The Effects of ABR Traffic on CBR Traffic

David Akehurst, Peter Linington, Andrew Symes and Gill Waters

Computing Laboratory
University of Kent at Canterbury
Canterbury
CT2 7NZ

1   Introduction

One recent development by the ATM Forum is the Available Bit Rate [1], or ABR, service class. The motivation behind this was to develop a service class that would make effective use of the remaining bandwidth unused by other traffic sources. A Source End System (SES) that requests an ABR connection is able to change its transmission rate according to the network loading, during the connection.

Two mechanisms for controlling the transmission rate of the SES that were considered were the Explicit Forward Congestion Indication (EFCI) mechanism and the Explicit Rate (ER) mechanism [1, 2]. The former uses an additive constant to increase its transmission rate or a multiplicative constant to decrease it. The latter relies on an explicit rate being set by the network that will utilise the remaining available bandwidth. These two mechanisms are detailed in the following sections.

Ideally, ABR traffic would make efficient use of the network without affecting other network traffic. However, the control mechanisms are not instantaneous, and may influence the quality of service of other network traffic. Whilst one study has investigated the performance of both ABR and Variable Bit Rate (VBR) traffic [3], this did not address the issue of Constant Bit Rate traffic (CBR). Therefore it is the intention of this paper to attempt to predict the potential effect on CBR traffic.

Section 2 of the paper discusses the ABR rate control and feedback mechanisms in detail. Section 3 describes the simulation models used. Simulation results are presented in Section 4, and conclusions in Section 5.

2   ABR Mechanisms

2.1   Explicit Forward Congestion Indication

The Explicit Forward Congestion Indication (EFCI) mechanism relies on individual switches on the path of a virtual circuit (VC) notifying the ABR SES of congestion. This information is conveyed in Resource Management (RM) cells that are generated periodically by the ABR SES

and injected into the normal ABR data stream. A count, based on the number of ABR data cells sent by the SES, is used to determine when RM cells are generated, as opposed to using a time based interval.

When a switch is congested, the congestion indicator bit in all forward travelling cells is set. A Destination End System (DES) receiving cells with this bit set must notify the ABR SES to reduce its transmission rate. RM cells returned by the DES are set to indicate that congestion has occured.

The ABR SES uses returned RM cells to determine whether the transmission rate should be increased or decreased. If the congestion bit in the RM cell is set, the transmission rate is decreased, and increased otherwise. Furthermore, if no RM cells are received within a specified period, congestion is assumed and the transmission rate is decreased. Again, this period is based on a count of the transmitted cells.

The ABR SES increases its transmission rate by a constant factor, the Additive Increase Rate (AIR), using the simple relationship:

new transmission rate = current transmission rate + AIR

The ABR SES decreases its transmission rate by a proportion of its current transmission rate. This proportion is referred to as the Multiplicative Decrease Factor (MDF) and is used in the relationship

new transmission rate = current transmission rate - (current transmission rate * MDF)

2.2 Explicit Rate

The Explicit Rate (ER) mechanisms also relies on individual switches on the path of a VC monitoring network loading. However, unlike the EFCI mechanism, a given switch calculates a transmission rate for the ABR SES that will meet a predefined target utilisation of the link. This information is propagated through the network using RM cells. If the explicit rate calculated by a switch is less than that in a received RM cell then it replaces the value in the RM cell. In this way, when the RM cell is finally returned by the DES to the ABR SES, the SES transmission rate is set to the minimum available bandwidth of all switches.

Several algorithms have been proposed to calculate the explicit rate [4, 5, 6, 7]. One of the algorithms considered here was developed at Ohio State University by Raj Jain, et. al. [8], and is referred to as the OSU scheme. The paper details several variations of this scheme, the simplest of which has been used in the simulation detailed here. The aim of the OSU scheme is for each link to be operating close to, but not at, 100% utilisation. This is defined as the *target utilisation* (e.g. 90%). For each link, given a target utilisation, a *target cell rate* can be calculated. In order to reach target utilisation, each switch calculates a load level from the input rate. If this is greater than an outward link's target utilisation, the link is overloaded. Conversely, if the load level is less than the link's target utilisation, the link is underloaded.

Efficient use of the network is achieved by avoiding bottlenecks at any switch. This means that the ABR transmission rate must be kept to the lowest value acceptable to any of the switches on the path of a VC. In terms of the OSU scheme this requires that the highest load level calculated by the switches on a VC path be returned to the ABR SES. The transmission rate of the SES can then be adjusted accordingly. Load levels are returned to the SES by RM cells that are generated on a time based interval as opposed to the cell based count used by EFCI mechanisms.

Once a link is operating close to its target utilisation, the OSU scheme attempts to share the used

bandwidth fairly amongst the VCs. A "fair share" of the bandwidth is determined simply to be the target cell rate divided by the number of active VCs. Within the OSU scheme, each ABR SES monitors its own transmission rate which is averaged over a given period. This information is recorded in any RM cells that are generated. If the transmission rate of an ABR SES is below the fair share then it is allowed to increase its transmission rate, and conversely it has to decrease its transmission rate if it is above the fair share.

A later adaptation of this scheme, ERICA+ [9], attempts to address some of the issues not covered by the OSU scheme [1]. The ERICA+ scheme indicates that only ABR sources contribute to the active VC total. Furthermore, the bandwidth available to ABR traffic is calculated by removing the VBR rate (averaged over some period) from the link bandwidth. The issue of CBR traffic is still not addressed in the scheme. Presumably this is also removed from the available bandwidth as with VBR traffic; however this is not explicitly stated. These adaptations of the OSU scheme were also implemented and simulations conducted.

2.3 Resource Management Control Loop

In order to implement either mechanism, information regarding network congestion and available bandwidth must be conveyed to the ABR SES. This is done by means of Resource Management (RM) cells that are generated by the ABR SES and sent to the terminus of the VC, or Destination End System (DES). En route the RM cells record network information. The DES returns RM cells to the ABR SES which then adjusts its transmission rate accordingly.

The path taken by RM cells in a VC is referred to here as the Resource Management Control Loop (RMCL). This has already been described as an *end-to-end* circular path from ABR SES to DES and back again to SES. However, the RMCL may also be segmented into a number of shorter RMCLs, referred to here as *virtual source and destination* RMCLs. In this situation a switch at the ingress of a segment acts as a virtual destination by returning any RM cells it receives. Similarly, a switch at the egress of a segment behaves as a virtual source, generating and inserting new RM cells into the data flow. In the limiting case, each switch on a VC can act as a virtual source and destination for RM cells.

3 Simulation Models

Four network models were designed to represent different rate control mechanisms and resource management control loops. These were:

EFCIETE: This model used EFCI to control the transmission rate of the ABR SES. An End-To-End RMCL was used, whereby RM cells generated by the ABR SES travelled to the DES where they were turned around and sent back to the SES.

EFCIVSD: This model also used the EFCI control mechanism at the ABR SES. However, a Virtual Source and Destination RMCL was employed. In this case each switch in the VC acted as a virtual source and destination for RM cells, using the same parameters as the ABR SES.

---

[1] Neither the description nor the pseudo code for the OSU scheme indicate whether all VCs (ABR, CBR and VBR) can be classed as active or whether the link bandwidth used to calculate the target utilisation refers to the entire bandwidth of the link or the remaining bandwidth after allocation to non-ABR sources. It was decided to implement the scheme verbatim, thus active VCs included ABR, CBR and VBR VCs, and the link bandwidth referred to the entire bandwidth.
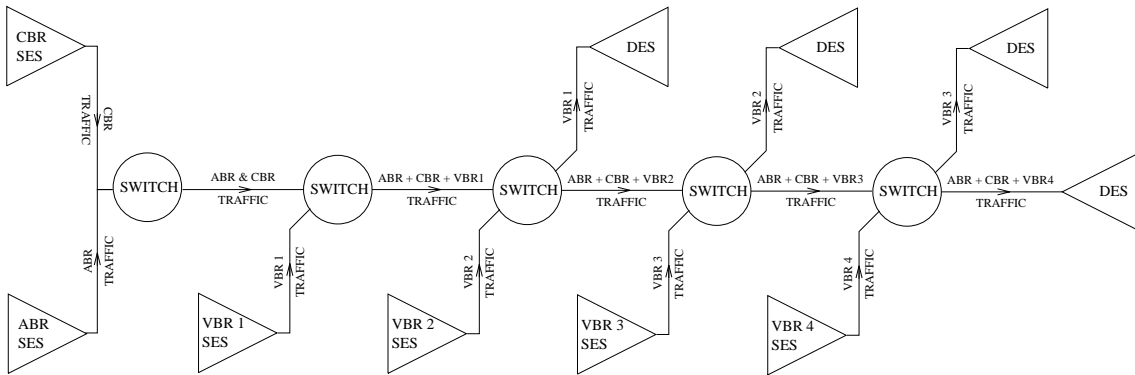
Figure 1: Network and Traffic Topology

At each virtual source, only ABR data cells were used to control the rate at which new RM cells were generated [2].

OSU: This model used explicit rate feedback to control the transmission rate of the ABR SES. The OSU scheme was used to determine the explicit rate calculated by switches within the VC. The RMCL was not segmented, with RM cells following an End-To-End, SES to DES, loop.

ERICA+: This model used the ERICA+ scheme to determine the ER calculation for control of the ABR SES transmission rate. An End-To-End RMCL was used to control RM cell routeing.

Each of the four models was developed using the simulation tool SES Workbench, which provides a design and specification environment (SES/$design^{TM}$) [10, 11]. Models were then simulated using the SES discrete event simulation engine (SES/sim^{TM}) [12].

A simple VC configuration was adopted for the ABR and CBR traffic that consisted of two separate SESs, five output buffered ATM switch nodes and a single DES. The VCs for both ABR and CBR traffic were assumed to consist of the same switches and links. Loading on the first switch consisted of only CBR and ABR traffic. Subsequent switches were also loaded with VBR traffic to represent additional background traffic. The VC for each VBR source was assumed to share only one link with CBR and ABR VCs. Thus VBR traffic entering a switch departed on the same link as the CBR and ABR traffic. The VBR traffic was then switched to a different link by the next switch. For clarification, this network and traffic topology is illustrated in Figure 1. Modelling of the background traffic was rudimentary and simply used an exponential distribution to represent the interarrival time of cells.

The reasons for including the VBR background traffic were twofold:

i. All the switches and links in the network topology had the same characteristics. Therefore, without the background traffic, more than one switch in the VC would simply increase the delay in CBR traffic by the line delay of each additional link. This would in turn have a deterministic effect on the delay variation experienced by CBR cells.

ii. It is unlikely that the only loading on a switch would be traffic with deterministic characteristics such as CBR and ABR traffic sources.

---

[2]It could be argued that a virtual source would be unable to distinguish ABR data cells from those generated by CBR, VBR or UBR sources. In this case all data cells would contribute to the rate at which RM cells were generated. However, it was found, particularly for high CBR and VBR transmission rates, that this approach produced an extremely high proportion of RM cells relative to the ABR loading, with RM cells contributing significantly to congestion.
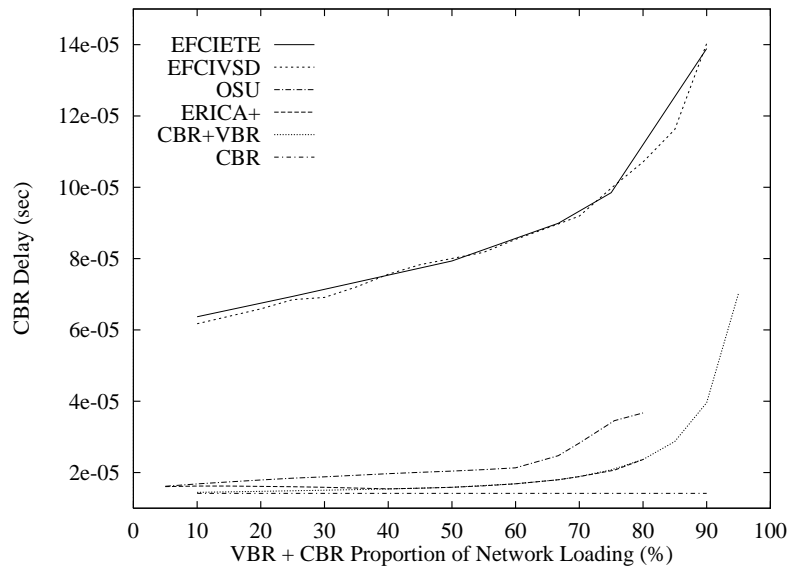
Figure 2: CBR Delay (seconds)

As the ABR and Management Control mechanisms were of interest it was decided to minimise other sources of variation. Hence, simulation results were generated for different CBR and mean VBR transmission rates (these rates were measured in terms of the proportion of link bandwidth used by each). For each simulation run, CBR and mean VBR transmission rates were the same, thus without ABR traffic present, network loading was divided equally between the the two sources. All results of Section 4 are plotted with respect to VBR and CBR loading.

Given the relationship between real and simulation time, it was necessary to maintain a relatively short simulation period. For each run a "warm up" interval was allowed during which time no statistics were collected. This was necessary to minimise any bias resulting from the initial period during which the ABR SES transmission rate was increasing up to the available bandwidth. A threshold was set on the queue size of each switch to determine when congestion occured.

4   Simulation Results

Results from the simulations conducted for the four models EFCIETE, EFCIVSD, OSU and ERICA+ are detailed in the following sections.

4.1   Mean CBR Delay

The mean delay in CBR traffic for different proportions of CBR plus VBR traffic is summarised in Figure 2. The delay statistics for the four simulation models were measured between ABR SES and DES. Also included in Figure 2 is the expected delay in CBR traffic without ABR or VBR traffic, and the observed delay with only CBR and VBR traffic.

The first point to note from Figure 2 is the plot for CBR and VBR traffic only. Even though the mean proportion of VBR and CBR traffic was below the link capacity, the VBR traffic has introduced a delay in the CBR traffic. Given the exponential distribution of a VBR source it is possible that its transmission rate will, on occasion, exceed the capacity of a link, resulting in delay for CBR traffic. However, when the combined CBR and VBR transmission rate is

Figure 3: Variation in CBR Delay (sec)

below the link capacity, the CBR traffic will experience the same delay as if no VBR traffic were present. Hence when the network carries additional VBR traffic, the mean CBR delay will always be equal to, or greater than, that when there is no VBR traffic. This disparity will obviously increase as the mean transmission rate of the VBR source increases. Therefore, any increase in CBR delay observed from the simulation models which include ABR traffic must be attributable in part to the VBR traffic.

The introduction of ABR traffic using the two EFCI models, EFCIETE and EFCIVSD, resulted in a significant increase in delay for the CBR traffic. It was observed that the delay induced by the ABR traffic increased as a function of the additional CBR and VBR traffic. This is indicated in Figure 2 by the divergence between the plot for CBR + VBR, and both EFCIETE and EFCIVSD plots.

Despite the significance of introducing the EFCI ABR traffic, neither RMCL mechanism appears to appreciably reduce the delay experienced by CBR traffic. Whilst shorter delays were observed for the virtual source and destination mechanism these are negligible in comparison to the overall delay incurred.

Unsurprisingly, the ER mechanisms appear to offer the optimum solution to congestion control of CBR traffic in the presence of an ABR source. The increase in CBR delay for the EFCI mechanisms was 9 - 17 times greater than the OSU scheme. The ERICA+ scheme appears to offer even better congestion control, with a delay in CBR traffic only slightly above that observed when no ABR traffic was present.

4.2   Variation in Mean CBR Delay

The mean delay variation of CBR traffic for different load values of CBR and VBR traffic is summarised in Figure 3. The variation was measured using the standard deviation of the mean delay. Also included in Figure 3 is the observed delay variation with only CBR and VBR traffic. With only CBR traffic present, no variation in delay would be expected.

The delay variation for all the plots of Figure 3 exhibit similar characteristics to those of Figure

Figure 4: ABR Transmission Rate (cells/sec)

2. The application of different RMCLs makes little difference to the variation in delay induced by the two EFCI mechanisms. Again the ER mechanisms appears to offer better congestion control, with variation in delay for CBR traffic reduced in comparison to the EFCI schemes. The delay variation for the EFCI mechanisms was seen to be some 9 - 17 times greater than that of the OSU scheme. Again the ERICA+ scheme only appeared to increase the variation in delay for low values of VBR and CBR traffic (VBR + CBR < 40%).

Also of note is the increase in delay variation for both EFCI plots. Whilst the mean delay of Figure 2 is a monotonically increasing function with VBR and CBR loading, the variation in delay has a minimum between approximately 10% and 40% network loading by VBR and CBR traffic.

4.3   Mean ABR rate

The mean ABR SES transmission rates for the models under different load values of CBR plus VBR traffic is summarised in Figure 4. Also included in Figure 4 is the total available bandwidth (AB) to ABR traffic and the target utilisation bandwidth (TUB) for the OSU and ERICA+ schemes.

Figure 4 illustrates that both EFCI mechanisms for ABR traffic make effective use of the remaining available bandwidth. There is very little difference in the mean observed transmissions rates for the two RMCL mechanisms, with only a slight improvement in utilisation attributable to segmentation when more bandwidth is available.

From Figure 4 it appears that the OSU scheme does not make as efficient use of available bandwidth as the EFCI mechanisms since it remains near to, or below, the lower bound of the target utilisation bandwidth (TUB). This is attributable to the OSU scheme's attempt to divide the available bandwidth evenly amongst active VCs. When VBR and CBR sources are below their fair share, adjusting the ABR to its fair share will result in under-loading of the link. The converse will occur when VBR and CBR are above their fair share. This can be observed in Figure 4 where VBR and CBR loads below 67% result in mean ABR rates below the TUB, and within the TUB for values greater then 67%.

Figure 5: ABR Transmission Rate Variation (cells/sec)

The most notable feature of Figure 4 is the dramatic decrease in transmission rate of the ERICA+ scheme. This is a result of the increasing proportion of VBR traffic. Recall that the VBR traffic has an exponential interarrival time with probability density function (pdf) $f(x) = \frac{1}{\theta} e^{\frac{-x}{\theta}}$, where θ is the mean interarrival time [13]. The ABR transmission rate is adjusted according to the most overloaded switch which in turn is relative to the switch with the highest VBR loading (i.e. the lowest remaining bandwidth available to ABR traffic). From the above pdf, the probability that the VBR transmission rate is above its mean is $1 - e^{-1}$ (i.e. the interarrival time is less than its mean). Now the probability that one VBR source is transmitting above its mean is 1 - *(the probability that all sources are transmitting less than or equal to their means)*. This gives a probability of $1 - e^{-n}$ (where n is the number of VBR sources). For this simulation where 4 VBR sources were used, the probability that one source is transmitting above its mean rate is approximately 0.98. Thus it can be seen that the remaining bandwidth available to ABR traffic will usually be less than that derived by simply removing the CBR and *mean* VBR rate.

As the mean VBR rate increases so the variance of the VBR rate will increase. Therefore, increasing the mean VBR rate will also increase the disparity between the mean rate and the VBR rate observed by the most overloaded switch. This results in the sharp drop in ABR transmission rate as VBR and CBR loading increases.

4.4 Variation in Mean ABR Rate

The variation in mean ABR SES transmission rates for the models under different load values of CBR plus VBR traffic is summarised in Figure 5. Again the standard deviation statistic was used as a measure of variation.

Both EFCI mechanisms demonstrate considerably less variation in transmission rate than the OSU schemes. For VBR and CBR loadings less than 5% and greater than 35%, the ERICA+ scheme appears to have the lowest variation in ABR. This is unsurprising given the mean transmission rate observed previously. Again the segmentation of the RMCL appears to make little difference to variations in the transmission rate of the EFCI mechanisms.

The large variation in delay observed for ABR traffic using the OSU scheme is an artifact of

employing the fair share allocation. When CBR and VBR loading is equal to its fair share (67%) the ABR rate will be operating close to its fair share and will only be influenced by changes in the VBR traffic (i.e. the minimum illustrated by Figure 5 at VBR and CBR loading of 67%). As the difference between CBR and VBR loading and their notional fair share increases, so the amount by which the ABR rate is changed to reach its fair share will also increase. Again this is illustrated in Figure 5 by the increasing delay variation either side of the minimum.

The variation in delay for the ERICA+ scheme peaks at approximately 15% network loading by VBR and CBR traffic. As the mean VBR transmission rate increases so the variation in this transmission rate will also increase. This, in turn, translates into more variation in the ABR transmission rate. As the mean VBR transmission rate increases it has also been shown that the remaining bandwidth available to ABR traffic decreases (see Section 4.3). Therefore the magnitude of the variation in the ABR transmission rate is decreased. It would appear that the interaction of these two factors produces a maximum at approximately 15% network loading.

5   Conclusions

Four different models for an ABR service class have been described. These were two ER feedback approaches (the OSU and ERICA+ schemes) using an end to end RMCL and two EFCI approaches, one using end to end resource management control, the other virtual source and destination. The resultant effect on CBR traffic and some aspects of ABR performance have been investigated, the results of which are summarised below.

- The two ER schemes appear to offer better congestion control for CBR traffic. Delay and variation in delay for the OSU scheme were some 9 - 17 times better than either EFCI mechanisms. The ERICA+ scheme appeared to offer good congestion control; however, this was due largely to the reduced transmission rate for higher VBR transmission rates.

- Both EFCI mechanisms appear to make better utilisation of the bandwidth available to ABR traffic than the ER mechanisms. In terms of the OSU scheme, this is the result of two factors. Firstly, the OSU scheme places a bound on the operating conditions of the network (i.e. the target utilisation). Secondly, the OSU scheme does not fully account for other traffic sources such as VBR and CBR traffic. The poor utilisation of bandwidth by the ERICA+ scheme is attributable to the presence of VBR sources with exponential interarrival times. Increasing the mean VBR rate dramatically reduced the transmission rate of the ABR traffic. However, the price to be paid for increased utilisation of available bandwidth by both EFCI mechanisms is significant increases in CBR traffic delay and delay variation.

- Segmenting the resource management control loop into virtual source and destination nodes had little effect on the performance of the EFCI mechanism. Only small improvements in CBR delay and variation, and transmission rates were observed.

It has been shown that the poor performance of the ERICA+ scheme in the presence of VBR traffic is a result of the positive skew in the pdf of the interarrival time. Reducing this positive skew would improve the ABR transmission rate of the ERICA+ scheme. Given the topology used in this simulation, the ABR transmission rate will also become worse as the number of switches with additional VBR traffic increases.

It should be noted that both EFCIETE and EFCIVSD models used the same parameters to control the ABR transmission rate for the different CBR and VBR transmission rates. Given that

different parameters will produce different performance results, optimising these parameters for different CBR and VBR values may result in improvement of the EFCI scheme over ER feedback.

The degradation in CBR delay in the presence of ABR traffic may be reduced by service class or VC queueing disciplines within switches. However, it is reasonable to assume that not all switches in a given VC path will implement these queueing disciplines. Hence, the results presented here represent a possible worst case scenario.

Given that the segmentation of the resource management control loop has produced little improvement in the performance of either CBR or ABR traffic this is clearly an area that warrants further investigation. In particular for segmentation at network edges rather than at individual switches within a network. Similarly, segmentation of the ER schemes could also be investigated.

References

[1] M. De Pryker. *Asynchronous Transfer Mode Solutions for Broadband ISDN*. Prentice Hall, 1995.

[2] F. Bonomi and K. Fendick. The rate-based flow control framework for the available bit rate ATM service. *IEEE Network*, 9(2):25–39, March/April 1995.

[3] R. Walthall and M. Clement. Simulation and analysis of the performance of EPRCA in a wide area ATM network consisting of both ER and EFCI switch mechanisms. In *4th International Conference on Telecommunications Systems, Modelling and Analysis*, Mar 1996.

[4] A. Charny. An algorithm for rate allocation in a packet-switching network with feedback. Master's thesis, Massachusetts Institute of Technology, May 1994.

[5] D. Clark A. Charny and R. Jain. Congestion control with explicit rate indication. In *ICC '95*, 1995.

[6] R. Jain. Congestion control and traffic management in ATM networks: Recent advances and a survey. *Submitted to Computer Networks and ISDN Systems*, February 1995.

[7] F. Lu. ATM congestion control. URL : http://www.cis.ohio-state.edu/~jain/cis788/ atm_cong/index.html, August 1995.

[8] S. Kalyanaraman R. Jain and R. Viswanathan. The OSU scheme for congestion avoidance in ATM networks using explicit rate indication. Technical Report OSU-CISRC-1/96-TR02, The Ohio State University, 1996.

[9] R. Jain S. Kalyanaraman R. Goyal S. Fahmy and F. Lu. ERICA+: Extensions to the ERICA switch algorithm. Technical Report ATM Forum/95-1346, ATM Forum, Oct 1995.

[10] Scientific and Engineering Sorftware, Inc, 4301 Westbank Drive, Bldg. A, Austin, TX78746, USA. *SES/workbench Creating Models*, release 3.1 edition.

[11] Scientific and Engineering Sorftware, Inc, 4301 Westbank Drive, Bldg. A, Austin, TX78746, USA. *SES/workbench Sim Language Reference*, release 3.1 edition.

[12] Scientific and Engineering Sorftware, Inc, 4301 Westbank Drive, Bldg. A, Austin, TX78746, USA. *SES/workbench Simulating Models*, release 3.1 edition.

[13] R. V. Hogg and E. A. Tanis. *Probability and Statistical Inference Third Edition*. Collier Macmillan, 1988.