Kluwer's Telecommunications Systems 23:1,2, p. 95-122, 2003

Measurement and Statistical Analysis of Asymmetric Multipoint Videoconference Traffic in IP Networks†

C. Skianis, K. Kontovasilis, A. Drigas, and M. Moatsos

National Center for Scientific Research "DEMOKRITOS"
P.O. Box 60228, GR-15310 AG. PARASKEVI ATTIKIS, GREECE
Emails: {skianis,kkont}@iit.demokritos.gr,
{adrigas,moatsos}@ariadne-t.gr

Abstract

The paper contributes results on the modeling of videoconferencing traffic over IP networks. The study is based on extensive data, gathered by tracing the actual packet exchange during a comprehensive set of realistic teleconferencing sessions over an asymmetric platform, in which commercial H.261-compliant terminal clients were communicating through a Multipoint Control Unit (MCU) at 'continuous presence' mode. Analysis of the data suggests that the video traffic from the client terminals can always be represented at the frame level as a stationary stochastic process with an autocorrelation function of exponentially fast decay and a marginal frame size distribution of approximately Gamma form. The video traffic from the MCU to the clients is again stationary and with exponentially decaying correlations, while the corresponding marginal frame-size PDF has the form of an appropriately weighted sum of Gamma components, the number of terms in the sum always being equal to the number of conferring terminals. The paper discusses methods for correctly matching the parameters of the modeling components to the data and for combining these components into complete traffic models that have been proposed in the literature.

Keywords: Videoconference systems, video bridging, multipoint control units (MCU), H.261 video coding, IP-traffic measurements, traffic modeling

1. Introduction

Videoconference is an increasingly popular network-based service. Its success is a result of reduced costs, continuous quality improvements and tailor made communication standards. Since videoconferencing relies on the exchange of bandwidth demanding video information, extensive deployment of this service calls for careful modeling of the associated network traffic, so that the appropriate amount of resources may be anticipated for by the network. Furthermore, and from a complimentary viewpoint, successful traffic modeling can provide valuable insights about the resulting network load, and these may be used towards an efficient

† Work partly funded by the Greek General Secretariat for Research and Technology (GSRT), through the research project ODISSEAS in the EPET II framework.

and more economical network usage, ultimately leading to lower communication costs and a more affordable service to the end-users.

Partly due to the reasons just mentioned, the problem of modeling video traffic, in general, and the video component of teleconferencing, in particular, has been extensively studied in the literature. A major surge of interest in the topic appeared with the advent of ATM, which offers the potential for efficient multiplexing of variable-bit-rate video streams. Relevant early studies [4,18,28,32] examined various characteristics of VBR video traffic, such as differences in successive frame sizes and cluster lengths [4], scene duration distributions [28,32], bit rate variations among scenes [32], and packet generation intervals (at various levels of video activity) [18]. Results from these and other works indicate that the histogram of frame sizes (often used as a proxy for the instantaneous bit-rate requirements) exhibits a (somewhat non-symmetric) bell-shape [15,18,28,21,22]. Furthermore, correlations in the video bit rate (again, usually assessed at the frame level) are found to decay exponentially [5,13,15,20,21], while other studies [22,23,24,25,26] observe a more complex phenomenon, in which the correlation decay is rapid for the initial lags, then continues at a lower rate.

The characteristics just mentioned provided a basis for the modeling of video traffic through relatively simple means, such as autoregressive (AR) models of various orders or Markov Chain (MC) based models. For the modeling of 'viewphone'-type video, reference [21] employs an AR(1) process with Gaussian residuals, which matches well to the bell-shaped density and exponential autocorrelations of the bit rate. The same modeling approach (again for teleconference-type video, coded according to various generic coding schemes) has been followed by [22], while [30] employs an AR(1) sequence with Gamma-distributed variables, in order to cope with the asymmetry in the rates histogram. As an alternative to AR-based models, [21] employs a MC-based approach, which corresponds to the homogeneous

superposition of a number of fictitious on/off Markovian sources, is more amenable to analytic treatment, and maintains the key features of exponentially decaying autocorrelations and bell-shaped density (actually binomial) for the bit rate. An extension of this MC model that can capture multiple video activity levels and scene changes (more frequent in TV-type content than teleconferencing) has been proposed as well [25].

Other works, in an attempt to track more closely the complexities observed in the autocorrelation decay, employ more elaborate means, such as autoregressive moving average (ARMA) models [12], Gaussian AR(1) models (sometimes more than one, in combination) with parameters modulated according to the transitions of a MC [23,26,33], Markov Renewal processes with sojourns fit to compound geometric distributions [20], or Markov Modulated Rate Processes (MMRP) with parameters fitted to the rates histogram [27]. From a different point of view, [7] promotes performance studies relating to video traffic on the basis of 'glitch' statistics, while [24] makes use of generic, coding independent, indices, as a means to characterize video traffic in general terms, not affected by coding idiosyncrasies.

Of particular relevance to our work is the approach in [15], where observations on a rather long sequence of frames from teleconferencing video provide evidence that the density of frame sizes has a Gamma shape (a fact that had also been noted by [18] and was subsequently employed in [30]) and that AR models of at least order two are required (and are sufficient) for a satisfactory match with the statistical characteristics of the sample. However, the same study observed that the AR(2) model cannot capture accurately the queueing performance of video, while a simple DAR(1) model, based on a discrete-time, discrete-state MC (with parameters matched to the frames size distribution and the exponential autocorrelation decay rate), is very well suited to this task. This modeling approach has been found appropriate for describing teleconferencing video coded by a variant of the H.261 standard [16] (being directly into the present paper's scope) and has been successfully employed in simulation-

based [5] and analytic/asymptotic [8] performance studies. A continuous-time variant of the DAR(1) model has been proposed [31], aiming at a model suitable for analytic queueing studies. In another direction, a four-parameter non-Markovian extension of DAR(1) has been used for capturing scene changes in television-type video content [11]; the study reports that in the absence of abrupt scene changes (as in typical teleconferencing) the results of the extension are in accordance with the simpler Markovian framework [15].

All the models reviewed up to now are short-range dependent, exhibiting rate correlations that decay at an exponential rate. Although there have been claims that video traffic possesses properties reminiscent of long-range dependence [1], teleconferencing video has been found to be only asymptotically self-similar [9], at a time-scale not affecting queueing. A study specifically focused on this issue [14] reports that indeed long range dependence does not affect the accuracy of queueing performance, as predicted by the DAR(1) model in [15], for all busy period lengths actually observed in practice.

The research results outlined in the previous paragraphs certainly constitute a valuable body of knowledge. However, many of these studies were undertaken at a time when the videoconference practice was still in early stages. As a consequence, many instances of traffic data that have been used for producing and/or validating results in earlier investigations relate to short video traces only a few seconds long (as in, e.g., [7,12,20,21,22]). Even when appropriately long traces exist, they are sometimes coded in nonstandard or simplified variants of standard coding schemes. For example, the long teleconferencing video trace of [15] has been employed for the validation of the DAR(1) model for H.261 video [16] through a open-loop (no rate control) coding variant that uses a fixed quantizer step.

Today, a large number of actual videoconference platforms exist, the majority of them operating over IP-based networking infrastructures and using practical implementations of the H.261 standard [3] for video coding, and it is important to know whether the models

established in the literature are appropriate for handling this contemporary setting in general. Furthermore, asymmetric teleconferencing platforms have been introduced, employing centralized management through multipoint control units (MCUs) [2,19],[1] with the aim of providing higher quality and better control over the sessions. When operating in the so-called 'continuous presence' mode, the MCU combines input from several client-terminals and sends the output back to these terminals. It is important to know if the more complex MCU traffic stream can be modeled by the known methods and, if not, to propose suitable models for this case too.

In addressing the context just stated, the research reported in this paper undertook extensive measurements of the IP traffic that was being generated during the course of actual videoconferencing sessions, hosted by a modern asymmetric platform that included client-terminals coordinated through an MCU. The experiments covered various cases with differing number of terminals (always conferring in continuous presence mode), and different quality-related MCU-parameters. The collected data were analyzed, providing results on the characterization of the video traffic originating both from the clients and the MCU.

These results suggest that, in all cases, the traffic from the clients can be represented—at the frame level—as a stationary stochastic process with an autocorrelation function of exponentially fast decay. Although the autocorrelation cannot be directly fitted to a single exponential term, careful choice of the decay rate allows the construction of a conservative (but asymptotically tight) exponential approximation. Furthermore, the distribution of sizes for the client-originated frames can always be satisfactorily approximated by a PDF of Gamma form, although an unconventional fitting of the model parameters, based on the shape of the histogram, is seen to be more appropriate than the usual moments matching approach.

[1] See also [29] for a comparison with other multimedia multiparty teleconference approaches.

These terminal-related findings, being the outcome of a comprehensive set of realistic experiments on an actual system using a practical implementation of H.261, can be seen as validation and reinforcement of earlier results ([5,13,15,20,21,22,23,24,25,26] for the autocorrelation, [15,16,18] for the rates distribution) and as additional justification for the modeling practices arising from them, in particular those based in the DAR(1) model [8,15,31].

The results for the traffic originating from the MCU (a topic not previously studied in the literature, to the best knowledge of the authors) are more complex, still however possessing a clearly identifiable structure. The traffic is again stationary and the autocorrelation function still exhibits exponentially fast decay and can be conservatively approximated by an appropriately chosen single exponential term. However, the distribution of the frame sizes is not conformant to a simple Gamma PDF. Instead, it is found that this distribution is closely approximated by a weighted sum of Gamma PDFs, the terms in this sum being always equal to the number of client-terminals in the corresponding videoconference session. The paper discusses methods for calculating the appropriate model parameters from the observed traffic data and proposes a simple adaptation of the DAR(1) model for usage with the MCU video traffic.

The rest of the paper is structured as follows: Section 2 discusses the videoconferencing platform employed for experimentation, describes the characteristics of the experiments' suite and outlines the procedures used for the measurement of data. Section 3 presents the results relating to the video traffic from the client terminals, and discusses methods for parameters matching and for usage of the resulting values in a traffic model. Along similar lines, Section 4 presents the results for the MCU traffic. Finally, Section 5 culminates with conclusions and pointers to further research.

2. The videoconference platform and experiment characteristics

The paper focuses on multipoint videoconference, i.e., an environment where the client-terminals communicate through a multipoint control unit (MCU) that coordinates the session. There are two main modes of operation: 'switched presence' [2] and 'continuous presence' [19]. In switched presence mode the MCU sends to all terminals the output from one participant, designated as "currently active", determined so either by human selection, or by an automatic detection of activity on the respective audio channel. In continuous presence the MCU combines the signal from up to four terminals and sends the resulting output to all the participants, who, in this way, become able of continuously viewing each other. The H.261 standard [3] allows the direct combination of four QCIF videos into a CIF video, without the need for decoding to the pixel level followed by a re-encoding, thus allowing an efficient handling of the continuous presence mode by the MCU. It is noted that in continuous presence the terminals operate asymmetrically [19], sending to the MCU QCIF pictures at some rate R and receiving CIF pictures at rate $4R$ (actually iR, for i participating terminals).

The experiments for the study reported herein were realized on a platform that consisted of standard personal computers running H.261-compliant commercial videoconferencing software (MS NetMeeting) and an MCU (a Cisco IP/VC 3510 unit), all networked over an IP-based LAN. The platform's topology is sketched in Fig. 1. Twelve different modes of videoconferencing, at various quality levels, were possible on this particular MCU; relevant characteristics of these modes are summarized in Table 1. Since the focus of the study was on continuous presence, Modes 3 and 4 (see Table 1) were used in the experiments, both employing a nominal frame rate of 15 frames/sec and corresponding to videoconferencing at high and low video bit rates, respectively. Two experiments were performed under Mode 3, one of them (referred to as 'Case 1' in the sequel) involving two conferring terminals, the

other (Case 2) involving four terminals. Another experiment (referred to as 'Case 3') employed Mode 4 with two terminals.

All experiments traced corresponding videoconferencing sessions of realistic content (reminiscent of a tele-working environment) and of sufficiently long duration (from 30 min to 2.5 hr, depending on the video bit rate used in each experiment). In each case, the IP packets exchanged in both directions between the terminals and the MCU were captured by traffic monitoring software. The collected data were further post-processed, in order to identify the packets carrying part of the same video frame (by tracing a common packet timestamp) and then calculate the sizes of individual successive frames. The final output of every experiment consisted of frame size sequences (their length ranging from about 18000 frames, to about 30000 frames, depending on the experiment), each corresponding to the frame sequence produced by either a terminal, or by the MCU. These sequences were subsequently used for the analysis reported in Sections 3 and 4.

Table 2 provides a synopsis of some experiments-related quantities. With reference to this table it may be observed that, in all cases, the bit rate from the MCU is a multiple of the terminals' bit rate, the multiplication factor being equal to the number of terminals. This is in accordance with the mechanics of the continuous presence videoconference and has the effect of maintaining an approximately equal frame rate in both directions between the MCU and the terminals. It may also be observed that the values of the bit rate achieved are in all cases much lower than the respective maximum specifications of the corresponding MCU modes (compare Tables 1 and 2), reflecting the fact that the content of the videoconference did not exhibit dramatic scene changes, frequent zooms, or other such effects.

Finally, it is noted that the columns of Table 2 for the 1st experiment (Case 1) include information for both conferring terminals and this information provides a rough evidence for the similarity of the respective traffic patterns. The issue is further elaborated upon in

Section 3. In general, close statistical proximity in the characteristics of all terminals within each session was observed throughout all experiments. For this reason, the information of Table 2 relating to the other experiments (Cases 2 and 3) refers to only one terminal for each case, as a representative of the video traffic from all terminals in the respective sessions.

3. Analysis of the video traffic from the terminals

Given the relevant published research on the topic, the following questions arose naturally (and their answer was pursued) during data analysis:

- Can the traffic be reasonably considered stationary?
- What is the form of the autocorrelation function? Is the associated decay exponentially fast? Is a steady decay rate maintained throughout, or does it vary in different lag-ranges?
- What is the form of the frame-size histogram? Can it be reasonably approximated by a Gamma (or normal, for non-pronounced asymmetry) density?
- Are the answers to the previous questions invariant among the terminals participating in the same videoconference scenario? More generally, are the general properties (not parameter values) characterizing the answers to the previous questions invariant among the traffic from terminals participating in different videoconference scenarios (i.e., different MCU modes)?

In brief, the answers to these questions, as supported by consistent evidence from the experiments' results (to be discussed in detail shortly), are as follows: the sequence of frame-sizes from a client terminal can be represented as a stationary stochastic process, with an autocorrelation function of exponentially fast decay (and finer properties to be commented upon later) and a marginal frame-size distribution of approximately Gamma form. These characteristics remain invariant for all MCU modes employed in the videoconference

experiments. Moreover, for the terminals participating in the same videoconferencing session the invariance in characteristics extends, beyond qualitative aspects, to the approximate equality of parameter values, so that the traffic from clients in the same session can be reasonably captured by a common model. The rest of the section provides individual details for each of these basic properties.

In checking for stationarity, each frame sequence corresponding to a terminal was split in ten windows and the empirical density function for the frame size was calculated from the samples in each window. These window-densities were found very much alike (see Fig. 2a,b), a property directly suggesting that the sequence is stationary. As a further test, the convolutions of these empirical densities were constructed for pairs of windows. Again, these convolution densities were almost identical across window combinations (Fig. 2c,d), reinforcing the previous result.

In the next stage of data analysis the autocorrelation function was directly sampled from the whole sequence of data. The relevant graphs appear in Fig. 3 (ignore the curves corresponding to model fitting results for the time being), on the basis of which two observations may be made. Firstly, the autocorrelation graphs for the two terminals in experiment Case 1 (subfig. a, b) are seen to match very closely and this fact supports the claim on the statistical identity of these terminals' characteristics. Secondly, the graphs exhibit a reduced decay rate beyond the initial lags, a behavior also noted in earlier studies [22,23,24,25,26]. In principle, this phenomenon can be captured by a weighted sum of two geometric terms, i.e.,

$$\rho_k = w\lambda_1^k + (1-w)\lambda_2^k, \qquad \text{with } |\lambda_2| < |\lambda_1| < 1. \tag{1}$$

The appeal of (1) stems from the fact that it corresponds to the general form of the autocorrelation function for AR(2) models (see, e.g., [6]), which have been observed [15] to

match well (and better than the simpler AR(1) models) some aspects of videoconferencing traffic.

In order to further check the suitability of (1), the relevant parameters were estimated through a least-squares fit to the autocorrelation samples for the first 500 lags. Numerical values for the results appear in Table 3, while the graphs of the fitted models are compared to the sample autocorrelations in Fig. 3. As it can be evidenced, the visual similarity of the sample autocorrelation curves for the two terminals in Case 1 is seconded by a close match to the corresponding parameter values, further reinforcing the claim on statistical identity. Furthermore, in all cases, the matched model (1) is able of capturing the long-term trends of the autocorrelation decay (see Fig. 3a-d), and this property is reflected most intensely in Fig. 3e for Case 2, where the model is compared against the samples over a wider range of lags (up to 5000). The success of (1) may be called upon for establishing that, in agreement with many previous studies, as reviewed in the introduction, (a) the autocorrelation decay is indeed exponentially fast and (b) that the long-term decay rate is significantly slower than the one in the first few lags.

As it can be seen from Fig. 3c-e, the sample autocorrelation function for the terminals in Cases 2 and 3 is non-monotonic, but this fact is not reflected in the respective values for $\lambda_{1,2}$, which are both positive. Fluctuations of this kind can also be obtained when the lambda-parameters in (1) are complex (necessarily conjugate, since ρ_k is real), in which case the autocorrelation function can equivalently be expressed as a geometrically dumped sinusoid [6] of the form

$$\rho_k = \lambda^k \cos(\theta k + \psi)/\cos\psi$$

for appropriate parameter values. These values were also estimated from a least squares fit to the autocorrelation samples (again for the first 500 lags) of Cases 2 and 3 and the results for

Case 2 are displayed in Table 3 and Fig. 3f, the figure displaying the graph of the fitted model plotted against the samples for the first 100 lags. The results suggest that the dumped sinusoid has a decay rate much faster than the actual one and is thus worse than the previous model. This failure promotes the conjecture that the autocorrelation function in cases 2 and 3 is actually closer to a compound geometric containing more than two terms and that the lower order terms are those giving rise to dumped sinusoids that account for the non-monotonicity.

However, there is no point in further pursuing the issue towards more complex models, as the property of the autocorrelation that is most important for queueing is the long-term decay rate and this is very well captured (in all cases) by the value of parameter λ_1 in the model (1). In fact, by adopting a single exponential model of the form $\hat{\rho}_k = \lambda_1^k$, it is guaranteed that, for all lags, $\rho_k \leq \hat{\rho}_k$, and the simpler approximation provides a conservative (from the queueing point of view) but asymptotically tight model, since stronger positive correlation results in more pronounced buffer occupancies, hence more probable overflows and longer queueing delays.

When constructing the single exponential approximation, care must be exercised during the choice of the parameter. It is reminded that this should be set equal to the value of the parameter λ_1, as determined by the least squares fit of the data to the model (1) and is *not* equal to the sample autocorrelation at lag-1. The validity of this approach is further highlighted by observing that the values for λ_1 in Table 3 are in all cases around 0.99, a value close to the lag-1 autocorrelation used in previous studies, e.g., [15,20].

We now turn to the discussion of the frame-size distribution. The relevant (smoothed) histograms, as obtained from the data, are displayed in figures 4a-b (for Case 1) and 5a-b (for Cases 2 and 3). All empirical densities feature a bell-like shape and a slight asymmetry around their maximum, thus being reminiscent of a negative binomial density, of the form

$$f_k = \binom{k+p-1}{k} a^p (1-a)^k, \quad 0 < a < 1,\ p > 0, \quad k = 0,1,\mathrm{K}\ . \tag{2}$$

Instead of directly dealing with the discrete density (2), in the following we work with its continuous counterpart, namely the Gamma density

$$f(x) = \frac{(x/\mu)^{p-1}}{\mu\Gamma(p)} e^{-x/\mu}, \quad \mu, p > 0,\ x \geq 0, \quad \Gamma(p) = \int_0^\infty t^{p-1} e^{-t}\, dt\ . \tag{3}$$

These distributions share all their distinctive characteristics and it is readily possible to derive the parameters of one of them from those of the other, by equating the respective expressions for the first two moments. An important characteristic of the Gamma distribution is that, for integral p, it corresponds to the distribution of the sum of p iid exponentially distributed random variables, each with parameter $1/\mu$. Another related feature is that the n-fold convolution of a Gamma density, of parameters μ and p, is again a Gamma density, of parameters μ and np, and this holds even for non-integral p. This last property was employed for checking that the histograms for the two terminals of Case 1, which are somewhat 'distorted' in the vicinity of their expected maxima (see Fig. 4a-b), should indeed by modeled by densities of Gamma form. Specifically, the convolution of each empirical density with itself was taken and these convolutions were clearly possessing a Gamma-like shape. (See Fig. 6 for the convolution corresponding to Terminal 1 of Case 1.)

For the purpose of fitting the Gamma model to the histograms, three different parameter-matching methods were tried. The first one was the usual moments approach, which makes use of the fact that the Gamma distribution has mean $p\mu$ and variance $p\mu^2$. By equating to the sample mean and variance, m and v respectively, one obtains $\mu = v/m$ and $p = m^2/v$. The second method (called LVMAX) relates the histogram's peak to the location at which the Gamma density achieves its maximum and to the value of this maximum. Indeed,

when $p > 1$, by employing (3) and requiring $f'(x) = 0$, it follows that the Gamma density has a unique maximum at

$$x^* = (p-1)\mu. \tag{4}$$

Through (4) and (3) one may obtain a relation that connects the location and value of the maximum to the shape parameter *p*, namely

$$x^* f(x^*) = \frac{(p-1)^p e^{-(p-1)}}{\Gamma(p)}. \tag{5}$$

Furthermore, when *p* is integral, $\Gamma(p) = (p-1)!$ and, by Stirling's formula (see, e.g., [10]), $\Gamma(p) \sim \sqrt{2\pi}(p-1)^{p-1/2} e^{-(p-1)}$, which, combined with (5) gives

$$x^* f(x^*) \approx \sqrt{\frac{p-1}{2\pi}}. \tag{6}$$

Although this last equation has been derived on the assumption of integral and large *p*, it provides a very good approximation for all values of interest (with less than 8.5% relative error for $p \geq 2$ and less than 1% for $p \geq 10$).

On the basis of (6) and (4), measurements of both the location and value of the histogram's peak, $\hat{x}$ and $\hat{f}$ respectively, determine the Gamma parameters as

$$\mu = \frac{1}{2\pi \hat{x} \hat{f}^2} \qquad \text{and} \qquad p = 2\pi \hat{x}^2 \hat{f}^2 + 1.$$

Note that LVMAX does not preserve, neither the sample mean, nor the sample variance.

Finally, the third method (called C-LVMAX) exploits the convolution property of Gamma densities mentioned earlier, through an application of LVMAX[2] to the self-convolution of the

[2] There's no point in employing convolution with the moments matching method, as it gives results identical to those produced by an application of the method on the original empirical density.

histogram, followed by a division of the resulting *p*-value by 2, in order to recover the parameters corresponding to the original setting.

Numerical results from the application of the parameters-matching methods on the data appear in Table 4, while the graphs for the corresponding Gamma densities are compared against the histograms in figures 4a-b and 5a-b. (In these figures, the curve for the LVMAX fit is the one, which, by design of the method, tracks exactly the histogram's peak.) Somewhat surprisingly, the best results were obtained for the two terminals of Case 1, i.e., those corresponding to histograms of a somewhat 'irregular' shape. Indeed, the contents of Table 3 show that, for Case 1, the three sets of parameters values are close to each other and that the mean and variance corresponding to the Gamma densities resulting from LVMAX and C-LVMAX are close to the sample counterparts. Furthermore, the parameter values for the two terminals are close enough to allow a single model for the description of both histograms and this fact provides another bit of evidence towards the statistical identity of the traffic profiles from all terminals in a videoconferencing session. Note, however, that this invariance is not carried over to the histograms corresponding to different sessions over the same MCU mode but with a differing number of terminals, as the histograms of Case 1 (MCU mode 3 with two terminals) are significantly different from the histogram of Case 2 (MCU mode 3 with four terminals) and the same holds for the corresponding sample means and variances.

Compared to Case 1, the results for the other two experiments present somewhat greater diversity. The respective histograms are very peaked (see Fig. 5a-b) but also contain a non-negligible probability mass corresponding to smaller frame sizes. Due to the second characteristic the Gamma density resulting from the moments fit is more spread out (and matches by design the sample mean and variance), while, due to the peakedness, LVMAX tends to produce a higher value of *p* and this leads to a variance considerably different from the sample value. C-LVMAX, by virtue of the convolution's "smoothing" effect, produces

intermediate results. In general, the variances resulting from the three parameter-matching approaches exhibit greater diversity than the corresponding mean values.

However, the importance of matching the sample variance should not be overemphasized. More important for queueing is the tail of the distribution function, since this determines the probabilities with which large frames are produced and correspondingly high traffic rates are attained, leading to congestion. A particularly effective way of comparing the tails of the matched Gamma models to the histogram is through quantile-vs-quantile plots (q-q plots) of the respective (cumulative) distributions. When there is a perfect match, the q-q plot runs across a straight line bisecting the angle between the coordinate axes. A deviation towards one of the axes indicates that, in the relevant region, quantiles of the distribution corresponding to that axis are larger.

The q-q plots for the Gamma models obtained from the three matching methods are displayed in figures 4c-e,f-h (for the two terminals of Case 1) and 5c-e,f-h (for Cases 2 and 3). As expected, the uniformity of results previously observed for Case 1 is also reflected in the q-q plots, which indicate that almost all matching methods track very closely the quantiles of the respective histogram. The sole exception is presented by the LVMAX fit on Terminal 2, where the anomaly around the expected maximum prevents its accurate detection, thus reducing the precision of the method. Even in this case, the convolution preprocessing of C-LVMAX corrects the estimation, giving the most precise results for this set of data.

For the other two cases, the q-q plots on the moments fit reflect clearly the effect of the small probability mass borne by smaller-than-mean frame sizes, by indicating dominance of the fit for low quantiles and respective lagging behind in the important higher ones governing the distribution’s tail. The Gamma models produced by LVMAX and C-LVMAX are better in this respect, tracking closely the histogram in all quantile ranges and maintaining tail dominance.

Overall, the C-LVMAX method performed best in all cases, by managing to follow closely the histogram in all quantiles, while also providing (tight) tail dominance and a close upper bound to the sample mean (see Table 4). Consequently, it is the method proposed for matching the Gamma parameters. Although it has the drawback of requiring full histogram information (instead of just the sample moments) the experiments suggest that the frame size distribution depends primarily on the characteristics of the MCU mode and the number of conferring terminals, remaining invariant for all (typical) terminal-originated traffic under these conditions. Thus, it appears possible to obtain the appropriate Gamma model from off-line measurements and the drawback is no longer relevant.

As a final comment on the parameters, it is noted that the values of p in Table 4 are high (always greater than 17) and this fully justifies the approximation (6) used by (C-)LVMAX. Moreover, these high values indicate that the asymmetry in the Gamma densities is weak, so that an equivalent normal density could also be used instead, a choice made in several previous studies [21,22,23,26,33].

We close this section by exploiting its main conclusions towards a full model for the terminal-originated traffic. Since the terminals exhibit correlations that can be bounded by a simple geometric function and Gamma distributed frame sizes, the approach of [15], namely that of modeling videoconferencing traffic through a DAR(1) model, is directly applicable. Furthermore, its success in predicting the queueing performance of videoconference traffic [5,14,15], including traffic resulting from video coded according to a variant of H.261 [16], makes it a natural choice. This model, which originally appeared in [17], produces a sequence of frame sizes according to the transitions of a discrete-time discrete-state MC, of the form

$$P = \rho I + (1-\rho)Q, \tag{7}$$

where I is the identity matrix, ρ is the autocorrelation coefficient at lag-1 and Q is a rank-one stochastic matrix with all rows equal to the probabilities resulting from the (suitably truncated at some maximum size) negative binomial density (2) corresponding to the Gamma fit for the frame size distribution. The DAR(1) has an exponentially decaying autocorrelation function equal to ρ^k and a marginal frame-size density with probability masses equal to the elements of the common rows in Q. In light of the autocorrelation properties highlighted by the experiments' results, it follows that ρ should be chosen equal to the parameter λ_1 in the model (1), as determined from the least squares fit discussed earlier in this section, *not* equal to the sample autocorrelation at lag-1, if the long term trends in the correlation decay are to be preserved. Similarly, the elements for the rows of Q should be determined through the Gamma fit produced by the C-LVMAX method. These two details are important for ensuring that the resulting model will provide a conservative (but also closely accurate) traffic characterization during queueing studies. Note that, once the discrete DAR(1) model is constructed, the discrete-state continuous-time variant proposed by [31] can be readily applied in the standard way, towards obtaining a model more amenable to analytic treatment.

4. Analysis of the video traffic from the MCU

The MCU-related data were analyzed by following a methodology similar to the one used for the terminals' traffic. Again, it was found that the sequence of frame sizes is characterized by a number of properties, which are invariant to the particular MCU mode employed. More specifically, the data produced by the experiments indicate stationarity (checked for by the techniques discussed in Section 3) and correlations with exponentially fast decay.

The sample autocorrelation graphs for the three experiments appear in Fig. 8. As previously, model (1) was fit to these samples, using least squares estimation, and the relevant numerical values are displayed in Table 3. Again, the fit follows satisfactorily the long-term trends of the

decay. This time, the λ_2 parameters for the rapidly decaying term are negative, accounting somewhat for the non-monotonic behavior. The fit to a geometrically dumped sinusoid was also tried and was found again to match well the samples for the first few lags, but to have an overly fast decay rate (see Fig. 8d, corresponding to Case 2). On account of the discussion in Section 3, an upper bounding approximation for the autocorrelation function through a single geometric term with parameter equal to λ_1 seems appropriate for the MCU traffic too, under all operation modes.

The feature that distinguishes most the MCU data from the terminal-related ones is the form of the marginal frame-size distribution. The relevant histograms are displayed in Fig. 9 (not including Case 3, which is, in all qualitative characteristics, analogous to Case 1 and is thus omitted). As it can be verified, the empirical densities can no longer be approximated by a simple Gamma form; instead, their shape resembles a weighted sum of Gamma components, i.e.,

$$f(x) = \sum_{i=1}^{k} a_i f_i(x), \qquad a_i > 0, \quad \sum_{i=1}^{k} a_i = 1, \tag{8}$$

where each $f_i(\cdot)$ is of the form (3). Furthermore, the number k of terms in the sum is always equal to the number of terminals participating in the teleconferencing.

In order to fit model (8) to the samples, two variants of a parameter-matching method were developed and applied to the data. The first variant relies on the assumption that the individual peaks x_i^* of the various Gamma components in the sum (8) are sufficiently far apart, with the implication that in the neighborhood of the i-th component's peak the contribution of the other terms will be negligible, so that effectively $f(x) \approx a_i f_i(x)$. On this assumption, the local maxima of f will approximately occur at x_i^*, $i = 1, \mathrm{K}, k$, and each maximum will attain a value approximately equal to $a_i f_i(x_i^*)$. In light of these properties, the

method uses the empirical density values around the local maxima for estimating the corresponding Gamma scale and shape parameters, μ_i and p_i respectively. The estimation relies on the observation that, by (3) and (4),

$$\log\frac{f_i(x)}{f_i(x_i^*)} = (p_i - 1)\left(\log\frac{x}{x_i^*} - \frac{x}{x_i^*} + 1\right), \qquad \text{for any } i, x. \tag{9}$$

Let now $\hat{x}_i$ and $\hat{f}_i$ stand for the location and value of the empirical density's i-th local maximum, and further let (x_n, f_n), $n = 1,\mathrm{K}, N$ be N other points on the density's graph, in the vicinity of the same local maximum. Transform the data by setting

$$Y_{n,i} = \log(f_n / \hat{f}_i) \qquad \text{and} \qquad X_{n,i} = \log(x_n/\hat{x}_i) - x_n/\hat{x}_i + 1. \tag{10}$$

Then, due to the "good separation" of the peaks and (9),

$$Y_{n,i} \approx \log\frac{f_i(x_n)}{f_i(x_i^*)} = (p_i - 1)X_{n,i}$$

and the transformed samples around the maximum are approximately linearly related. Through a least squares fit to these samples, one may directly obtain the shape parameter for the i-th Gamma component as

$$p_i = \frac{\sum_{n=1}^{N} Y_{n,i} X_{n,i}}{\sum_{n=1}^{N} X_{n,i}^2} + 1. \tag{11}$$

Once the shape parameter p_i has been computed the corresponding scale parameter may be retrieved through (4), as $\mu_i = \hat{x}_i/(p_i - 1)$.

Repetition of this procedure on samples around each local maximum provides estimates for the parameters of all Gamma components in the superposition (8). The final step of the method relates to the computation of the weight-factors a_i. In principle, these may be chosen

so as to match the first $k-1$ sample central moments m_i, $i = 1,\text{K}\,,k-1$. Indeed, since the *j*-th central moment of a Gamma distribution is equal to $p(p+1)\Lambda\ (p+j-1)\mu^j$, the matching provides $k-1$ linear equations

$$m_j = \sum_{i=1}^{k} a_i\, p_i(p_i+1)\Lambda\ (p_i+j-1)\mu_i^j, \qquad j = 1,\text{K}\,,k-1,$$

for the weights, which, together with the normalization condition in (8), provide a full system on the unknowns a_i. However, this approach does not work well, either yielding a nearly singular system of equations, or leading to negative weight values. For this reason, it is better to follow a more direct approach, based on the fact that the value of the histogram's *i*-th peak is $\hat{f}_i \approx a_i\, f_i(x_i^*)$. Then, by (4) and (5), the value of the weight should be equal to

$$a_i = G\,\hat{f}_i \mu_i \Gamma(p_i)\left(\frac{p_i-1}{e}\right)^{-(p_i-1)}, \tag{12}$$

where *G* is a proportionality constant enforcing the normalization condition in (8). This way of computing the weights, besides being more stable, is also easier to compute. However, it doesn't guaranty the preservation of any sample moments.

As it will be discussed shortly, the matching method just discussed works very well, even for the distribution corresponding to Case 1, where the two peaks are not very far apart (see Fig. 9a). Nevertheless, in order to ensure that the matching procedure will remain robust even in future applications, with empirical distributions featuring peaks even more closely spaced, another variant of the method, involving an iterative refinement, was also developed. The iterations proceed by initializing the various parameters of the density in (8) by the method just described, thus producing $a_i^{(0)}, p_i^{(0)}, \mu_i^{(0)}$, for $i = 1,\text{K}\,,k$. Then, at the *l*-th iteration:

- For each component $i = 1,\text{K}\,,k$:

- The effect of the other Gamma components, as currently estimated, is "removed" from the global empirical density $f(x)$, by forming the residual $f_i^{(l)}(x) = f(x) - \sum_{j \neq i} a_i^{(l-1)} f_i(x; p_i^{(l-1)}, \mu_i^{(l-1)})$, taken to stand for a multiple of the *i*-th Gamma component.

- The location $\hat{x}_i^{(l)}$ and value $\hat{f}_i^{(l)}$ of the (now global) maximum of $f_i^{(l)}(x)$ are determined and further samples on the graph of this function are gathered, in the neighborhood of the peak.

- The samples are transformed according to (10) and are used in the least-squares estimator (11) for the determination of $p_i^{(l)}$. The new estimate for the scale parameter is obtained as $\mu_i^{(l)} = \hat{x}^{(l)}_i \big/ (p_i^{(l)} - 1)$.

▪ The new set of weight factors $a_i^{(l)}$ is obtained through (12).

Both variants of the method were applied to the data from Cases 1 and 2 and Table 5 contains the parameter values after the initialization step (first variant) and after subsequent iterations. It may be observed that the iterations converge quite rapidly even in the less favorable Case 1 (for Case 2 convergence is reached at the first step). Note also that, although the weight fitting (12) is not designed to match any sample moments, the resulting model (8) maintains a mean and variance that are close to the sample counterparts. It should be observed that the values for the *p*-parameters are high (and quite high for Case 2, due to the intense peakedness exhibited by the corresponding histogram). This fact suggests that the histograms could have been fitted to a weighted sum of normal components, with almost equivalent results.

Further evidence for the success of the matching procedure is provided by Fig 9, which contains graphs of the resulting Gamma superpositions and compares them against the

corresponding histograms. As it can be observed, the fitted models follow closely the empirical densities. The same figure also includes q-q plots for the respective distributions, which also indicate a satisfactory behavior. The plot for Case 2 tracks almost perfectly the histogram's quantiles, while that of Case 1 is less tight.

In parallel to the concluding remarks of Section 3, it is noted that the characteristics of the MCU data suggest the potential of using a modified DAR(1) model to describe the relevant traffic. The ρ parameter in (7) can be matched as for the terminals, while the probabilities in the rows of *Q* should be set equal to a weighted sum of negative binomial probabilities, using the weights a_i, $i = 1,\text{K},k$, from (8). In this arrangement $Q = \sum_{i=1}^{k} a_i Q_i$, where each Q_i employs only the negative binomial density equivalent to the *i*-th Gamma component. The resulting DAR(1) model has the same complexity with the models describing terminal-originated traffic.

5. Conclusions

This paper reported on an experimental study of H.261-encoded video traffic during the course of realistic videoconferencing sessions, hosted by a modern asymmetric platform that included client-terminals communicating over an IP-network and being coordinated through an MCU, which operated in continuous presence mode. The experiments covered various cases with differing number of terminals and different quality-related MCU-parameters. Analysis of the collected data established general results about the video traffic originating both from the clients and the MCU.

Traffic from the terminals was seen to be stationary and to possess an autocorrelation function decaying exponentially fast. Although the correlations are more complex than a simple geometric model, careful choice of the decay rate allows the construction of a conservative (with respect to queueing behavior) but asymptotically tight such approximation. Another

property of the terminals' traffic is that the distribution of frame sizes may be satisfactorily approximated by a PDF of Gamma form, although an unconventional fitting of the model parameters, according to the C-LVMAX method, was seen to be more appropriate (again, with respect to queueing) than the usual moments matching approach. These observations are useful for the appropriate parameterization of DAR(1) models, for the purpose of describing the traffic from videoconferencing clients in network performance studies.

In general, the terminal-related findings of the paper, being the outcome of a comprehensive set of realistic experiments on an actual system using a practical implementation of H.261, can be regarded as validation, reinforcement and extension of earlier results and as additional justification for the modeling practices arising from them, in particular those based in the DAR(1) model.

The MCU traffic was also observed stationary and possessing an autocorrelation function that decays exponentially fast and can be conservatively approximated by an appropriately chosen single exponential term. However, the distribution of the frame sizes is not conformant to a simple Gamma PDF. Instead, it was found that this distribution is closely approximated by a weighted sum of Gamma PDFs, the terms in this sum always being equal to the number of client-terminals in the corresponding videoconference session. The paper discussed methods for calculating the appropriate model parameters from the observed data and proposed a simple adaptation of the DAR(1) model for usage with the MCU video traffic.

The more complex traffic pattern of the MCU is due to the fact that the latter combines on output several video streams. The combination pattern would be different if the MCU were operating in a combined continuous/switched presence mode (i.e., concurrently presenting the output from up to four selected terminals and switching to another one whenever it becomes active). Study of the traffic produced by the MCU under this more complex mode will be the subject of future research.

References

[1] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger, "Long-range dependence in variable-bit-rate video traffic", *IEEE Trans. Commun.*, 43 :1566–1579, 1995.

[2] CCITT Study Group XV – Report R93, Recommendation H.231, *Multipoint control units for audiovisual systems using digital channels up to 2 Mbit/s*, May 1992.

[3] CCITT Study Group XV – Report R95, Recommendation H.261, *Video codec for audiovisual services at p×64 kbit/s*, May 1992.

[4] H. S. Chin, J. W. Goodge, R. Griffiths, and D. J. Parish, "Statistics of video signals for viewphone-type pictures", *IEEE JSAC*, 7(5):826–832, 1989.

[5] D. M. Cohen and D. P. Heyman, "Performance modeling of video teleconferencing in ATM networks", *IEEE Trans. Circuits Syst. Video Technol.*, 3(6):408–422, 1993.

[6] D. R. Cox and H. D. Miller, *The Theory of Stochastic Processes*, Chapman & Hall, London, 1965.

[7] I. Dalgic and F. A. Tobagi, "Performance evaluation of ATM networks carrying constant and variable bit-rate video traffic", *IEEE JSAC*, 15(6):1115–1131, 1997.

[8] A. Elwalid, D. Heyman, T. V. Lakshman, D. Mitra, and A. Weiss, "Fundamental bounds and approximations for ATM multiplexers with applications to video teleconferencing", *IEEE JSAC*, 13(6):1004–1016, 1995.

[9] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing analysis with long-range dependent packet traffic", *IEEE/ACM Trans. Networking*, 4(2):209–223, 1996.

[10] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. 1, 3rd Ed., Wiley, 1968.

[11] M. R. Frater, J. F. Arnold, and P. Tan, "A new statistical model for traffic generated by VBR coders for television on the broadband ISDN", *IEEE Trans. Circuits Syst. Video Technol.*, 4(6):521–526, 1994.

[12] R. Grunenfelder, J. P. Cosmas, S. Manthorpe, and A. Odinma-Okafor, Characterization of video codecs as autoregressive moving average processes and related queueing system performance", *IEEE JSAC*, 9(3):284–293, 1991.

[13] B. G. Haskell, "Buffer and channel sharing by several interframe picturephone coders", *Bell System Tech. J.*, 51(1):261–289, 1972.

[14] D. P. Heyman and T. V. Lakshman, "What are the implications of long-range dependence for VBR-video traffic engineering?", *IEEE/ACM Trans. Networking*, 4(3):301–317, 1996.

[15] D. P. Heyman, A. Tabatabai, and T. V. Lakshman, "Statistical analysis and simulation study of video teleconference traffic in ATM networks", *IEEE Trans. Circuits Syst. Video Technol.*, 2(1):49–59, 1992.

[16] D. Heyman, A. Tabatabai, T. V. Lakshman, and H. Heeke, "Modeling teleconference traffic from VBR coders", in Proc. *ICC'94*, pp1744–1748, 1994.

[17] P. A. Jacobs and P. A. W. Lewis, "Time series generated by mixtures", *J. Time Ser. Anal.*, 4(1):19–36, 1983.

[18] R. Kishimoto, Y. Ogata, and F. Inumaru, "Generation interval distribution characteristics of packetized variable rate video coding data streams in an ATM network", *IEEE JSAC*, 7(5):833–841, 1989.

[19] S.-M. Lei, T.-C. Chen, and M.-T. Sun, "Video bridging based on H.261 standard", *IEEE Trans. Circuits Syst. Video Technol.*, 4(4):425–437, 1994.

[20] D. M. Lucantoni, M. F. Neuts, and A. R. Reibman, "Methods for performance evaluation of VBR video traffic models", *IEEE/ACM Trans. Networking*, 2(2):176–180, 1994.

[21] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J. D. Robbins, "Performance models of statistical multiplexing in packet video communications", *IEEE Trans. Commun.*, 36(7) :834–843, 1988.

[22] M. Nomura, T. Fujii, and N. Ohta, "Basic characteristics of variable rate video coding in ATM environment", *IEEE JSAC*, 7(5):752–760, 1989.

[23] G. Ramamurthy and B. Sengupta, "Modeling and analysis of a variable bit rate video multiplexer", in Proc. *7th Intl. Teletraffic Congress Seminar*, Morristown, NJ, 1990.

[24] R. M. Rodriguez-Dagnino, M. R. K. Khansari, and A. Leon-Garcia, "Prediction of bit rate sequences of encoded video signals", *IEEE JSAC*, 9(3):305–314, 1991.

[25] P. Sen, B. Maglaris, N.-E. Rikli, and D. Anastassiou, "Models for packet switching of variable-bit-rate video sources", *IEEE JSAC*, 7(5):865–869, 1989.

[26] C. Shim, I. Ryoo, J. Lee, and S. Lee, "Modeling and call admission control algorithm of variable bit rate video in ATM Networks", *IEEE JSAC*, 12(2):332–344, 1994.

[27] P. Skelly, M. Schwartz, and S. Dixit, "A histogram-based model for video traffic behavior in an ATM multiplexer", *IEEE/ACM Trans. Networking*, 1(4):446–459, 1993.

[28] W. Verbiest, l. Pinnoo, and B. Voeten, "The impact of the ATM concept on video coding", *IEEE JSAC*, 6(9):1623–1632, 1988.

[29] M. H. Willebeek-LeMair and Z.-Y. Shae, "Videoconferencing over packet-based networks", *IEEE JSAC*, 15(6):1101–1114, 1997.

[30] S. Xu and Z. Huang, "A Gamma autoregressive video model on ATM networks", *IEEE Trans. Circuits Syst. Video Technol.*, 8(2):138–142, 1998.

[31] S. Xu, Z. Huang, and Y. Yao, "An analytically tractable model for video conference traffic", *IEEE Trans. Circuits Syst. Video Technol.*, 10(1):63–67, 2000.

[32] Y. Yasude, H. Yasuda, N. Ohta, and F. Kishino, "Packet video transmission through ATM networks", in Proc. *IEEE GLOBECOM 1989*, pp.876–880.

[33] F. Yegenoglu, B. Jabbari, and Y.-Q. Zhang, "Motion-classified autoregressive modeling of variable bit rate video", *IEEE Trans. Circuits Syst. Video Technol.*, 3(1):42–53, 1993.

Table 1. Videoconferencing modes on the Cisco IP/VC 3510 MCU

Mode	Description	Frame Rate	Max. Video Bit Rate for terms/MCU (Kbit/s)	Max. # Parties	Output Picture Format
1	Low Quality	7.5	110/400	5	CIF
2	Low Quality	7.5	110	5	QCIF
3	High Quality	15	320/1280	4	CIF
4	Low Quality	15	100/400	5	CIF
5	Super Quality	30	720	5	QCIF
6	Medium Quality	30	320	5	QCIF
7	Voice Only	—	—	4	—
8	Super Quality	30	720	4	QCIF
9	Voice Only	—	—	4	—
10	Dynamic Rate	30	~	5	QCIF
11	Full Room	15	110	14	QCIF
12	Good Quality	7.5	220	3	QCIF

Table 2. Summary of relevant quantities for each experiment

Case No	**1**			**2**		**3**	
Term/MCU	Term (1)	Term (2)	MCU	Term	MCU	Term	MCU
MCU Mode	3			3		4	
Parties	2			4		2	
Experiment Duration (sec)	3386			1888		9293	
Total # Video Frames	26334	27016	28652	18543	18619	28632	31661
Video Bit Rate (Kbits/s)	204	209	408	208	812	63	127
Frame rate (fps)	8.2	8.0	8.5	9.8	9.9	3.1	3.4
Average Frame Size (Bytes)	2988	3062	5779	2581	10081	2353	4520
Variance of Frame Size (Bytes^2)	527161	525127	2927935	165174	9734163	301882	1423909

Table 3. Model parameters fitted to sample autocorrelation graphs for traffic from terminals and MCU

	Experiment	Parameters for compound exponential fit $\rho_k = w\lambda_1^k + (1-w)\lambda_2^k$			Parameters for exponentially dumped sinusoid $\rho_k = \lambda^k \cos(\theta k + \psi)/\cos\psi$		
		w	λ_1	λ_2	λ	θ	ψ
Terminals	Case 1 Term 1	0.5857	0.9980	0.5115			
	Case 1 Term 2	0.5190	0.9985	0.6261			
	Case 2	0.1579	0.9944	0.6009			
					0.5089	0.0068	0.0421
	Case 3	0.2103	0.9936	0.4522			
MCU	Case 1 MCU	0.3319	0.9985	-0.0352			
	Case 2 MCU	0.0462	0.9989	-0.2124			
					-0.6364	0.8926	3.3086
	Case 3 MCU	0.1795	0.9968	-0.0058			

Table 4. Gamma parameters for the various fitting methods applied to terminals' data

Experiment	Fitting Method	Gamma Param.		Moments	
		p	μ	mean	variance
Case 1 Term 1	Sample			2988.88	527161.50
	MOM	16.95	176.37	2988.88	527161.50
	LVMAX	23.07	133.60	3081.60	411710.61
	C-LVMAX	17.31	176.33	3053.17	538375.47
Case 1 Term 2	Sample			3062.39	525127.92
	Mom	17.86	171.48	3062.39	525127.92
	LVMAX	20.59	131.34	2704.34	355193.13
	C-LVMAX	18.59	167.55	3114.28	521807.77
Case 2	Sample			2581.22	165174.29
	Mom	40.34	63.99	2581.22	165174.29
	LVMAX	124.42	21.22	2640.22	56024.61
	C-LVMAX	97.71	27.17	2655.09	72144.41
Case 3	Sample			2353.02	301882.91
	Mom	18.34	128.30	2353.02	301882.91
	LVMAX	37.26	66.19	2466.19	163224.58
	C-LVMAX	25.66	94.93	2435.46	231194.00

Table 5. Parameter matching of weighted Gamma terms in MCU frame-size histograms

Experiment		Case 1				Case 2	
Iteration		Init	1	2	3	Init	1
Gamma Parameters	p_1	16.24	16.40	16.39	16.39	129.25	129.25
	p_2	32.23	32.46	32.46	32.46	309.81	309.81
	p_3	-	-	-	-	367.62	367.62
	p_4	-	-	-	-	657.52	657.52
	μ_1	163.07	161.38	161.45	161.45	17.29	17.29
	μ_2	181.39	180.11	180.11	180.11	17.15	17.15
	μ_3	-	-	-	-	25.63	25.63
	μ_4	-	-	-	-	20.56	20.56
weights	a_1	0.16	0.16	0.16	0.16	0.02	0.02
	a_2	0.84	0.84	0.84	0.84	0.14	0.14
	a_3	-	-	-	-	0.46	0.46
	a_4	-	-	-	-	0.38	0.38
Sample Moments	mean	5779.08				10081.69	
	variance	2927935.92				9734163.80	
Relative Difference	mean	-0.078	-0.078	-0.078	-0.078	+0.020	+0.020
	variance	-0.197	-0.200	-0.200	-0.200	-0.070	-0.070

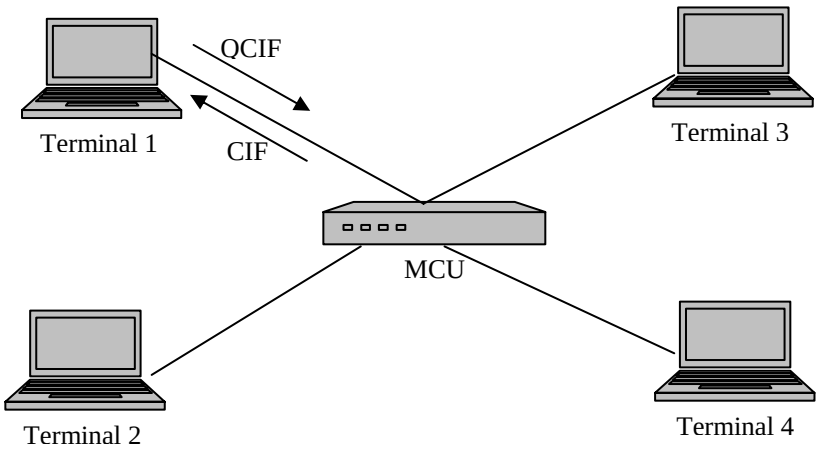

Figure 1. Testbed topology

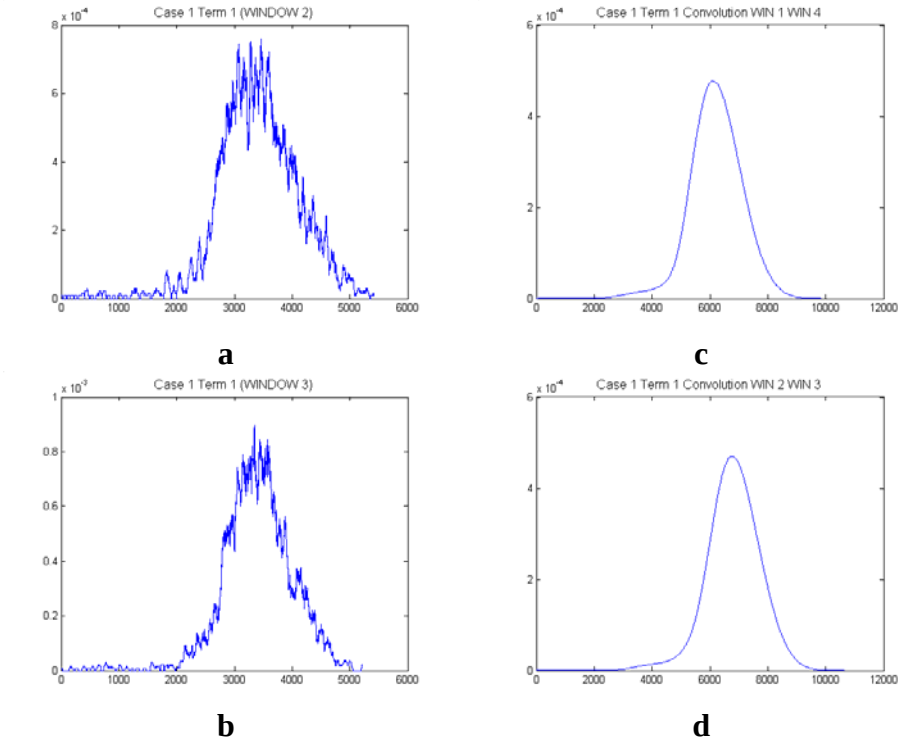

Figure 2. Frame size histograms in different windows (left) and convolutions of such histograms (right)

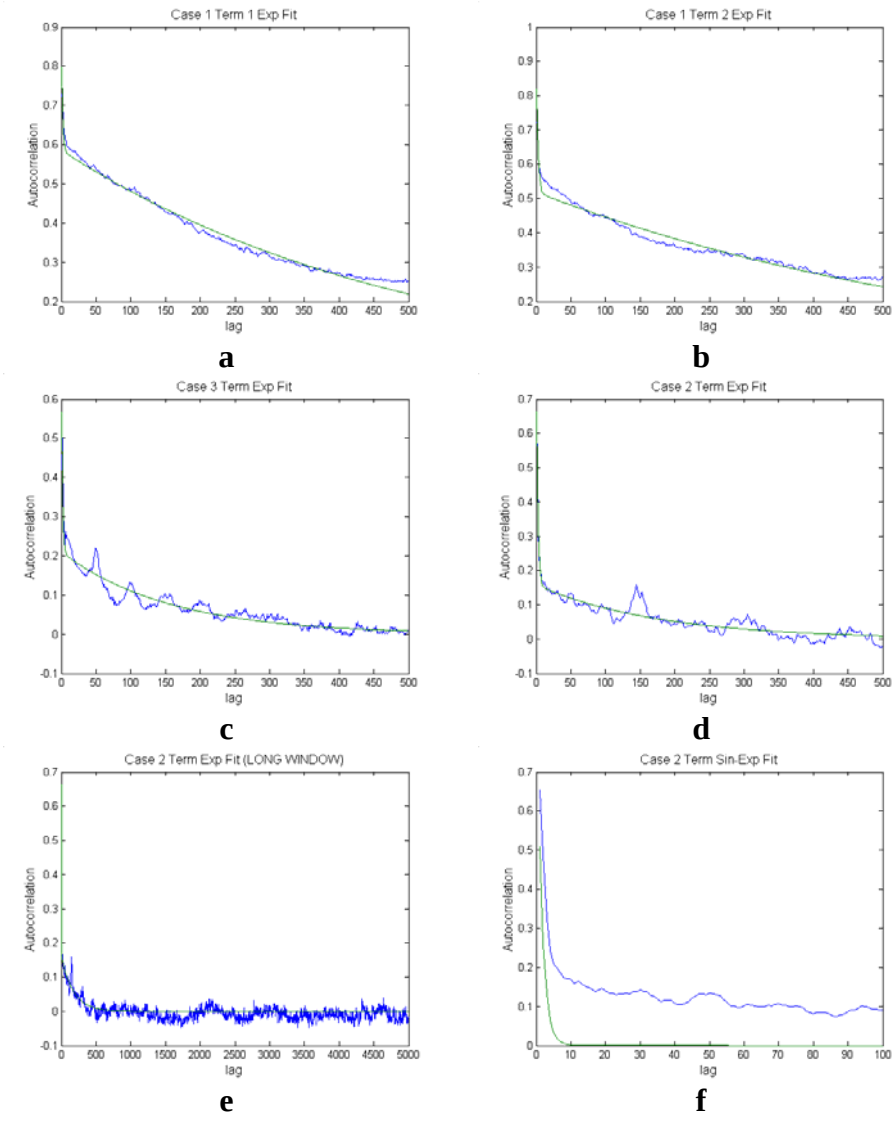

Figure 3. Autocorrelation graphs and fitted models for traffic from terminals: a,b,c,d: graphs and exponential fits for Term. 1&2 of Case 1 and terminals of Cases 3&2 (in stated sequence); e: wider range of fit for Case 2; f: dumped sinusoidal fit for Case 2

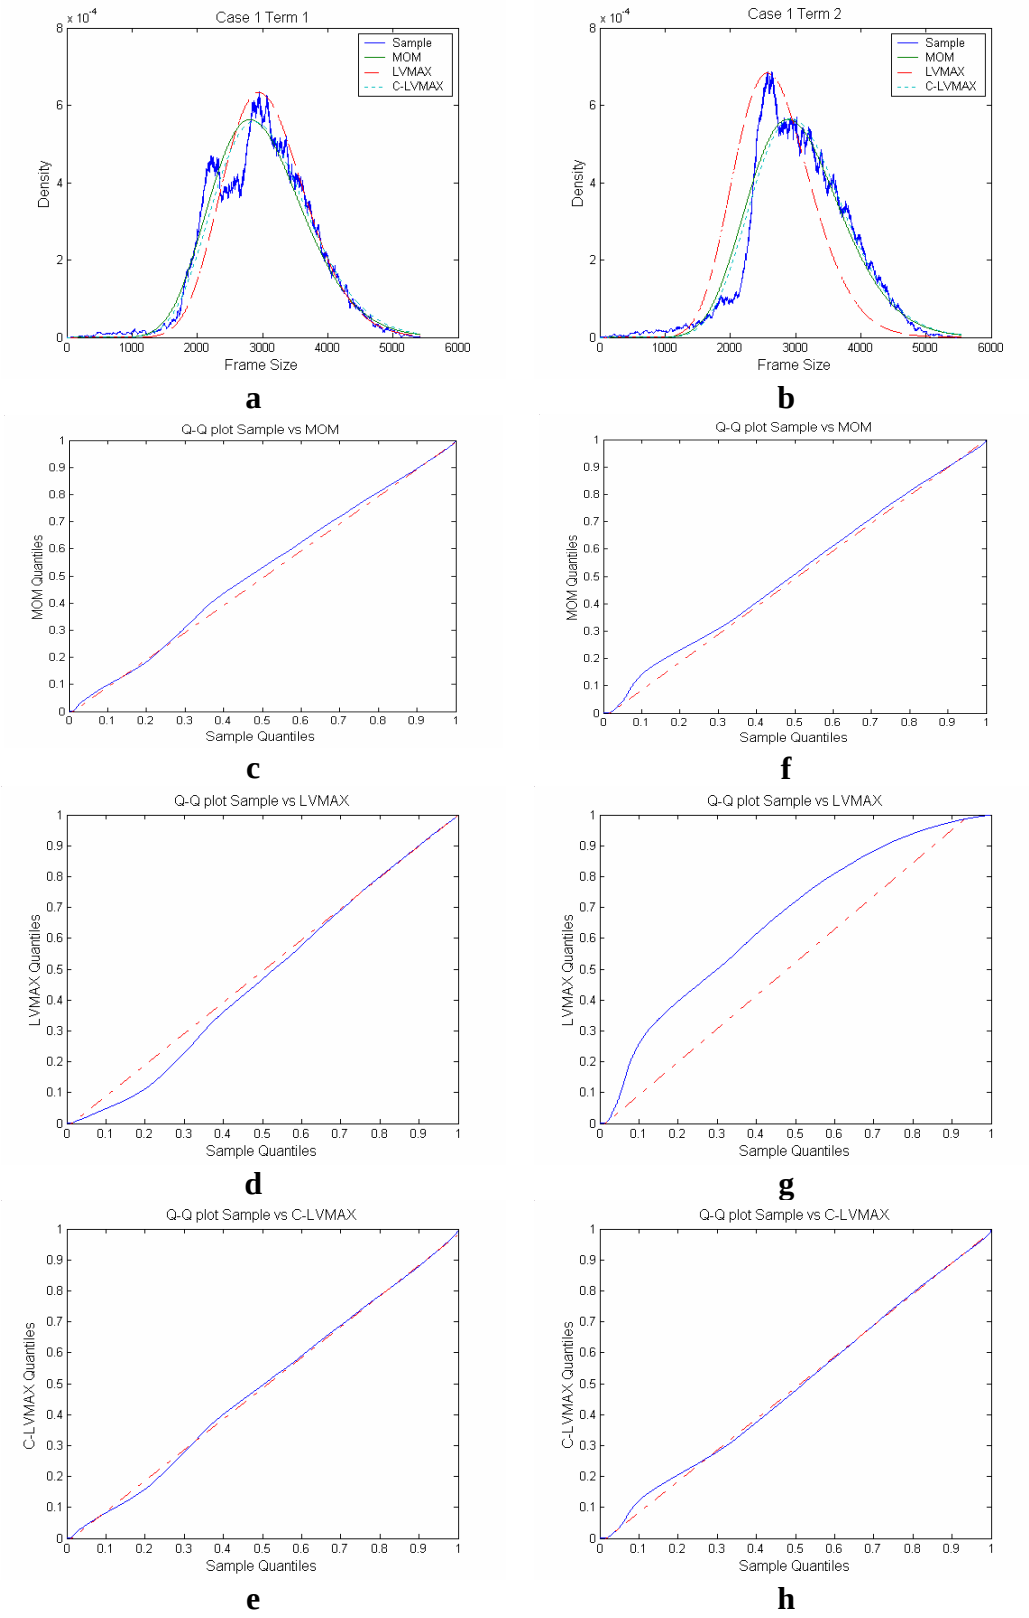

Figure 4. Frame size histograms and Gamma models for terminals in Case 1. Left Term 1, right Term 2. Top row for histograms and: moments fit (solid), LVMAX (dash-dot), C-LVMAX (dash). 2nd to 4th row for qqplots of histogram (hor. Axis) vs: moments fit, LVMAX, C-LVMAX (in stated order)

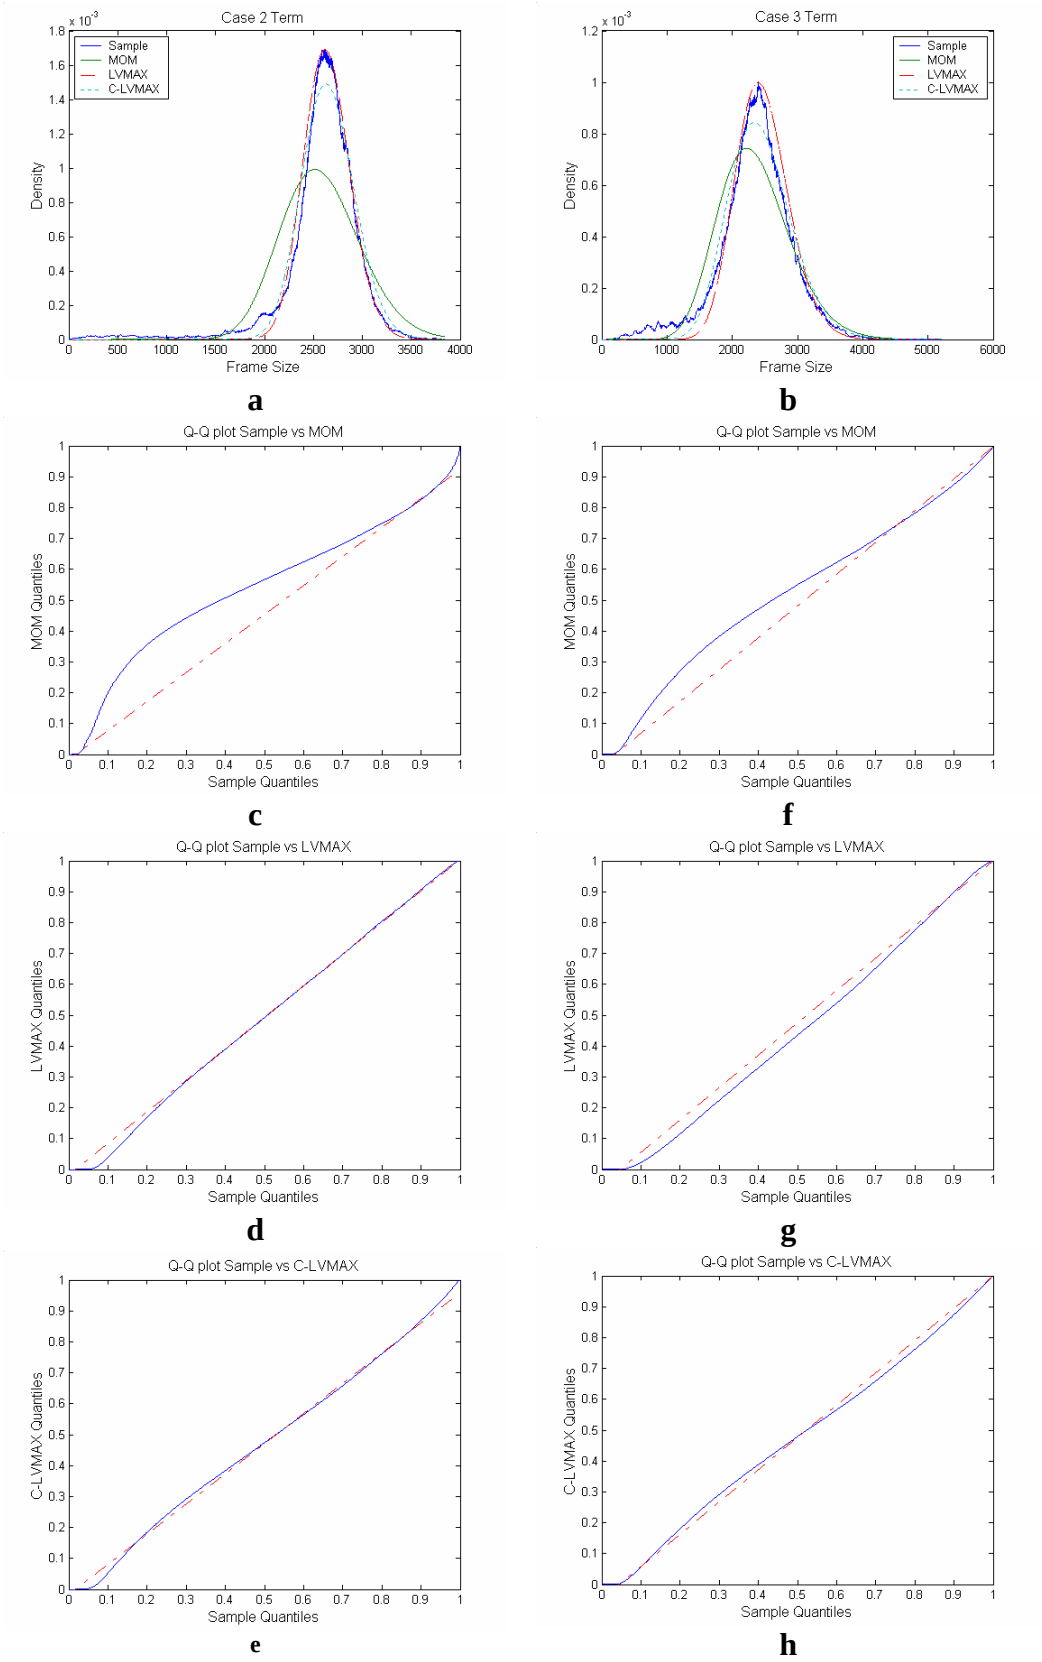

Figure 5. Frame size histograms and Gamma models for terminals in Cases 2 (left) and 3 (right). Top row for histograms and: mom. fit (solid), LVMAX (dash-dot), C-LVMAX (dash). 2nd to 4th row for qqplots of histogram (hor. Axis) vs: mom. fit, LVMAX, C-LVMAX (in stated order)

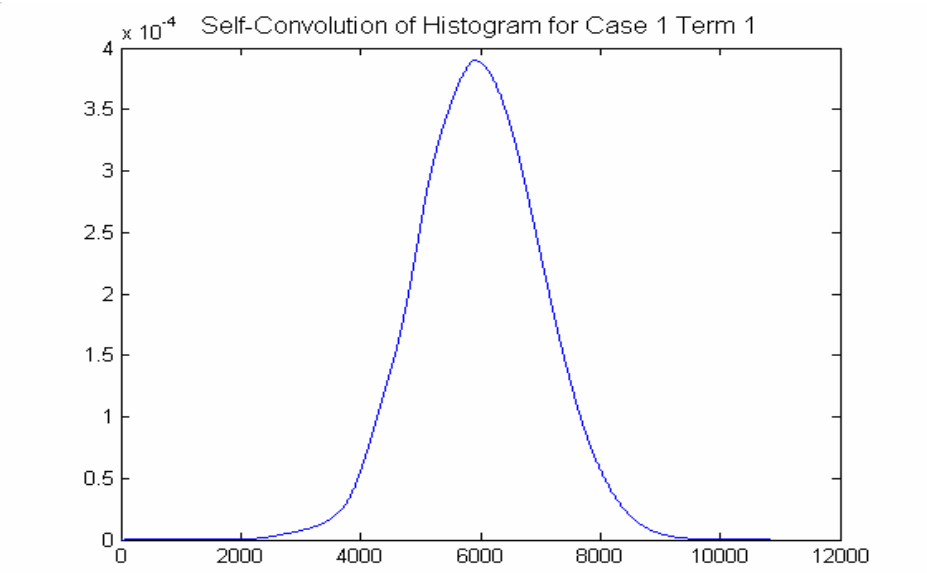

Figure 6. Convolution of the frame sizes histogram for Term 1 of Case 1 with itself

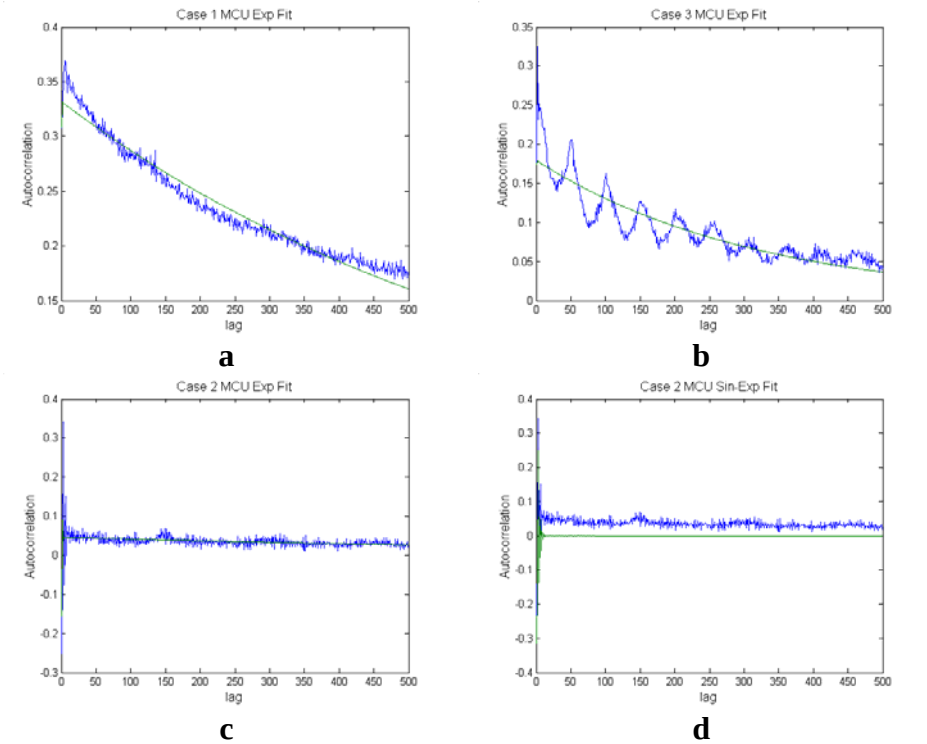

Figure 8. Autocorrelation graphs and fitted models for MCU traffic. a,b,c: graphs and exponential fits for Cases 1,3,2 (in stated sequence); d: dumped sinusoidal fit for Case 2

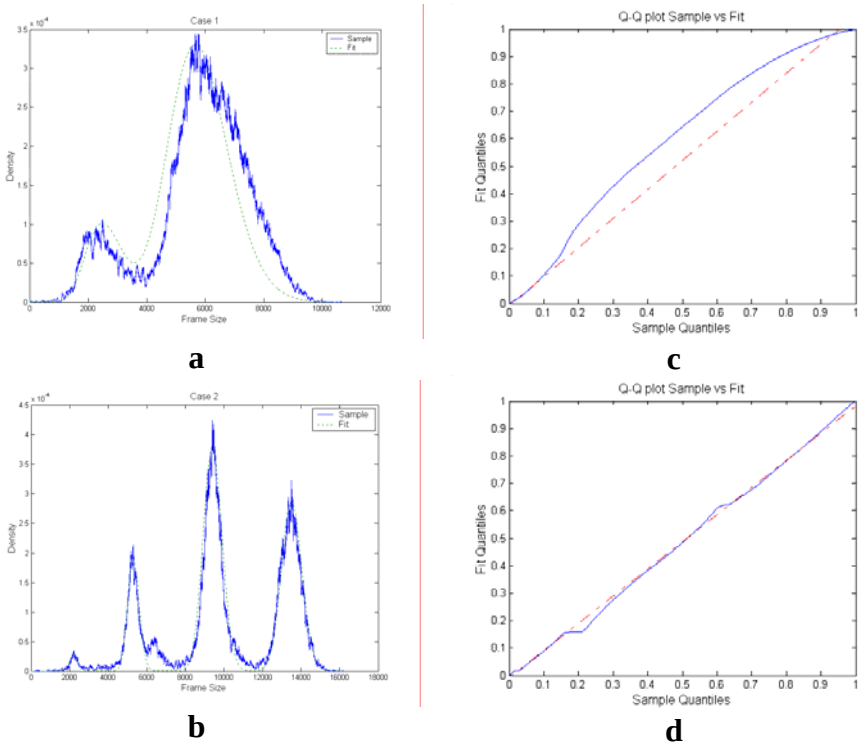

Figure 9. Frame size histograms and Gamma models for MCU in Cases 1 (top) and 2 (bottom). Right column for qqplots of histogram (hor. Axis) vs fit