Y}(t)\}$ and $\{\vec{Y}'(t)\}$ be two joint processes defined as $\vec{Y}(t) = (Y_1(t), \cdots Y_K(t))$, and $\vec{Y}'(t) = (Y'_1(t), \cdots Y'_K(t))$ respectively. Since $\{Y_k(t)\}$ and $\{Y'_k(t)\}$ are equivalent for all k, $\{\vec{Y}(t)\}$ and $\{\vec{Y}'(t)\}$ are also equivalent. The original finite capacity knapsack process $\{\vec{\mathcal{X}}(t)\}$ is nothing but $\{\vec{Y}(t)\}$ defined over the finite state space $\mathcal{S} = \{\vec{n} \in \mathcal{I}^K : \vec{b}.\vec{n} \leq C\}$. Since $\{\vec{Y}(t)\}$ and $\{\vec{Y}'(t)\}$ are equivalent, we approximate the long run behavior of $\{\vec{\mathcal{X}}(t)\}$ with that of $\{\vec{Y}'(t)\}$ over the finite state space $\mathcal{S}$.

The processes $\{Y'_k(t)\}$, $\forall k$ are all $G/M/\infty$ queuing processes whose steady state distribution can be obtained using the procedure outlined in [8]. As the processes $\{Y'_k(t)\}$, $\forall k$ are all mutually independent, the steady state distribution of the joint process $\{\vec{Y}'(t)\}$ is the product of the individual $G/M/\infty$ steady state solutions. The joint process $\{\vec{Y}'(t)\}$ is also Markovian and reversible, since the constituent processes $\{Y'_k(t)\}, \forall k$ are reversible, Markovian, and mutually independent. By using Kelly's theorem on solving for a truncated, reversible, Markovian process [4], we obtain the steady state distribution of $\{\vec{Y}'(t)\}$ over $\mathcal{S}$ from its distribution over $\mathcal{S}'$, and hence the finite capacity knapsack solution.

B. Equivalent per class reversible queuing process

For each of the K queuing processes $\{Y_k(t)\}, k = 1, \cdots K$, we construct a process $\{Y'_k(t)\}$, which is equivalent to $\{Y_k(t)\}$ in the following sense:

1) Let $\zeta_k(t)$ and $\zeta'_k(t)$ be the arrival processes corresponding to $\{Y_k(t)\}$ and $\{Y'_k(t)\}$ respectively. Let $A_k(t)$

Manuscript received May 9, 2004. The associate editor coordinating the review of this letter and approving it for publication was Prof. Changcheng Huang. This work was supported in part by the NSF ITR grant 0219747.

V. Sarangan is with the Computer Science Department at Oklahoma State University (e-mail: saranga@cs.okstate.edu).

D. Ghosh and R. Acharya are with the Department of Computer Science and Engineering at Pennsylvania State University e-mail: {dghosh, acharya}@cse.psu.edu).

N. Gautam is with the Department of Industrial Engineering at Pennsylvania State University (e-mail: ngautam@psu.edu). The first two authors contributed equally to this work.

Digital Object Identifier 10.1109/LCOMM.2005.02026.

and $A'_k(t)$ be the respective inter-arrival distributions of $\zeta_k(t)$ and $\zeta'_k(t)$. $A_k(t)$ and $A'_k(t)$ are such that they have the same mean and variance.
2) The service time distributions of $\{Y_k(t)\}$ and $\{Y'_k(t)\}$ are exactly the same.

In the arrival process $\zeta'_k(t)$, calls are drawn from an infinite population into an arrival holding station, where they are held for some random time. They are then sent to an infinite server system. The arrival holding station for class k is shown in Fig. 1(a). It is a series-parallel combination of two mutually independent holding boxes, named stage 1 and stage 2. At any point in time, only one arrival is held in the station. The parameters of class k's holding station are:

- σ_k: Probability of an arrival being drawn into stage 1.
- $1/\lambda_i^{(k)}$: Mean holding time in stage i ($i = 1, 2$) of the station; exponentially distributed.
- α_k (β_k): Probability that an arrival which is being held in the station, toggles from stage 1 to 2 (1 to 2), when a customer departs from the $G/M/\infty$ system.

The above parameters of class-k arrival holding station are chosen such that: (a) the mean and variance of $A'_k(t)$ are matched with those of $A_k(t)$, and (b) the class-k $G/M/\infty$ queuing process ($\{Y'_k(t)\}$) is *reversible*.

C. Mean and Variance of $A'_k(t)$

Let T_r^k be the time between the constructed $class - k$ arrivals r_k and $r_k + 1$, where $r_k = 1, 2, \cdots$. Let $\Psi_k(s) = E[e^{-sT_r^k}]$. The expression for $\Psi_k(s)$ can then be written as:

$$\Psi_k(s) = \sum_{n_k=1}^{\infty} q_{n_k-1} \left[\sigma_k F_{X_{n_k,1}}(s) + (1-\sigma_k) F_{X_{n_k,2}}(s)\right] \quad (1)$$

where q_{n_k} is the arrival time probability of seeing 'n_k' $class - k$ customers of the actual arrival process $\zeta_k(t)$ in the infinite server system, and can be obtained using the approach described in [8]. $F_{X_{n_k,i}}(s)$ is given by $E[e^{-sX_{n_k,i}}]$, $i = 1, 2$, where $X_{n_k,i}$ is the random variable denoting the time spent in the arrival holding station by a $class-k$ call drawn in to stage $i = 1, 2$. The details of calculating $F_{X_{n_k,i}}(s)$ can be found in [6]. Once $F_{X_{n_k,i}}(s)$ is known, the mean η_k and variance Ω_k of the inter-arrival times can be obtained from equation 1 as: $\eta_k = \Psi'_k(0)$ and $\Omega_k = \Psi''_k(0) - \eta_k^2$.

D. Reversibility of $\{Y'_k(t)\}$

Theorem 1: If the parameters of the arrival holding station are such that

$$\alpha_k + \beta_k = 1, \text{ and, } \lambda_1^{(k)} \sigma_k \alpha_k = \lambda_2^{(k)} (1-\sigma_k) \beta_k \quad (2)$$

then, the Markov process $\{Y'_k(t)\}$ is reversible.
The proof of this theorem can be found in [6].

E. Steady State Occupancy Distribution of $\{\vec{Y}'(t)\}$ over $\mathcal{S}$

Once $\{Y'_k(t)\}$ is constructed, we can obtain its steady state distribution over $\mathcal{S}'$ using the technique suggested in [8]. From the solution of K such processes, $\{Y'_k(t)\}, k = 1, \cdots K$, the distribution of the joint process $\{\vec{Y}'(t)\}$ over $\mathcal{S}'$ can be obtained. Then as outlined in section II-A, we can obtain the steady state solution for $\{\vec{Y}'(t)\}$ over $\mathcal{S}$. We note that, although the per-class queuing process shown in Fig. 1(b) is reversible when equation 2 is satisfied, there will be additional transitions between the boundary states of the knapsack with a finite capacity C, which may make the process $\{\vec{Y}'(t)\}$ irreversible. We neglect the effect of these transitions in our approach, and hence obtain only an approximate steady state occupancy distribution of $\{\vec{Y}'(t)\}$.

We are currently working on obtaining bounds to quantify the 'closeness' of the process $\{\vec{Y}'(t)\}$ to the original knapsack process $\{\vec{\mathcal{X}}(t)\}$. In order to obtain a bound, one has to acquire some handle on the process $\{\vec{\mathcal{X}}(t)\}$. It is precisely because of this lack of an elegant handle on $\{\vec{\mathcal{X}}(t)\}$, we resort to an approximation in the first place. This dependency makes the goal of obtaining a bound challenging.

III. Simulation Set-up and Results

A. Methods Compared and Performance measures used

We compare the following approaches of analyzing a knapsack with bursty call arrivals with the actual knapsack behavior obtained through simulations.

1) **Poisson Approach:** Given a bursty arrival process, we create an equivalent Poisson process with the same mean. Then we use the approach outlined in [3] to obtain the steady state occupancy distribution.
2) **Heuristic Approach:** In this method, given a bursty arrival process, we assume that the corresponding original queuing process is *reversible*. Then we use the approach outlined in [3] to obtain the solution. We study this approach to clearly understand the cost and benefits of approximating the original queuing process by an equivalent, *reversible* process.
3) **Proposed Approach:** Approach outlined in section II.

It is to be noted that the Poisson approach is a first order approximation, whereas the proposed approach is a second order approximation.

We use the following as the performance measures: (a) the L_1 norm between the actual knapsack distribution and the distribution obtained through the above approaches, and (b) the percentage errors in average number of calls *(AvgCall)* and average knapsack utilization *(AvgUtil)*. We simulated the actual multi-class stochastic knapsack system, in OPNET [1], a network simulation tool. For each class, we generate the bursty call arrival process by drawing the call inter-arrival times from an heavy tailed Weibull distribution. The results reported here are from a knapsack with capacity $C = 20$ units, and three classes with $b_1 = 1$, $b_2 = 2$, and $b_3 = 4$.

B. Results

We studied the behavior of the three schemes under varying loads and burstiness. Due to lack of space, we present only the results we obtained under call arrivals with high burstiness, and heavy knapsack loads. The performance of the three schemes under bursty arrivals (when the call inter-arrival times are heavy tailed) and heavy loads are shown in Figs. 2, and 3 respectively. The plots in these graphs are made with respect to varying heavy tailedness and varying loads of class 3 arrivals. We find that the gains achieved by our proposed scheme is significant, either when the traffic is bursty or when the

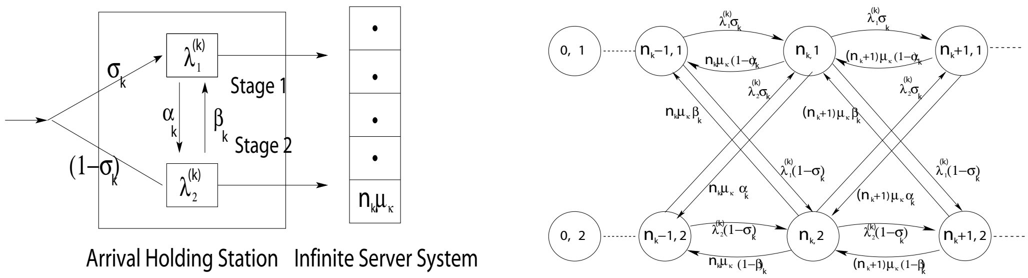

(a) Schematic representation of the constructed system. (b) State transitions of the constructed process.

Fig. 1. Equivalent class k queuing process with infinite servers.

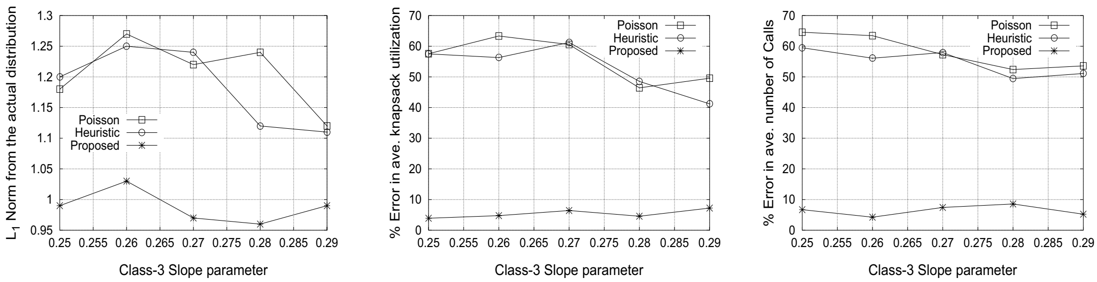

(a) Accuracy of the distributions. (b) % Error in ave. knapsack utilization. (c) % Error in ave. number of calls.

Fig. 2. Performance of the three schemes under high burstiness. Plotted against the slope parameter of class 3's inter-arrival distribution.

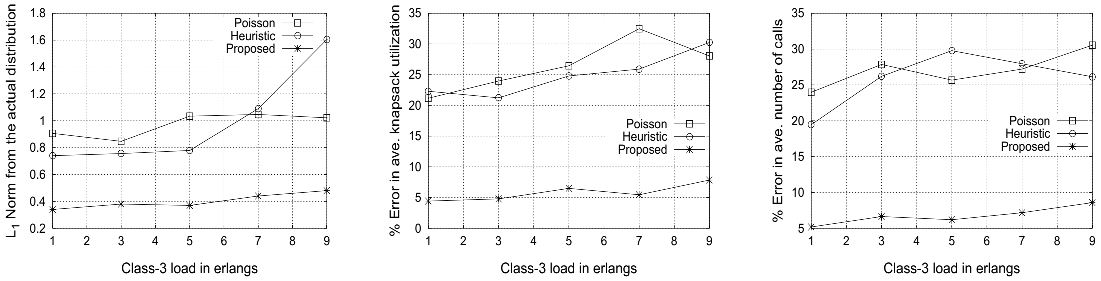

(a) Accuracy of the distributions. (b) % Error in ave. knapsack utilization. (c) % Error in ave. number of calls.

Fig. 3. Performance of the three schemes under heavy loads. Plotted against the load offered by class 3 arrivals.

knapsack is heavily loaded. Under other scenarios, the Poisson approximation gives an adequate performance.

Experimentally, we were not able to discover any monotonocity in any of the performance measures with respect to burstiness or system load. Since the heuristic approach does not perform as well as the proposed approach, we conclude that the improvement obtained by the proposed approach (in comparison to the Poisson approximation method), is not merely due to matching an extra moment while constructing the equivalent arrival process. We have also observed that between the Poisson and the heuristic approaches, there is no consistency in behavior.

IV. Conclusions

In this paper, we have proposed a method that analytically solves for the approximate steady state occupancy distribution of a stochastic knapsack with bursty call arrival processes, and exponential call holding times. Based on our studies, we conclude that whenever the arrivals are bursty or when the knapsack is heavily loaded, the proposed method gives significantly more accurate results than the Poisson approximation. We are currently investigating methods to obtain bounds on the performance of the proposed scheme for various inter-arrival distributions. We are also working on extending the proposed approach to analyze knapsacks with call holding times other than exponential.

References

[1] [1] http://www.opnet.com
[2] [2] A. Feldmann, *Characteristics of TCP Connection Arrivals, Self Similar Network Traffic and Performance Evaluation*, K. Park and W. Willinger, eds. New York: John Wiley and Sons Inc., pp. 367-400, 2000.
[3] [3] J. S. Kaufman, "Blocking in a shared resource environment," *IEEE Trans. Commun.*, vol. 29, pp. 1471-1481, Oct. 1981.
[4] [4] F. P. Kelly, *Reversibility and Stochastic Networks*. Chichester, UK: Wiley, 1979.
[5] [5] K. W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*. New York: Springer-Verlag, 1995.
[6] [6] V. Sarangan, "Multi-class Internet: models for routing and analysis," Ph.D. dissertation, Dept. of CSE, Penn State University, University Park, PA, Aug. 2003.
[7] [7] W. Stallings, *High Speed Networks: TCP/IP and ATM Design Principles*. Englewood Cliffs, New Jersey: Prentice Hall Inc., 1998.
[8] [8] L. Takacs, *Introduction to the Theory of Queues*. New York: Oxford University Press, 1962.