Network Protocol Performance Bounding Exploiting Properties of Infinite Dimensional Linear Equations*

Ioannis Stavrakakis

Department of Informatics and Telecommunications,
National and Kapodistrian University of Athens,
Ilissia, 157 84 Athens, Greece
ioannis@di.uoa.gr

Abstract. This paper presents a quite versatile and widely applicable performance analysis methodology that has been applied for the study of network resource allocation protocols in the past. It is based on the identification of renewal cycles of the operation of the system and the setting up of recursive equations with respect to quantities-indices defined over the renewal cycles and sessions that appear within. Application of the expectation operator on these equations leads to infinite dimensional systems of linear equations which are shown to posses certain properties leading to rigorous and almost arbitrarily tight bounds on various performance metrics of interest. The special case of a random access protocol is used as an example in order to illustrate the derivation of the recursive equations capturing the protocol dynamics and system inputs. Finally, some other examples of application of the methodology are briefly discussed, illustrating the versatility and powerfulness of the approach. This analysis methodology can be quite useful for understanding the behavior of current complex and large scale networking environments, as well as assessing their scalability, stability and performance.

1 Introduction

The field of computer / communication networking was basically born in the sixties following the inception of packet switching and the development of the first networking infrastructure to enable the transport of information between a few research sites in USA. Since then, the field has boomed and has changed drastically the way society functions in almost all aspects.

As it is the case with any scientific and engineering development, the efficiency and effectiveness of the designed networks has been central to their success. Networking is by definition about the *efficient* interconnection of (geographically) distributed entities, as simply providing (dedicated) links between any two entities would not be feasible or acceptable. The greatest challenge in

* This work has been supported in part by the IST-FET project RECOGNITION (FP7-IST-257756).

K.A. Hummel et al. (Eds.): PERFORM 2010 (Haring Festschrift), LNCS 6821, pp. 129–140. 2011

designing an efficient networking system is to devise efficient ways for utilising the available resources. These resources could vary from a single common resource (channel) shared by a number of distributed entities, to the distributed storing/processing/forwarding resources made available by independent mobile nodes.

The efficient (computer) resource utilization problem is one that received the attention of computer system designers before the beginning of data networking. For this reason, approaches and methodologies applied to the design and analysis of computer systems were widely imported into the networking arena. Nevertheless, the distributed nature of the available resources in a networking environment called for the development of new methodologies for designing and analyzing networking structures. Considering the fact that a network is a distributed system whose efficient design requires the incorporation of numerous disciplines and tools, networking was widely viewed as inter-disciplinary and hard to be considered as a separate discipline on its own. It is only very recently that the question as to whether a networking discipline or science can be defined, was posed, [1]. Although still borrowing from an ever increasing number of diverse disciplines (such as anthropology and biology, more recently), the twist of the problems networking is defining and approaches to addressing them are becoming quite distinct from the associated underlying disciplines. In addition, the traditional networking paradigm of connecting sites and transporting messages is rapidly changing into one of content accommodation and delivery in a networking environment collectively contributed through individually-owned disperse networking resources. The key resource is also changing from the transmission resource into a pool of enabling resources such as storage, power, mobility, processing, social, etc, [2].

Returning to the performance evaluation methodologies developed to help design and analyze the early packet networks of the sixties, one can clearly observe a focus on stochastic modelling, queueing theory and operations research, supplemented with mathematical theory and tools. This trend pretty match continued in the seventies and through the nineties when designing multi-user (random) resource access schemes (ALOHA, Ethernet type of networks), integrated networks (ATM) and the Internet. These studies were increasingly supplemented by powerful computer simulations, capturing system details that the analysis often neglects; nevertheless, the insight and guidance through an often immense design space that the analytical modelling provides is still invaluable. The paradigm shift that the Internet has brought towards more autonomic networking clouds, which has been further intensified with the proliferation of autonomic networks (P2P) and the shift of focus towards content generation, accommodation and provisioning, has brought into the performance evaluation arena new methodologies coming from disciplines such as biology, anthropology, social sciences, distributed algorithms, complexity, etc.

The focus of this paper is to present a modelling and performance evaluation approach that yields a tractable analysis methodology to a broad class of networking problems. It is based on a fairly versatile stochastic modelling methodology that leads to the formulation of a typically infinite-dimensional system of linear equations whose proper manipulation yields (in a computationally tractable way) tight lower and upper bounds on performance indices of interest. This approach is presented in the context of the analysis of a specific random access multi-user network that was proposed in the late eighties, since the particular example illustrates the applicability of the methodology to determining - among other performance indices - the stability region of the system, as well.

As already stated, the fundamental problem in networking is that of allocating resources to distributed entities. Nowadays, this competition for resources manifests itself in more complex and generalized manner: bandwidth is only one of a plethora of resources supporting networking; the scale of the resources and competition is global; resources and services are proliferating and are largely autonomic; the networking paradigm, priorities and commodities are changing, etc. For such environment - characterized largely by the generation of unprecedented amounts of content and services - the questions of stability, scalability and performance are central to their efficient design. Although new designs and approaches may be necessary to establish stability and scalability (e.g., by focusing on reducing the irrelevant and useless load (demand for resources) through information processing and effective decision making, [16]), the analysis methodology presented here can be quite useful for understanding the "laws and physics" underlying complex and large scale networking environments today, as well as assessing their scalability, stability and performance.

The paper is organized as follows. Section 2 presents the Limited Sensing Random Access Protocol (LS-RAP) that is used here to illustrate the methodology. In section 3 the stochastic modelling and associated recursive equations capturing the protocol and system dynamics are presented. Section 4 presents the infinite dimensional linear equations and describes how their structure and properties can be exploited by invoking relevant theory to derive bounds on the maximum stable throughput and mean delay. Section 5 presents a discussion on some ways to improve on the bounds, as well as on the applicability of the methodology for studying other protocols and queueing systems in the past. Finally, the contribution of this paper is summarized in section 6.

2 The Limited Sensing Random Access Protocol (LS-RAP)

We consider the classical problem of sharing a common communication channel by a large number of distributed users, as it is the case, for instance, in the classical Ethernet. The environment is assumed to be discrete-time and synchronized in the sense that the users are aware of the beginning of a fixed-size time slot whose length equals the packet transmission time. It is assumed that the unit of information is the message that consists of M packets and that the cumulative message arrival process is Poisson with intensity λ messages per message length. As it is widely understood, a random channel access protocol is appropriate for such a networking environment. As a result, collisions arise when more than one packets are transmitted in overlapping slots. A collision resolution algorithm that belongs to the class of the stack algorithms [5,6], and originally proposed in [3,4] is outlined here and will be denoted by LS-RAP (limited sensing random access protocol).

An active user (i.e., one that has a message to transmit) keeps sensing the channel starting from the slot that follows its message generation instant until this message is successfully transmitted; that is, earlier channel history is not needed (limited sensing). At the end of each slot ternary feedback information (F) is available to all active users revealing the channel status, i.e., whether the channel was idle (I), involved in a successful transmission (S), or in a collision (C). When a packet collision occurs, the sender aborts its transmission before the end of the current slot; thus, only one slot is wasted in a collision. A message transmission is successful once its first packet has been successfully transmitted. The users are geographically separated and their status (i.e., existence or nonexistence of a message to be transmitted) cannot be communicated to the rest of the users.

The concept of the stack is used to illustrate the operation of the algorithm and is not needed for its implementation. The content of a counter (assigned to each user) determines the cell of the stack which the particular user belongs in; it also determines the class of users in which it belongs. Let B_k denote the class of users whose counter content equals k, $k \geq 0$ and let B denote the class of inactive users. B_0 denotes the class of new active users, i.e., those users who have a message to transmit but no message transmission has been attempted so far. B_1 is the class of active users who attempt a packet transmission at the beginning of the slot that follows. Each user's counter content is updated at the end of the slots based on the channel feedback, the steps of the algorithm (see below) and the counter content itself; the new value of the counter determines the class which the user enters. Each user is also assigned a downcounter ω whose initial value is M and which decreases by one unit per slot, starting from the slot in which the first packet of the message was successfully transmitted; this counter determines the time when a successfully transmitting user completes the transmission of all M packets of the message and becomes inactive. LS-RAP can be described by the following updates of the user classes; see [3] for a detailed discussion on the implications of these updates.

1. If $F = S$ *then:* a) $B_0 \rightarrow B_0$ b) $B_1 \rightarrow B_1$, if $\omega > 0$; $B_1 \rightarrow B$, if $\omega = 0$ c) $B_k \rightarrow B_k$, $k \geq 2$.
2. If $F = C$ *then:* a) $B_0 \rightarrow B_1$ b) $B_k \rightarrow B_{k+1}$, $k \geq 2$ c) For each $b \in B_1$, $b \rightarrow B_2$ with probability $1 - p$; $b \rightarrow B_1$ with probability p.
3. If $F = I$ *then:* a) $B_0 \rightarrow B_1$ b) If last nonidle slot was involved in a collision, then for each $b \in B_2$, $b \rightarrow B_1$ with probability $1 - p$; $b \rightarrow B_2$ with probability p; $B_k \rightarrow B_k$, $k \geq 3$ c) If last nonidle slot was involved in a successful transmission, then c_i) if the current slot is the first idle slot after the successful one, then $B_k \rightarrow B_k$, $k \geq 2$ c_{ii}) if the current slot is not as in c_i), then $B_k \rightarrow B_{k-1}$, $k \geq 2$.

3 Renewal Cycles, Sessions and Associated Recursive Equations Capturing Protocol and System Dynamics

Key to the analysis of LS-RAP is the identification of renewal cycles and the computation of their (mean) lengths. This is done by deriving recursive equations with respect to the cycle length and auxiliary sessions (that are stochastically identifiable and contained within a cycle), capturing the protocol dynamics and the system characteristics and inputs. Although these recursive equations are typically infinite dimensional, their structure presents sparse couplings which make their description and (approximate) computation less difficult. Furthermore, one can also define other quantities associated with the renewal cycle and sessions and compute them through a largely similar set of recursive equations. Such equations are presented for the case of the cumulative delay experienced by all packets that were generated and were transmitted during a renewal cycle and specific sessions.

The aforementioned general approach is applied here to the case of LS-RAP. A technical definition of the renewal cycle can be given via the use of an imaginary marker. The marker is originally placed in cell 0 of a conceptual stack. Upon collision, the marker is placed in cell 2. The position of the marker changes in the same way in which the counter content of the users of class B_3 changes, depending on the channel feedback. The slot in which the marker returns to cell 0 is the first one of the renewal cycle that follows and it is always *idle*. When the marker is in cell 0 and a successful transmission occurs the marker is placed in cell 1 and moves up or down as described before. Idle slots do not move the marker from cell 0 and they result in renewal cycles of length one.

From the technical definition of the renewal cycle above it is clear that indeed the system stochastically regenerates itself after the first slot of such a cycle. The number of users k, $k \geq 0$, who attempt a packet transmission in the slot which follows the first idle one of a renewal cycle, determines the multiplicity of the session that follows. The length of a session of multiplicity k is defined as the time required until the conflict of multiplicity k is resolved. All users who attempt packet transmission in that slot plus all those entering the system before the end of this session, transmit successfully during that session. Notice that the sessions of multiplicity k are contained within renewal cycles and their stochastic behaviour is completely determined by the multiplicity k and independent from the position within the renewal cycle they appear.

Let l_k denote the length of a session of multiplicity k (in time slots), $k \geq 0$. The following recursive equations can be written with respect to l_k, $k \geq 0$:

$$l_0 = 1, \quad l_1 = 1 + M + l_{F_M} \tag{1a}$$

$$l_k = 1 + \{1 + l_{k,0}\} I_{\{I_1+F_1=0\}} + \{l_{I_1+F_1} + l_{k-I_1+F_3}\} I_{\{I_1+F_1 \neq 0\}} \;, k \geq 2 \tag{1b}$$

$$l_{k,0} = \{1 + l_{k,0}\} I_{\{I_2+F_2=0\}} + \{l_{I_1+F_1} + l_{k-I_1+F_4}\} I_{\{I_2+F_2 \neq 0\}} \;, k \geq 2 \tag{1c}$$

where $I_{\{.\}}$ is the indicator function and F_M, F_1, F_2, F_3, F_4, I_1, I_2 are independent random variables; F_1, F_2, F_3, and F_4 are Poisson distributed over one slot, F_M is Poisson over $M + 1$ and I_1, I_2 are binomial with parameter k and p.

The equations in (1) can be explained as follows. a) The session of multiplicity 0 consists only of the idle slot which marks the beginning of this session. b) The session of multiplicity 1 consists of the following parts. i) The idle slot which is always the first of the session. ii) The M slots involved in the successful transmission of the single packet attempted. iii) The length l_{F_M} which is the same as the length of a session of multiplicity F_M. c) For $k \geq 2$, the session consists of the following. i) The idle slot which is always the first of the session. ii) Since collision occurs we have to distinguish between two cases. Let I_1 be the number of users which remain in class B_1 after the splitting and let F_1 be the number of new messages which arrive in the slot before the collided one. ii_a) If $I_1+F_1 > 0$ we add $l_{I_1+F_1} + l_{k-I_1+F_4}$ to the length of the session since the original session is split into two with the corresponding multiplicities. ii_b) If $I_1 + F_1 = 0$, we add another slot to the session since no transmission takes place, plus $l_{k,0}$. The latter quantity is equal to the length of a session of multiplicity k without including the slot of the original collision, i.e., $l_{k,0} = l_k - 1$.

Equations (1) are a set of versatile recursive equations with respect to the length of sessions of various multiplicities that appear within a renewal cycle, which are needed to determine the length of a renewal cycle. Equations similar to (1) can be derived with respect to other quantities of interest, whose solution yield performance metrics of interest such as the cumulative delay experienced by all messages generated and transmitted during sessions of various multiplicities and ultimately (through averaging with respect to the session multiplicities) over a renewal cycle. In general, one can define any proper counting metric associated with the sessions (for instance, the number of messages experiencing a delay less than n) whose mean value would ultimately yield a performance metric of importance (message delay distribution). The recursive equations for the cumulative delay experienced by all messages generated and transmitted during sessions of various multiplicities and ultimately over a renewal cycle under LS-RAP may be found in [3,4].

4 Performance Evaluation by Exploiting the Structure and General Properties of Infinite Dimensional Linear Equations

The main contribution of the first part of this paper is the presentation of a methodology for deriving bounds on key performance indices (maximum throughput and delays) by exploiting results from infinite dimensional linear equations. These equations are derived by applying the expectation operator on the recursive equations on key random variables, presented earlier in section 3.

By applying the expectation operator on (1), we obtain the following infinite dimensional linear equations with respect to the expected length of a session of multiplicity k, L_k, where the coefficients h_k and α_{kj} depend on λ, p and M and can be found in [3,4].

$$L_k = h_k + \sum_{j=0}^{\infty} \alpha_{kj} L_j \ , k \geq 0 \ . \tag{2}$$

Similar equations can be obtained with respect to C_k (mean cumulative delay of messages over a session of multiplicity k), with *identical* coefficients α_{kj}.

4.1 Bounds on the Maximum Stable Throughput

Bounds on the mean session length and the mean cumulative message delay over a session of multiplicity k, $k \geq 0$, are needed in order to derive bounds on the maximum achievable throughput and mean message delay. For this reason we proceed with the derivation of such bounds by presenting the foundations and mechanisms for deriving them from the infinite dimensional linear equations as shown in (2). The proofs may be found in [3]. *The maximum stable throughput* S_{max} is defined as the maximum over all input traffic rates λ that induce $L_k < \infty$ for $k < \infty$. The following proposition provides for a lower bound on S_{max} (see [3] and [5] for its proof).

Proposition 1. ***Lower bound on*** S_{max}***:*** *If* $\{x_k^u\}_{k=0}^{\infty}$ *is an infinite sequence of real numbers which satisfy the following conditions: (1)* $0 \leq x_k^u < \infty$, $0 \leq k < \infty$*, (2)* $h_k + \sum_{j=0}^{\infty} \alpha_{kj} x_j^u \leq x_k^u$, $0 \leq k < \infty$*, and (3)* $h_k \geq 0$, $\alpha_{kj} \geq 0$, *for* $k \geq 0$, $j \geq 0$*, then:*
(a) the infinite dimensionality linear system of equations $h_k + \sum_{j=0}^{\infty} \alpha_{kj} x_j = x_k$ $, 0 \leq k < \infty$ *has a unique nonnegative solution* $\{x_k\}_{k=0}^{\infty}$ *that satisfies* $0 \leq x_k \leq x_k^u$, $0 \leq k < \infty$ *and*
(b) If $x_k \equiv L_k$ *satisfies the above conditions for some input traffic* λ^l*, then* λ^l *is a lower bound on* S_{max}*.*

For $\lambda < \lambda^l$ and for some value of the splitting probability p, $0 \leq p \leq 1$, a quantity $L_k^u = \beta(\lambda, p)k - \gamma(\lambda, p)$ is derived analytically in [3] which satisfies the conditions of Proposition 1 (i.e., $L_k \leq L_k^u$ for $\beta(\lambda, p)$, $\gamma(\lambda, p)$ and k finite).

Proposition 2. ***Upper bound on*** S_{max}***:*** *An upper bound on* S_{max}*,* λ_N^u*, can be obtained as the maximum over all Poisson rates for which the truncated up to* N *system in (2) has a unique nonnegative solution that satisfies the condition* $\lim_{M\to\infty} \max_{N>M}\{\sum_{j=N}^{\infty} \alpha_{kj} L_j\} = 0$*. Then,* λ_N^u *decreases monotonically as* N *increases and* $\lim_{N\to\infty} \lambda_N^u = S_{max}$*, [6].*

4.2 Bounds on the Message Delay

As discussed in [3], the total message delay consists of two components: the mean access delay and the mean in system delay. The former reflects the mean time spent by a message from the time it is generated until it makes the first transmission attempt (entering class B_1) and can be calculated easily based on conditional probabilities on the channel state and the input traffic rates (see [3]). The in system delay reflects the remaining time spent in the system till successfully transmitted and is more challenging to compute. We derive bounds on the mean in system message delay, D_s, for input traffic rates $\lambda < \lambda^l$ for which it is guaranteed that the delays will be bounded.

Since the operation of the system follows a renewal process and the multiplicities of the sessions are independent and identically distributed random variables, the following theorem is a direct application of the strong law of large numbers (p. 126, [9]) and is easily proved, [3].

Lemma 1. *For $\lambda < \lambda^l$, the mean in system delay D_s is given by $D_s = \frac{C}{\lambda L}$ with probability* 1 *(wp 1), where $L = E\{L_k\}$, $C = E\{C_k\}$ where $E\{\}$ is over all multiplicities k which are Poisson distributed over one slot.*

Proposition 3. ***Bounds on the mean in system delay***: *Upper and lower bounds on D_s can be obtained by deriving upper and lower bounds on L and C and invoking Lemma 1.*

4.2.1 Upper Bounds on L_k and C_k:

A linear with respect to k upper bound on L_k, L_k^u, was presented in section 4.1. The values of $L^u \equiv E\{L_k^u\}$ were found to be very close to those of L^l for $\lambda < 0.8\lambda^l$. As λ approaches λ^l, the upper bound increases rapidly and becomes loose; some approaches to calculating a tighter upper bound are discussed in section 5.

To derive an upper bound on C_k, C_k^u, we follow a procedure similar to the one employed in the derivation of L_k^u. By deriving similar recursive equations with respect to c_k, the cumulative in system delay of a session of multiplicity k, and applying the expectation operator, a similar to (2) set of infinite dimensional linear equations with identical coefficients α_{kj} can be derived. An upper bound on C_k $k \geq 0$, of the form

$$C_0^u = 0,\ C_k^u = v_1 k^2 + v_2 k + v_3,\ k \geq 1$$

was obtained for all input traffic rates $\lambda < \lambda^l$ where v_1, v_2, v_3 are some finite constants depending on λ and p and are analytically derived in [3].

4.2.2 Lower Bounds on L_k and C_k:

The theorem that follows the definition below proves that lower bounds on L_k can be obtained by solving truncated (finite) versions of the infinite dimensional equations in (2).

Definition: Majorant/Minorant systems If $x_k = A_k + \sum_{j=0}^{\infty} B_{kj} x_j$, $0 \leq k \leq \infty$ (MAJ) and $y_k = \alpha_k + \sum_{j=0}^{\infty} b_{kj} y_j$, $0 \leq k \leq \infty$ (MIN) are infinite dimensionality linear systems of equations with $A_k \geq |\alpha_k|$ and $B_{kj} \geq |b_{kj}|$, $0 \leq k \leq \infty$, $0 \leq j \leq \infty$, then we say that the system in (MAJ) is a majorant for the system in (MIN); similarly, the system in (MIN) is a minorant for the system in (MAJ).

Theorem 1. *If a majorant for a given system of linear equations has nonnegative solutions x_k, $k \geq 0$, then the given system has the solution y_k which satisfy $|y_k| \leq x_k$, $0 \leq k \leq \infty$.*

The proof of the theorem can be found in [10]. Note that since the infinite dimensionality linear system of equations in (2) has a nonnegative solution for every $\lambda < \lambda^l$ and it is a majorant for its truncated version, Theorem 1 implies that, for every $\lambda < \lambda^l$, the solutions L_k^l of the truncated system (2) are lower bounds on L_k. Lower bounds on C_k are obtained as for L_l by solving truncated version of the corresponding linear equations.

5 Discussion on the Performance Bounding Approach and Related Work

The methodology presented in sections 3 and 4 is fairly powerful and has been applied for the study of processes, protocols and queueing systems appearing in networking. As long as the operation of the system presents renewal cycles (which is typically the case under stability), one can write recursive equations for the length of statistically identifiable sessions appearing within the renewal cycle. Applying the expectation operator on these recursive equations leads to the infinite dimensional system of the form of (2), with non-negative coefficients and constants. Then, the properties of their solutions and the bounds presented in section 4 can be exploited to obtain performance metrics such as the system stability region and mean delay bounds.

First it should be noted that a truncated version of (2) for large N is computationally tractable, especially if one adopts the solution approach based on successive substitution. That is, if one substitutes an arbitrary non-negative sequence L_k^0, $k \geq 0$, to the right-hand side of (2) and derive the sequence L_k^1 in the left-hand side and then repeat the procedure by using the derived sequence L_k^1 in the right-hand side, and so on. It turns out that the solution is reached fairly fast and is of arbitrary accuracy depending on the value of N. By trying an even larger value of N, one can observe a negligible change in the solution for L_k, $k \geq 0$, especially for small k which are heavier contributing (through the Poisson probability of the multiplicities) to the mean renewal cycle length, L, we are interested in. Thus, L is very well approximated by solving a relatively small number N of linear equations.

Deriving upper bounds on L_k, $k \geq 0$, is more cumbersome. In section 4 we presented a linear with respect to k bound on L_k by calculating its coefficients analytically in [3]; a quadratic bound on the cumulative delay is also analytically derived in [3]. As the solutions of the truncated systems of equations for modest N yield a lower bound that is very close to the actual value, the cumbersome derivation of upper bounds on L_k and C_k is practically not necessary. The upper bound on L_k is necessary though in order to derive a lower bound on the maximum throughput, to determine the minimum support capabilities of the system. Although the lower bound on the throughput is very tight away from the instability region, it becomes very loose when approaching the instability region.

The following approaches can help derive tighter upper bounds on L_k (and similarly on C_k). The basic idea behind the first approach is to set L_k for values of k higher than the truncation value N, equal to the (looser) upper bound L_k^u and then solve the resulting finite system of equations to obtain a tighter bound on L_k, for $0 \leq k \leq N$, which are the L_k's contributing heavier to the mean L. The following theorem proves that the solution of these truncated equations constitute indeed an upper bound.

Theorem 2. *Let* $\{x_k^u\}_{k=0}^{\infty}$ *be a sequence of real numbers which satisfies* (α) $0 \leq x_k^u < \infty$, $0 \leq k < \infty$ *and* (β) $h_k + \sum_{j=0}^{\infty} \alpha_{kj} x_j^u \leq x_k^u$ *with* (γ) $h_k \geq 0$, $\alpha_{kj} \geq 0$, *for* $0 \leq k, j \leq \infty$. *Then the following hold.*

a) The finite dimensionality system of linear equations

$$x_k^{ut} = h_k + \sum_{j=N+1}^{\infty} \alpha_{kj} x_j^u + \sum_{j=0}^{N} \alpha_{kj} x_j^{ut} \tag{3}$$

has a nonnegative solution x_k^{ut} *which satisfies* $x_k^{ut} \leq x_t^u$, $0 \leq k \leq N$.

b) If x_k *is a nonnegative solution of the system* $x_k = h_k + \sum_{j=0}^{\infty} \alpha_{kj} x_j$, $k \geq 0$, *then* $x_k \leq x_k^{ut}$, $0 \leq k \leq N$.

By employing the sequences $\{L_k^u\}_{k=0}^{\infty}$ and $\{C_k^u\}_{k=0}^{\infty}$ in the place of the sequence $\{x_k^u\}_{k=0}^{\infty}$ in the above theorem and solving the resulting finite dimensionality linear systems of equations given by (3), tight upper bounds on L_k $((L_k^{ut}))$ and C_k (C_k^{ut}) were obtained, for $k \leq N = 24$. By considering these tight upper bounds, tight upper bounds on L (L^{ut}) and C (C^{ut}) are obtained.

The lower bounds on S_{max}, λ^l, derived under Proposition 1 turned out not to be close to the upper bound, λ^u, derived under Proposition 2. The fact that the value of λ^u obtained by solving $N = 5$ linear equations was the same (up to the third decimal point) with that obtained by solving for $N = 24$ linear equations, suggests that λ^u is very close to S_{max} and that the lower bound is very loose. This belief can be substantiated and justified by re-deriving a lower bound as under Proposition 1 but using the following expressions for upper bounding L_k (suggested in [7] for the first time):

$$x_0^u = 1, \; x_k^u = (1+\epsilon) L_k^l, \; 1 \leq k \leq 7 \tag{4a}$$

$$x_k^u = \beta(\lambda, p) k - \gamma(\lambda, p), \; 8 \leq k \leq \infty \tag{4b}$$

where ϵ is an arbitrary small positive number. A sequence $\{x_k^u\}_{k=0}^{\infty}$ as in (4) which satisfies the conditions of Proposition 1, was possible to obtain for $\lambda < \lambda^{lt}$ (where $\lambda^{lt} \approx \lambda^u$) up to the third decimal digit), and thus $S_{max} \approx \lambda^u$.

The analysis methodology presented here has been applied to the study of other protocols and queueing systems in networking. In [11] tight bounds on the maximum stable throughput and mean packet delay are derived under a random access protocol that can accommodate packets from 2 distinct priority classes. In [12,13], the process of the successfully transmitted packets over a random access channel employing LS-RAP is analytically matched to a first and second order markov model, by constructing recursive equations describing the renewal cycles induced by the protocol, as well as equations counting the occurrences of specific blocks of 2 or 3 consecutive slots over a cycle, to be used to calculate the transition probabilities of the approximating Markov model.

Finally, the methodology has been applied for the delay analysis of complex priority queueing disciplines in [14,15], which were developed in order to model and evaluate the performance of the Distributed Queued Dual Bus (DQDB) metropolitan networks. Complex (priority) queueing disciplines can be relatively easily studied by formulating recursive equations for the mean cumulative delay of all packets from a given priority class that were transmitted over a renewal cycle. The mean delay is calculated by the ration $\frac{C^i}{\lambda^i X}$, where X is the mean renewal cycle length and C^i and λ^i are the cumulative delay over a cycle and arrival rate of the packets of priority i. It should be noted that X is known under *any* queueing discipline and is given by $\frac{1}{1-\rho}$, where ρ is the utilization of the queueing system. Thus, the renewal cycle length and stability conditions (given by $\rho \leq 1$) are already available. Lower bounds on C^i can be computed by solving truncated version of the resulting infinite dimensional linear equations. An upper bound on the induced mean packet delay for class i, D^i_{up}, can be obtained (assuming that the priority discipline is non-preemptive and work-conserving) in terms of their lower bounds D^i_{lo}, and the mean delay, D^{FIFO}, induced by the equivalent FIFO queueing system (which is typically known or easy to derive), from the following expression assuming K priority classes [14]:

$$D^i_{up} = \frac{1}{\lambda^i}[\lambda D^{FIFO} - \sum_{k=1,k\neq i}^{K} \lambda^k D^k_{lo}] \tag{5}$$

6 Conclusions

The contribution of this paper is about a comprehensive presentation of a quite versatile and widely applicable methodology for evaluating the performance of resource allocation protocols and queueing systems appearing in networking.

The methodology consists of two steps. First, the renewal cycles of the system under study are identified and recursive equations that invoke properly identified sessions contained within the renewal cycles are derived. In addition to ultimately yielding the length of the renewal cycle, these recursive equations can help construct (similarly) additional ones with respect to random variables that count key events of interest over such sessions, such as the total delay incurred by the packets, the number of packets experiencing a delay exceeding a threshold, etc; performance metrics of interest can then easily be derived by considering the expected values of such random variables.

The second step of the methodology considers a system of infinite dimensional linear equations that is formulated by taking the expectation on the recursive equations developed in the first step. A rich theory on the properties of the solutions of such equations is presented along with techniques to obtain tight bounds on their solution. Through the rigorous propositions presented, the additional approaches for tightening further the upper bounds, the connection to the stability of the system, and the discussion on different applications of the approach, the potentially wide applicability and effectiveness of the presented methodology have been established. Although high performance computing machines can solve nowadays large dimensionality systems of linear equations, there are at least two reasons for which the presented methodology for coping with infinite dimensional linear equations can still be valuable: for establishing the stability region of a system and for modelling complex systems requiring the formulation of multi-dimensional random variables that lead to an exploding dimensionality of the system equations.

References

1. INFOCOM Network Science Wkshp (NetSciCom 2010), San Diego, USA (2010)
2. Stavrakakis, I.: The Internet of the Future, presentation in session "WP Objective 1.1, WP 2009-2010: Network of the Future". In: ICT 2008, Lyon (2008)
3. Stavrakakis, I.: Random Access Multi-user Communication Networks and Approximation of their Output Processes, Ph.D. dissertation, University of Virginia, UMI order No. 8913912, Ann Arbor Michigan (1988)
4. Stavrakakis, I., Kazakos, D.: A Limited Sensing Protocol for Multiuser Packet Radio Systems. IEEE Transactions on Communications 37(4), 353–359 (1989)
5. Tsybakov, B.S., Vvedenskaya, N.D.: Random Multiple Access Stack Algorithm. Translated from Problemy Peredachi Informatsii 16(3), 80–94 (1980)
6. Vvedenskaya, N.D., Tsybakov, B.S.: Random Multiple Access of Packets to a Channel with Errors. Translated from Problemy Peredachi Informatsii 19(2), 52–68 (1983)
7. Georgiadis, L., Papantoni-Kazakos, P.: Limited Feedback Sensing Algorithms for the Packet Broadcast Channel, Special Issue on Random Access Communications. IEEE Trans. Inform. Theory IT.31, 280–294 (1985)
8. Capetanakis, J.: Tree Algorithm for Packet Broadcast Channel. IEEE Trans. Inform. Theory IT.25, 505–515 (1979)
9. Chung, G.L.: A Course in Probability Theory. Academic, New York (1974)
10. Kantorovich, L.V., Krylov, V.I.: Approximate Methods of Higher Analysis. Interscience, New York (1958)
11. Stavrakakis, I., Kazakos, D.: A Multi-user Random Access Communication System for Users with Different Priorities. IEE Transactions on Communications 39(11) (1991)
12. Stavrakakis, I., Kazakos, D.: On the Approximation of the Output Process of Multiuser Random Access Communication Networks. IEEE Transactions on Communications 38(2) (1990)
13. Stavrakakis, I., Kazakos, D.: Performance Analysis of a Star Topology of Interconnected Networks under 2nd -order Markov Network Output Processes. IEEE Transactions on Communications 38(10) (1990)
14. Stavrakakis, I.: Delay Bounds on a Queuing System with Consistent Priorities. IEEE Transactions on Communications 42(2) (1994)
15. Landry, R., Stavrakakis, I.: Queuing Study of a 3-Priority Policy with Distinct Service Strategies. IEEE/ACM Transactions on Networking 1(5) (1993)
16. http://www.recognition-project.eu/