14

On Compartmental Modelling of Multi-Service Communication Networks

M. Sivabalan, H.T. Mouftah and G. Takahara[†]

Department of Electrical and Computer Engineering
† Department of Mathematics and Statistics
Queen's University
Kingston, Ontario, Canada K7L 3N6

telephone : (613) 545-2934 fax : (613) 545-6615
e-mail : mouftah@eleceng.ee.queensu.ca

Abstract

Compartmental modelling, a technique that is widely used in areas such as bioscience, process control engineering, may serve as the basis for the performance evaluation of systems that are too complex to be handled by exact analytical/computational methods. With this technique, the dynamics of a system are described in terms of a set of differential or difference equations which permits the investigation of both the steady-state and transient behaviour of the system. In this paper, we describe how this technique can be used to model connection-oriented symmetric communication networks supporting multiple service classes, where each class is characterized by its bandwidth requirement, arrival and departure processes. We then demonstrate an application of this modelling technique to the steady-state analysis of a network and examine the accuracy of this modelling technique via simulations.

Keywords
Multi-service networks, network modelling.

1   INTRODUCTION

Compartmental modelling technique may serve as the basis for the performance evaluation of the systems that are too complex to be handled by exact analytical/computational methods. This technique is widely used in areas such as biology, medicine, and ecology. Recently, Garzia and Lockhart (1989) have studied its applicability in understanding the behaviour of connection-oriented communication networks supporting a single class of service. The results of this study indicate that compartmental models are computationally efficient and complex enough to capture the steady-state and transient behaviour of the networks. The goal of this paper is to show how this modelling technique can be applied to multi-service networks.

The rest of the paper is organized as follows: section 2 gives a brief review of compartmental modelling. Section 3 describes the multi-service network model that we consider in this paper. Section 4 explains the formulation of the compartmental model for the multi-service network. In section 5, we present and discuss the numerical results obtained using the compartmental model from the steady-state analysis of a network with two service classes. Section 6 assesses the accuracy of the compartmental model through comparison with simulations, and section 7 concludes this paper.

2   REVIEW OF COMPARTMENTAL MODELLING

Consider a system consisting of a finite number of distinct material types. In such a system, a *compartment* can be considered as a sub-system containing a single material type (Godfrey, 1983). In compartmental modelling, the system is visualized as an interconnected system of compartments. Compartments may exchange materials with each other and with the environment in accordance with the law of conservation. Transformation of a material type i into another material type j can be represented by a flow from compartment i to j. Also, the arrival of material type i into the system and the departure of material type i from the system can be represented by a flow from the environment to the compartment i and a flow from the compartment i to the environment respectively.

A system with n compartments can be described by the following differential equations:

$$\frac{dx_i}{dt} = f_{i0} + \sum_{\substack{j=1 \\ j \neq i}}^{n} (f_{ij} - f_{ji}) - f_{0i}, \tag{1}$$

where $i = 1, 2, \ldots n$, x_i is the amount of material in the compartment i and f_{ij} is the rate of flow from the *donor* compartment j to the *acceptor* compartment i, with 0 denoting the environment (see Figure 1).
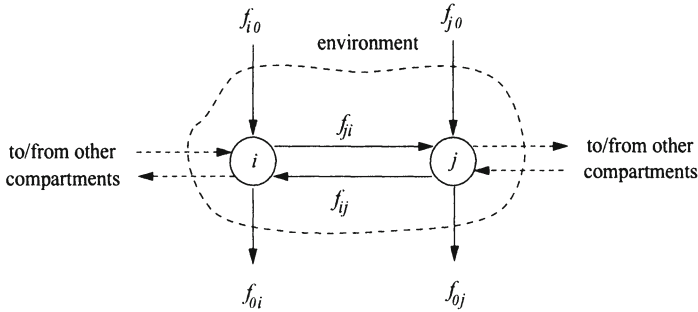
Figure 1 Two compartments of an n-compartment $(n > 2)$ system.

In the simplest case, it is adequate to model the compartments as donor-controlled systems. In such systems, the rate of flow f_{ji} is related, linearly or nonlinearly, to the amount of material only in the donor compartment i. However, in a more complex model, the flow rate depends on the amount of material in not only the donor compartment but also in one or more other compartments. In general, the rate coefficients can be chosen to capture the non-linear, time-dependent, and state-dependent nature of a very complex system as follows:

$$
\begin{aligned}
f_{ij} &= a_{ij}(\mathbf{x})x_j(t) \\
&\quad i = 0, 1, \ldots, n \\
&\quad j = 1, 2, \ldots, n; \ j \neq i \\
f_{i0} &= b_i(\mathbf{x}) \\
&\quad i = 1, \ldots, n,
\end{aligned}
\tag{2}
$$

where $\mathbf{x} = (x_1(t), x_2(t), \ldots, x_n(t))$ is the system state vector at time t, $a_{ij}(\mathbf{x})$ is the rate coefficient of the flow from the compartment j to i, and $b_i(\mathbf{x})$ is the flow rate from the environment to the compartment i (i.e., input). The equation 1 and 2 can be combined to form the following set of first-order differential equations:

$$
\frac{dx_i(t)}{dt} = b_i(\mathbf{x}) + \sum_{\substack{j=1 \\ j \neq i}}^{n} \left(a_{ij}(\mathbf{x})x_j(t) - a_{ji}(\mathbf{x})x_i(t) \right) - a_{0i}(\mathbf{x})x_i(t),
\tag{3}
$$

where $i = 1, 2, \ldots, n$. We refer to these equations as *state equations*. By solving these equations, we can obtain the steady-state as well as the transient behaviour of the system in terms of the quantity of interest $\mathbf{x}$.

Garzia *et al.* (1989) have applied the compartmental modelling technique to investigate the behaviour of a single service, circuit-switched networks. Because of the state-dependent nature of such networks, the resulting state equations are very complex and hard to solve. To overcome this difficulty, the authors have transformed the differential equation 3 into the difference equation 4 and analyzed the system in discrete-time domain.

$$\frac{x_i(t + \tau) - x_i(t)}{\tau} = b_i(\mathbf{x}) + \sum_{\substack{j=1 \\ j \neq i}}^{n} (\ a_{ij}(\mathbf{x})x_j(t) - a_{ji}(\mathbf{x})x_i(t)\ ) - a_{0i}(\mathbf{x})x_i(t), \quad (4)$$

where $i = 1, 2, \ldots, n$, and τ is the time-step. In section 4, we will show how the single service compartmental model developed by Garzia *et al.* (1989) can be extended to the symmetric multi-service network to be described in the following section.

3 NETWORK MODEL

Consider a fully connected symmetric network supporting multiple service classes. Each pair of network nodes is assumed to be connected by a pair of uni-directional links (one for each direction). Each service class is characterized by its bandwidth requirement, average arrival rate, and average holding time. We assume connection-oriented information transfer mode, e.g., ATM (de Prycker, 1991). In such a mode, upon arrival of a call, the network selects a path between the source and destination and reserves its resources along the chosen path. The information transfer for the call takes place on the chosen path. When the call is terminated, the resources reserved for that call is released. In this paper, we consider the link bandwidth as the only resource required for setting up calls, i.e., a call is blocked on a link if the link does not have enough bandwidth.

We assume that the link bandwidth is "quantized" into Basic Bandwidth Units (BBU's) and all links are of equal bandwidth. The bandwidth required to support a call from any class can be specified as an integral multiple of BBU. The calls are assumed to be bi-directional. We consider paths of at most two hops. The single and two hop paths are referred to as *direct* and *alternate* paths respectively. When a call set-up request arrives, the network first attempts to connect the call on the direct path. If this attempt fails, the network will try an alternate path to set up the call. When probing an alternate path, hops corresponding to that path are probed one after the other. We also assume that the time required to probe and establish a connection on a hop is constant, and both links corresponding to a hop are probed simultaneously.

The network is assumed to be spatially homogeneous, i.e., the network load is uniformly spread across the network. So, each link is offered an identical load. We further assume that the connection requests arrive at a link according to Poisson process. With these assumptions, the network blocking problem is reduced to the link blocking problem where the links block the calls independently (link independence assumption). Calls from a given service class are blocked with identical probabilities, and different classes could experience different blocking probabilities.

We will use the following notations for the network and call related parameters throughout the rest of this paper:

- N : number of nodes in the network.

- K : number of service classes supported by the network.

- C : link bandwidth in BBU's.

- τ : time to probe and establish connection on the pair of links between two nodes. This is what is going to be used as the time-step (see equation 4) in the state equations to be derived in section 4.

- b_i : bandwidth required (in BBU's) for a class i call.

- λ_i : external arrival rate (in calls/sec) of class i calls.

- $\frac{1}{\mu_i}$: average call holding time (in seconds) of a class i call.

- $\rho_i(t)$: link offered load at time t due to class i calls.

- $q_i(t)$: link blocking probability that class i calls experience at time t. Also, $p_i(t) = 1 - q_i(t)$.

- r_i: number of alternate paths between a source-destination pair used for class i calls.

4 FORMULATION OF THE COMPARTMENTAL MODEL

Between its arrival and departure, a call can be in a number of distinct phases (see Lee, Hluchyi and Humblet, 1995 for a detailed description of these phases). In this paper, we consider only two such phases. The first one is the probe phase in which the call is attempting to get a path, and the second one is the information transfer phase in which the probe is successful and the call is in the process of transferring information. However, it is possible to incorporate the other phases into the compartmental model to be developed later in this section. The compartmental modelling requires the selection of appropriate compartments. The network is treated as a "black box", and the compartments are chosen in such a way that each compartment contains the calls that are in a particular state.

As in Garzia *et al.* (1989), we define the following compartmental variables for class $i, i = 1, 2, \ldots, K$, at time t:

- $P_i(t)$: average number of calls probing the direct paths.

- $P_i^{1,j}(t)$: average number of calls probing the 1 st hop of the j th alternate path, where $j = 1, 2, \ldots, r_i$.

- $P_i^{2,j}(t)$: average number of calls probing the 2nd hop of the jth alternate path, where $j = 1, 2, \ldots, r_i$.

- $C_i^1(t)$: average number of established calls using direct paths.

- $C_i^2(t)$: average number of established calls using alternate paths.
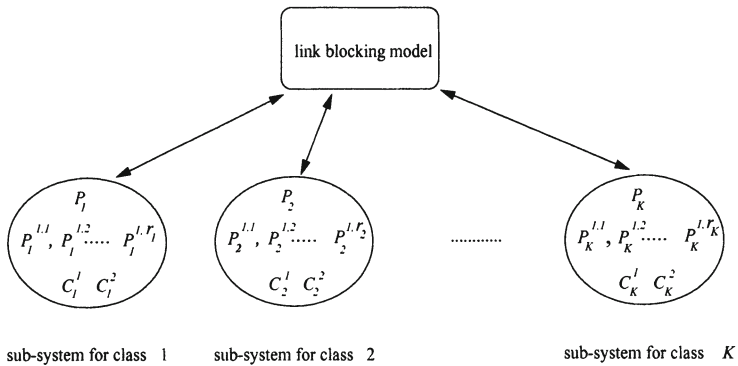
Figure 2 Compartmental sub-systems.

The compartments belonging to a given service class collectively form a compartmental sub-system as shown in Figure 2. In deriving the state equations, a sub-system is treated independently of others assuming that the link blocking probability of the corresponding service class is known. Moreover, depending on the bandwidth sharing policies (to be discussed later), the blocking probability of one service class may depend on the link offered load from the other classes. We account for the possible interactions among the compartmental sub-systems through link blocking as shown in Figure 2.

Difference Equations

Figure 3 illustrates the flow among the compartments of service class i. In deriving the state equations, the system is assumed to be in equilibrium during the interval $(t, t + \tau]$ (different equilibriums for different values of t), and the steady-state value of a variable at time $t + \tau$ is computed using its value at time t. By inspecting Figure 3, we can obtain the difference equations 5 through 10 relating the compartments of class i calls.

$$P_i(t + \tau) \;=\; \lambda_i\, \tau \tag{5}$$

$$C_i^1(t + \tau) \;=\; C_i^1(t) - \tau\, \mu_i\, C_i^1(t) + p_i(t)\, P_i(t) \tag{6}$$

$$P_i^{1,1}(t+\tau) = \begin{cases} q_i(t)\, P_i(t) & , \ r_i \geq 1 \\[2mm] 0 & , \ otherwise \end{cases} \tag{7}$$
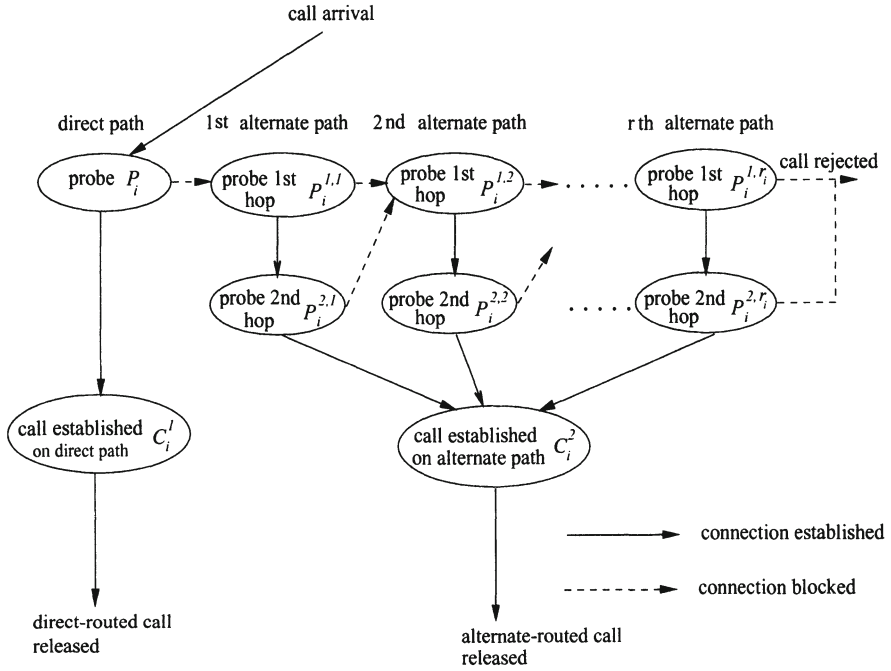

Figure 3 Flow diagram for service class i.

$$P_i^{1,j}(t+\tau) = \begin{cases} q_i(t)\, (P_i^{1,j-1}(t) + P_i^{2,j-1}(t)) & , \ r_i \geq 2, \ j = 2, \ 3, \ldots, r_i \\[2mm] 0 & , \ otherwise \end{cases} \tag{8}$$

$$P_i^{2,j}(t+\tau) = \begin{cases} p_i(t)\, P_i^{1,j}(t) & , \ r_i \geq 1, \ j = 1, \ 2, \ldots, r_i \\[2mm] 0 & , \ otherwise \end{cases} \tag{9}$$

$$C_i^2(t+\tau) = C_i^2(t) - \tau\, \mu_i\, C_i^2(t) + p_i(t) \sum_{j=1}^{r_i} P_i^{2,j}(t) \tag{10}$$

The above difference equations involve the link blocking probability $q_i(t)$, which depends on the link offered load $\rho_i(t)$. The latter is the ratio between the arrival and departure rates of the connections over a link. At time t, the arrival rate $L_i^a(t)$ of connections to any link from service class i can be computed as follows: each of the calls in the states $P_i(t)$, $P_i^{1,j}(t)$ and $P_i^{2,j}(t)$, where $j = 1, 2, \ldots, r_i$, attempts to capture b_i BBU's on both links belonging

to a hop (since the call is bi-directional) in time τ. Since the total number of links in a fully connected network with N nodes is $N(N-1)$, the arrival rate is obtained as:

$$L_i^a(t) = \frac{2}{N(N-1)\,\tau}\,(P_i(t) + \sum_{j=1}^{r_i} P_i^{1,j}(t) + \sum_{j=1}^{r_i} P_i^{2,j}(t)). \tag{11}$$

In order to find the departure rate $L_i^d(t)$ of class i connections from a link, we note that the calls in state $P_i^{2,j}(t)$ will release the bi-directional connections established on the 1 st hop of the j th alternate paths in time τ. Further, calls in state $C_i^1(t)$ and $C_i^2(t)$ will release one and two, respectively, bi-directional connections in time $\frac{1}{\mu_i}$ (time required to tear down a connection is neglected). The average departure rate is obtained as:

$$L_i^d(t) = \frac{\sum_{j=1}^{r_i} \dfrac{P_i^{2,j}(t)}{\tau} + \mu_i\, C_i^1(t) + 2\,\mu_i\, C_i^2(t)}{\sum_{j=1}^{r_i} P_i^{2,j}(t) + C_i^1(t) + 2\, C_i^2(t)}. \tag{12}$$

The link offered load from service class i is given by $\rho_i(t) = \frac{L_i^a(t)}{L_i^d(t)}$.

The equations 5 through 10 can be used to compute instantaneous values of the state variables at any given time t. Steady-state values of the variables can also be obtained by iterating these equations for sufficiently long time. Now all what we need is the blocking probability for each class for solving the state equations. The following section describes the procedures for computing the blocking probabilities.

Link Blocking Probability

We model the link as a multi-service loss system (Ross, 1995) in which an arriving connection request is either admitted or blocked. An admitted connection remains in the system for the duration of its holding time. The blocked connection is assumed to be lost. The bandwidth sharing policy dictates how link bandwidth is shared among different services. The blocking probability depends on the bandwidth sharing policy used. In this paper, we consider two popular sharing policies, namely *complete sharing* (CS) and *complete partitioning* (CP). During the iteration of the state equations, the blocking probabilities are recomputed at each step. Therefore, these quantities must be calculated using computationally efficient procedures to minimize the time required to carry out iterations. We now outline such procedures for CS and CP policies.

The connection requests from service class i arrive with traffic intensity ρ_i according to a Poisson process. Let us denote the state of the system

by $\mathbf{n} = (n_1, n_2, \ldots, n_K)$, where n_i is the number of connections from class i. Under CS policy, a class i connection request is accepted on the link in state $\mathbf{n}$ only if $\sum_{i=1}^{K} b_i \, n_i + b_i \leq C$. Using the recursive algorithm developed independently by Kaufman (1981) and Roberts (1981), the probability P_n that exactly n BBU's are occupied on the link can be computed as follows:

$$n \, P_n = \sum_{i=1}^{K} \rho_i \, b_i \, P_{n-b_i}, \qquad n = 1, 2, \ldots, C, \tag{13}$$

where $\sum_{n=0}^{C} P_n = 1$, and $P_n = 0$ if $n < 0$. The steady state blocking probability for class i connections is given by:

$$q_i = \sum_{n=C-b_i+1}^{C} P_n, \qquad i = 1, 2, \ldots, K. \tag{14}$$

Under CP policy, link bandwidth is divided into K partitions with capacities $(C_1, C_2, \ldots, C_K)$ where partition C_i is dedicated only to class i. So, in state $\mathbf{n}$, a new class i connection request is accepted only if $(n_i + 1)b_i \leq C_i$. Since there is no overlap among the partitions, the blocking probability for class i connections can be computed independently of other classes using the Erlang-B formula:

$$q_i = B(M_i, \rho_i) = \frac{\frac{\rho_i^{M}}{n!}}{\sum_{i=0}^{M_i} \frac{\rho_i^{M}}{n!}}, \tag{15}$$

where M_i is the largest integer less than or equal to $\frac{C_i}{b_i}$. $B(M_i, \rho_i)$ can be computed fast using the following recursive formula:

$$\frac{1}{B(n, \rho_i)} = 1 + \frac{n}{\rho_i} \frac{1}{B(n-1, \rho_i)}, \qquad i = 1, 2, \ldots, M_i, \tag{16}$$

with $B(0, \rho_i) = 1$.

5 A NUMERICAL EXAMPLE

In this section, we use the compartmental model to evaluate the steady-state performance of a fully connected 10-node network. The network supports two service classes ($K = 2$), namely *class 1* and *class 2* services. Each link has a bandwidth of 100 BBU's. A class 1 call requires 1 BBU and has a mean holding time of 1 minute, whereas a class 2 call requires 3 BBU's and has a mean holding time of 3 minutes. We further assume that 75% of the arriving calls are from class 1 and the remaining 25% are from class 2. Let's denote the aggregate call arrival rate from both class 1 and 2 to the network by λ^a. Then, $\lambda_1 = 0.75\lambda^a$ and $\lambda_2 = 0.25\lambda^a$. These parameters are selected in such

a way that each service class offers equal load to the network, i.e., $\frac{\lambda_1}{\mu_1} = \frac{\lambda_2}{\mu_2}$.
We set the number of alternate paths to 6 for both classes ($r_1 = r_2 = 6$),
and the network load is changed by varying λ^a. The time step τ used for this
study is 0.1 second. The system is brought to the steady-state after iterating
the state equations for 5000 seconds (requires 50 000 iterations). We use the
average number of calls in the network as the performance metric. For class
i calls, the value of this metric is obtained as a summation of C_i^1 and C_i^2.
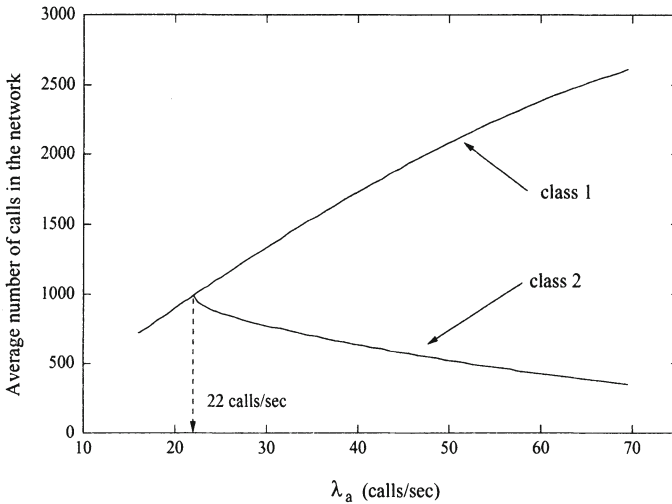
Figure 4 Variation of average number of calls in the network with λ^a under
CS policy

Figure 4 shows the impact of λ^a on the number of calls in the network
under CS policy. As λ^a increases, the number of calls from both classes
increases steadily until λ^a reaches approximately 22 calls/s. As λ^a increases
further, the number of class 1 calls continues to increase whereas that of class
2 calls decreases. We observe that, under heavy load ($\lambda^a > 22$ calls/s), the
link bandwidth is monopolized by class 1 calls, i.e., the network is unfair to
class 2. This observation supports the well known fact that CS policy favours
services with smaller bandwidth requirements.

Class 2 calls can be protected using other bandwidth sharing policies. For
example, let us consider CP policy. Since there is no interference between
the service classes, competition for link bandwidth is avoided. Figure 5 shows
the variation of the number of calls in the network with λ^a under CP policy.
Two partitioning cases are considered. In the first case, link bandwidth is
equally divided between the two classes, and in the second case, 75% of the
link bandwidth is allocated for class 2 and the rest is allocated for class 1.
Note that, in the second case, a partition for a class is proportional to the
bandwidth requirement of the class. Unlike CS policy, CP policy keeps the
number of class 2 calls steady at higher load. However, from Figure 4 and 5,
we clearly see that the network performance with respect to class 2 service

is improved at the cost of blocking more class 1 calls. Note also that, under CP policy, the number of calls is sensitive to the partition setting. With the first setting, the number of class 1 calls is much higher than that of class 2, whereas the second setting results in equal number of calls from each class throughout the entire loading region.
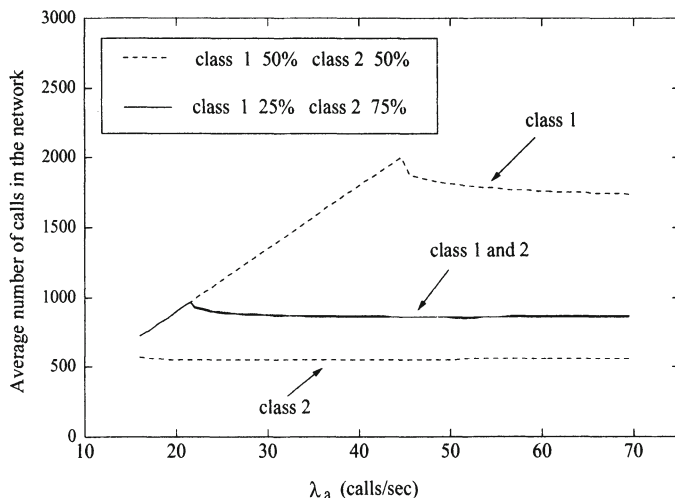

Figure 5 Variation of average number of calls in the network with λ^a under CP policy.

We will now have a closer look at the behaviour of the network when λ^a approaches 22 calls/sec (in the increasing direction) under CS policy (see Figure 4). For the purpose of this discussion, let us call this rate the critical rate. Figure 6 illustrates the impact of λ^a on the link blocking probability. Observe that, below the critical rate, both classes experience very small blocking. But when λ^a slightly exceeds the critical rate, the link blocking probability suddenly becomes very high for both classes, with class 2 calls experiencing higher blocking than class 1 calls. This phenomena can be explained as follows: as λ^a approaches the critical rate, the link blocking becomes significant and hence calls probing the direct path experience more blocking. Consequently, the number of alternate-routed calls begins to increase. The increasing alternate-routed calls offer more load to the links (because they use two hops) and therefore lead to a further increase in the link blocking. This "domino" effect causes the drastic increase in the link blocking probability and the number of alternate-routed calls. Since an alternate-routed call uses twice the amount of resources as a direct-routed call and the amount of network resources is fixed, the domino effect causes the total number of calls to drop quickly as λ^a exceeds the critical rate. This is evident from the slope of the curve for class 2 in Figure 4 in the vicinity of the critical rate. Similar observation has been made by Akinpelu (1984) in telephone networks.

For the selection of the network parameters used in this example, class 1

calls are not affected by the domino effect. This is because under CS policy, class 1 service gets the extra bandwidth required to accommodate its increasing alternate-routed calls by further depriving the class 2 service of the link bandwidth. Under CP policy, class 1 service also exhibits a sudden reduction in the number of calls due to domino effect (see Figure 5). In this case, however, the value of λ^a at which the domino effect begins is different from that under CS policy, and depends on how the link is partitioned. Class 1 service now obtains the additional bandwidth required to accommodate its increasing alternate-routed calls from its own partition as the links switch to a higher blocking state. So, the increase in the alternate-routed calls causes a sudden drop in the total number of class 1 calls in the network.
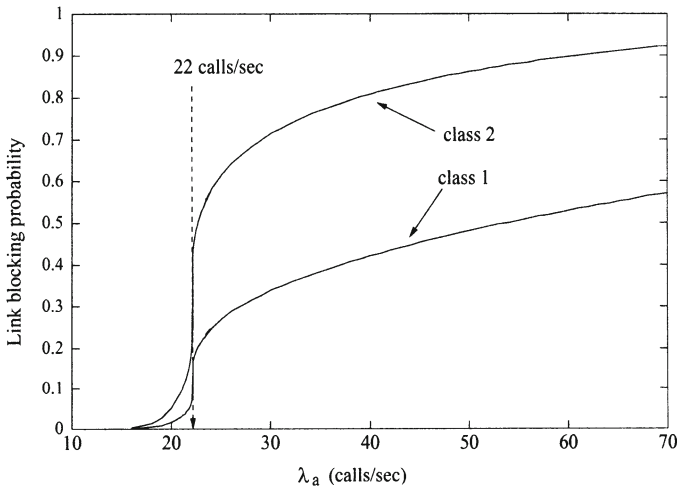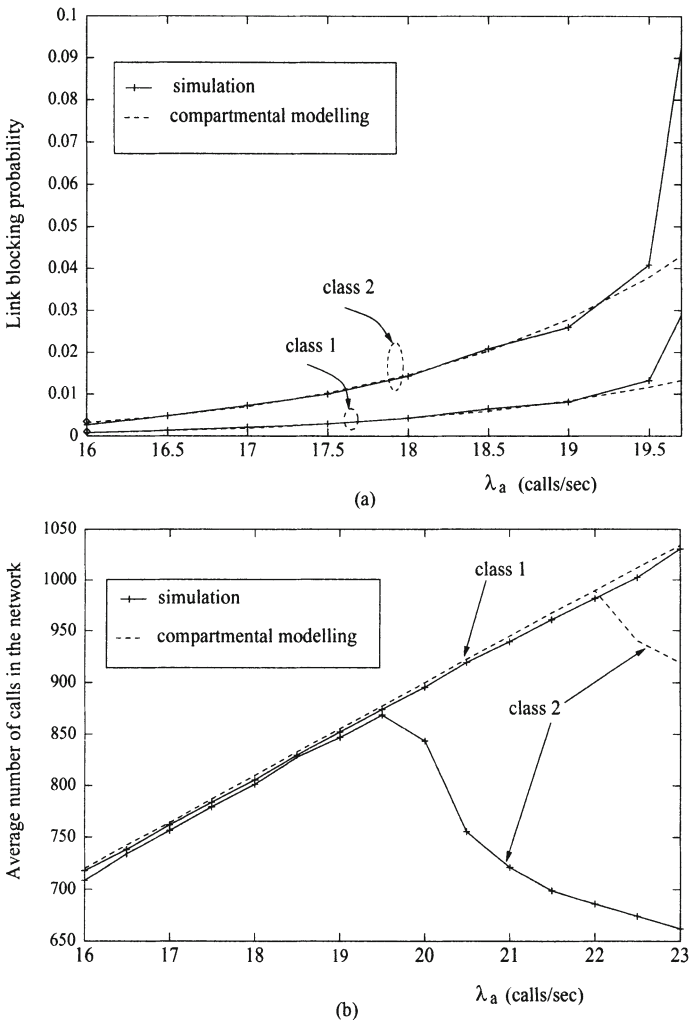

Figure 6 Variation of link blocking probability with λ^a under CS policy.

6   ACCURACY OF THE COMPARTMENTAL MODEL

The compartmental model aggregates the system description into just a few state variables and is based on the solution of the classical single-link blocking model. These features reduce the complexity of solving the state equations. Due to somewhat oversimplifying assumptions, we cannot expect the compartmental models to yield very accurate results. In this section, we use call level discrete-event simulations to assess the the results obtained using the compartmental model.

For the simulations, we consider an alternate routing scheme (responsible for selecting alternate paths) known as the *least busy alternate routing* in which the alternate path with the highest residual capacity is chosen first (see Ross, 1995 for details). We consider only CS policy. Figure 7(a) and (b) shows the variations of link blocking probability (averaged over all links) and the number of calls respectively with the load. In the region where $\lambda^a \leq 19.5$, the compartmental model yields results that are in good agreement with those of

simulations. As the load increases the simulation results quickly deviates from the analytical results. However, the critical aspects of the network behaviour remain the same. For instance, according to the simulation results, the number of class 2 calls begins to drop when λ^a approaches 19.5 calls/sec (in the increasing direction) and the onset of this event is associated with a drastic increase in the link blocking. The compartmental model also predicts a similar trend. In this case, the number of class 2 calls begins to drop when λ^a slightly exceeds 22 calls/sec (note also the drastic increase in the link blocking probability in Figure 6).

(a)

(b)

Figure 7 Comparison of compartmental modelling with simulations with respect to (a) link blocking probability (b) average number of calls.

In compartmental modelling, we do not consider a particular routing

scheme. Instead we assume that the routing scheme operates in such a way that the load is uniformly spread across the network. This condition may be difficult to achieve by a real routing scheme such as the one used for our simulations. We have also carried out additional simulation with different parameters, e.g., network size, bandwidth sharing policy, etc. The main observation is that the compartmental model predicts the critical aspects of the network behavior.

7   CONCLUSIONS

The performance of communication networks can be evaluated through simulation. It is easy to develop simulation models that capture the fine aspects of network elements. A major concern with this approach is the long simulation time. On the other hand, it is very difficult to develop analytical models which accurately represent all the characteristics of a network. Furthermore, many analytical models developed are suitable only for steady-state analysis.

Compartmental modelling technique yields a good compromise between accuracy and complexity of analytical approach, and can be used for both steady and transient analysis of complex networks. This technique has been applied to connection-oriented symmetric networks carrying a single service class. We have shown how this technique can be easily adapted to multiservice environment, and demonstrated its application to the steady-state performance analysis of a 10-node network supporting two service classes. Comparison with simulation results indicates that, the compartmental models can be useful for qualitatively investigating the critical aspects of symmetric networks.

8   REFERENCES

Akinpelu, J. M. (1984) The Overload Performance of Engineered Networks with Non-hierarchical Routing, *AT & T Bell Labs. Tech. J.*, Vol. **63**, No. **7**,1261-81.

de Prycker, M. (1991) Asynchronous Transfer Mode: A Solution for Broadband ISDN, Ellis Horwood Ltd., New York.

Garzia, M. R. and Lockhart, C. M. (1989) Nonhierarchical Communications Networks: An Application of Compartmental Modeling, *IEEE Transactions on Communications*, Vol. **37**, No. **6**, 555-64.

Godfrey, K. (1983) Compartmental Models and Their Application, Academic Press, London.

Kaufman, J. S. (1981) Blocking in a shared resource environment, *IEEE Transactions on Communications*, Vol. **29**, No. **10**, 1474-81.

Lee, W. C., Hluchyi, M. G., and Humblet, P. A. (1995) Routing Subject to Quality of Service Constraints in Integrated Communication Networks, *IEEE Network Magazine*, Vol. **9**, No. **4**, 46-55.

Roberts, J. W. (1981) A service system with heterogeneous user requirements, *Performance of Data Communications Systems and their Applications*, 423-31, North-Holland.

Ross, K. (1995) Multiservice Loss Models for Broadband Telecommunication Networks, Springer, New York.

9 BIOGRAPHY

M. Sivabalan received the B.Sc. degree in Electrical Engineering from the University of Peradeniya, Sri Lanka, and the M.Sc. degree in the area of Cryptography from the Department of Electrical and Computer Engineering, Queen's University, Kingston, Ontario, Canada, where he is now a Ph.D. candidate. His current research area is routing in high speed networks.

H. T. Mouftah is a Professor in the Department of Electrical and Computer Engineering, Queen's University, Kingston, Ontario, Canada since 1979. Prior to that, he worked with the Data Systems Planning Department at Nortel Technology (then BNR) of Ottawa for three years. He has also spent his sabbatical years 1986/87 and 1993/94 with Nortel Technology working in the area of high speed integrated networks, routing, and traffic management. Since 1989, he is a Principal Investigator for the Telecommunications Research Institute of Ontario (TRIO), a government Centre of Excellence in Communications responsible for a project on Broadband Packet Switching Networks.

Dr. Mouftah is the recipient of the 1989 Engineering Medal for Research and Development of the Professional Engineers of Ontario (PEO). He is the joint holder of an Honorable Mention for the Frederick W. Ellersick Prize paper award for the best paper in the Communications Magazine in 1993. He is also the joint holder of the Outstanding Paper Award for a paper presented at the IEEE 14th International Symposium on Multiple-Valued Logic (ISMVL'84). He is the recipient of the IEEE Canada (Region 7) Outstanding Service Award (1995). Dr. Mouftah is a fellow of the IEEE (1990), and Member of the PEO and the Canadian Association of University Teachers.

G. Takahara received the B.A. degree in Mathematics from the University of British Columbia, Vancouver, B.C., and the M.S. and Ph.D. degrees in Statistics from Carnegie Mellon University. Since 1994, he has been an Assistant Professor in the Department of Mathematics and Statistics at Queen's University in Kingston, Ontario, Canada. His research interests are in applied probability modelling in communications networks, loss networks, retrials, circuit switching, and queueing theory. He is a member of the IEEE Communications Society.