

# Measuring Network Traffic: The Hidden Distribution

(or 'Nothing is ever quite what it seems in data communications!')

Vic Grout, CMath MIMA  
Senior Lecturer in Computing  
North East Wales Institute of Higher Education, Wrexham

## ABSTRACT

This piece, by way of an example, introduces a few questions encountered in data communications analysis and offers some (not quite so many) answers. A close inspection of figures resulting from a period of traffic measurement in a network reveals an interesting anomaly. Whilst the root cause is easily identified, following the path to its end leads to a complete and altogether more satisfactory solution.



unimportant, except to note that larger traffic samples across larger portions of a network are more likely to correlate with our predictions, assuming those predictions to be correct.

Traffic within a network or part of a network will take the form of a number of circuits established between end-points for the purposes of sending and/or receiving data. If we measure and log such activity at these end-points, we obtain a sequence of call connect/disconnect (CCD) signals that may be used to describe the underlying traffic.<sup>1</sup> If we can assume that a circuit established across one or more links dedicates those links to the call they carry (i.e. no data between other end-points can be carried at the same time) then we have all the information we need to describe the network usage.<sup>2</sup>

Two distributions, both available from the CCD signals, will tell us all we need to know about the network traffic. These are:

- the distribution of call arrivals – the *inter-call distribution* (ICD) – each measurement calculated as the difference in time between the arrival of one call and the arrival of the next.
- the distribution of call lengths – the *call duration distribution* (CDD) – each measurement calculated as the difference in time between the arrival of a call and its termination.

The two distributions have many features in common and we can, where appropriate, refer to a general *call distribution* (CD) where the distinction is unimportant.

In principle, both distributions are continuous so a *theoretical CD* would have a form as shown in Figure 1. Any sequence of

## Introduction

Communication networks today are big business. However, if you've got one, it pays to understand it! The proliferation and expansion of voice, data and integrated (voice and data) systems continues apace but they don't manage themselves. Networks have to be continually 'tweaked' to provide optimum performance, to offer the highest levels of service. A key feature in providing the right service is an awareness of what customers want in terms of the information they wish to send and receive. The slightest misunderstanding as to the nature of the traffic your network carries, the smallest misconception, can be costly. Seconds count in networking. This article examines a simple example, trivial in itself, but nevertheless a useful exercise in never taking anything at face value.

As with any large system, there are essentially two ways of trying to understand what is going on in a communications network: modelling and observation – *theory* and *practice* in other words. The ideal situation arises when both methods agree. A network in which theoretical predictions are borne out by empirical observations or in which observed peculiarities can be explained on paper is likely to be well-understood and consequently well-managed. However, this relationship is not always easily achieved.

## The Scene

When we speak of the *traffic* in a network, we essentially mean data flowing in an encoded form between pairs of points within it. These points may be simply the end points of a single transmission link, a client/server pair in a local area network (LAN) or two hosts joined by a wide area network (WAN) or internet. For the purpose of this simple discussion, the distinction is




Figure 1. Theoretical CD




Figure 2. In of theoretical CD

events, recorded with respect to the time at which they occur, and in which any two events, and hence their timings, are taken as independent (not always a justified assumption but let's work with it for now) will give rise to a negative exponential probability distribution with probability density function (pdf) given by

$$f(t) = \theta e^{-\theta t}$$

where  $t$  represents the time between two consecutive events. In such cases there is obviously a higher probability of short intervals than longer ones ( $f(a) > f(b)$  for  $a < b$ ) with the mean given by  $\mu$  where

$$\mu = \frac{1}{\theta}$$

Taking the natural logarithm of the theoretical CD pdf gives

$$\begin{aligned} L(t) &= \ln(f(t)) \\ &= \ln(\theta e^{-\theta t}) \\ &= \ln\theta - \theta t \\ &= K - \theta t \end{aligned}$$

(where  $K$  is a constant) which is linear in  $t$  as shown in Figure 2. This provides a useful and convenient check for observed data. Data taken from a small college LAN is presented in Figure 3 for comparison. Naturally, we expect a degree of real-world deviation from the theoretical model but some particular causes are discussed in the next section.

As a slight aside, if the ICD follows this model and we let  $c$  represent the number of calls arriving in a given unit of time, then  $c$  will follow the *Poisson distribution* (as shown in Figure 4) with pdf

$$p(c) = \frac{\lambda^c e^{-\lambda}}{c!}$$




Figure 3. In of observed data (college LAN)




Figure 4. Poisson call arrivals

where  $\lambda$  is the mean number of call arrivals in unit time so that

$$\lambda = \frac{1}{\mu} = \theta$$

For this reason, phenomena such as calls within a communications system found to have negative exponential interval pdfs are commonly referred to as having Poisson characteristics or obeying the Poisson model.

### So Much For The Theory!

So call distributions in a network will always look like this, will they? Well, no! Actually there's quite a lot that can happen. Different types of system, using various technologies, will have different CDs. Even within a network, there will be distributed deviations from the norm, reflecting local configurations or problems such as bottlenecks and equipment failure ('outs'). Measuring – and *understanding* – network traffic can be something of a Forth Bridge exercise. Even if you can get a grip on what's happening, the smallest reconfiguration can throw it all off and you have to start investigating all over again. There's a lifetime's work involved in trying to work out what's going on in a network so it's probably as well for us that the networks themselves rarely last that long! They disappear or are replaced or evolve into the next version. Whether any of them are ever truly understood before they go is debatable.<sup>3</sup>

In fact, even if a network *does* behave the way we would like it to – that is, if it follows our neat models, it often hides it rather well! The example that follows is trivial but should be instructive.

Suppose we monitor and log a (preferably large) session of CCD signals. We tabulate the results and they look something like Figure 5.<sup>4</sup> This almost looks like the negative exponential distribution from Figure 1, but is it? If it is then what has happened

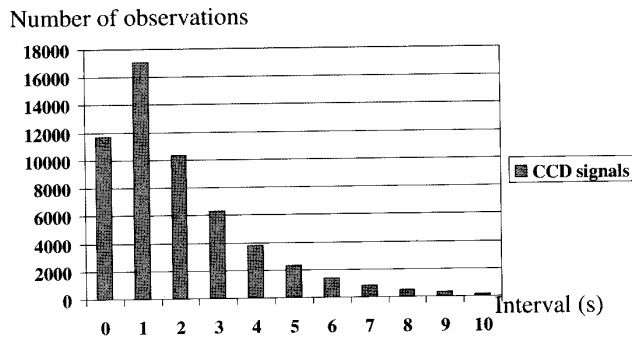


Figure 5. Observed data (test network)

to the first value? That doesn't look right. Are there fewer calls of the shortest length being carried on this network? Should we tweak the network parameters to accommodate?

To do so would be an overreaction (particularly if we note that a similar anomaly appears to be present in Figure 3) but worse things have happened in practice! However, if we take a look at where these figures came from, we should be able to get an idea of why they look as they do.

In fact, it's not difficult to see where the discrepancy originates. The problem stems from the limited accuracy with which we measure events. The theoretical CD is continuous whereas the observed CD cannot be. Typical call logging will be accurate to the nearest second<sup>4</sup> so our intervals will be recorded as integers. This is not of purely academic interest. Calls of less than a second in duration are commonplace today. Control packets are typically short and even large data packets take little time to transmit. A standard E1 link transmits at 2.048Mbps – over two million bits per second – and much higher data rates are available. You can transmit a lot of data in a second – and it can go a long way!

So we shouldn't expect our CD to look exactly as Figure 1. Fine, but what should it look like? What would we expect to see from our observed CD if the underlying theoretical CD is negative exponential? Our observations suggest a distribution very similar to the negative exponential in all but the first value. Is there such a distribution? Well, if there is, we can derive it from first principles ...

The interval we are trying to measure can be considered as the difference between times  $t_i$  and  $t_{i+1}$  ( $t_i < t_{i+1}$ ) at which events  $\Phi_i$  and  $\Phi_{i+1}$  occur ( $1 \leq i \leq n$ ;  $n$  the number of events over some period). The elapsed time between events  $\Phi_i$  and  $\Phi_{i+1}$  is given by

$$\Delta_i = t_{i+1} - t_i$$

so that  $\Delta_i$  is a positive real. If we could approximate  $\Delta_i$  (in seconds say) by  $\lfloor \Delta_i \rfloor$ , the integer part of  $\Delta_i$ , we would not have a problem. Instead, however, it is the values of  $t_i$  and  $t_{i+1}$  that are truncated on logging to  $\lfloor t_i \rfloor$  and  $\lfloor t_{i+1} \rfloor$  respectively and it is this that leads to the distorted distribution. The apparent elapsed time between events  $\Phi_i$  and  $\Phi_{i+1}$  will be recorded as

$$\delta_i = \lfloor t_{i+1} \rfloor - \lfloor t_i \rfloor$$

so that  $\delta_i$  is a non-negative integer. In general of course  $\delta_i$  may not be equal to  $\lfloor \Delta_i \rfloor$  as suggested in Figure 6 so that, while the distribution of  $\delta$  may be related to the distribution of  $\Delta$ , the two are not identical. (In proceeding to discuss the distributions of  $\Delta$  and  $\delta$ , the subscripts have been dropped for simplicity.)

We can however establish a connection. Obviously

$$\Delta = \lfloor \Delta \rfloor + \{\Delta\}$$

where  $\{\Delta\}$  is the fractional part of  $\Delta$  with  $0 \leq \{\Delta\} < 1$ . For any given value of  $\Delta$ ,  $\delta$  will equal either  $\lfloor \Delta \rfloor$  or  $\lfloor \Delta \rfloor + 1$  and we can calculate probabilities exactly as follows.

Define the conditional probability function  $p(\delta | \Delta)$  to be the probability that  $\lfloor \beta \rfloor - \lfloor \alpha \rfloor = \delta$  given that  $\beta - \alpha = \Delta$ . Then, assuming the fractional parts of the times  $t_i$  to be uniformly distributed over the interval  $0 \leq \{t_i\} < 1$ , we have

$$\begin{aligned} p(\lfloor \Delta \rfloor - j | \Delta) &= 0 & \text{for } j \geq 1, \\ p(\lfloor \Delta \rfloor | \Delta) &= 1 - \{\Delta\}, \\ p(\lfloor \Delta \rfloor + 1 | \Delta) &= \{\Delta\}, \\ p(\lfloor \Delta \rfloor + j | \Delta) &= 0 & \text{for } j \geq 2 \end{aligned}$$




Figure 6. Truncation and subtraction

as shown in Figure 7. The shape of  $p(\delta | \Delta)$  for different values of  $\delta$  is illustrated in Figure 8 and from this we can see that the  $p(\delta | \Delta)$  function can be written in rearranged form as

$$p(\delta | \Delta) = \begin{cases} 0 & \text{for } \Delta < \delta - 1, \\ \Delta - \delta + 1 & \text{for } \delta - 1 \leq \Delta \leq \delta, \\ \delta - \Delta + 1 & \text{for } \delta \leq \Delta \leq \delta + 1, \\ 0 & \text{for } \Delta > \delta + 1. \end{cases}$$

If  $f(\Delta)$  is the pdf of the continuous variable  $\Delta$  then  $g(\delta)$ , the pdf of the discrete variable  $\delta$ , is given generally by

$$g(\delta) = \int_0^{\infty} f(\Delta) p(\delta | \Delta) d\Delta$$

If  $\Delta$  follows the negative exponential distribution with  $f(\Delta) = \theta e^{-\theta \Delta}$ , the pdf  $g(\delta)$  is given by

$$g(\delta) = \int_0^{\infty} \theta e^{-\theta \Delta} p(\delta | \Delta) d\Delta$$

We need to consider two cases separately since  $p(\delta | \Delta)$  is symmetric about  $\delta$  at each point except  $\delta = 0$  (see Figure 8). This indeed is the root cause of the anomaly in figure 5. For  $\delta = 0$ , we have

$$\begin{aligned} g(0) &= \int_0^1 \theta e^{-\theta \Delta} (1 - \Delta) d\Delta \\ &= 1 + \frac{e^{-\theta} - 1}{\theta} \end{aligned}$$

and, for  $\delta \geq 1$ ,




Figure 7. The function  $p(\delta | \Delta)$




Figure 8.  $p(\delta | \Delta)$  for various  $\delta$

$$g(\delta) = \int_{\delta-1}^{\delta} \theta e^{-\theta \Delta} p(\Delta - \delta + 1) d\Delta + \int_{\delta}^{\delta+1} \theta e^{-\theta \Delta} (\delta - \Delta + 1) d\Delta$$

$$= \frac{2e^{-\theta \delta} (\cosh \theta - 1)}{\theta}$$

Before we go any further, we can check that this is a valid pdf as follows:

$$\sum_{\delta=0}^{\infty} g(\delta) = 1 + \frac{e^{-\theta} - 1}{\theta} + \frac{2(\cosh \theta - 1)}{\theta} \sum_{\delta=1}^{\infty} e^{-\theta \delta}$$

$$= 1 + \frac{1}{\theta} \left[ e^{-\theta} - 1 + \frac{(e^{\theta} + e^{-\theta} - 2)e^{-\theta}}{1 - e^{-\theta}} \right]$$

$$= 1$$

as required. Values of  $f(\Delta)$  and  $g(\delta)$  for different  $\theta$  are given in Tables 1 and 2 and a graphical comparison in Figure 9. There is an obvious difference at zero and then this reduction is made up almost imperceptibly over the values that follow. The shape of  $g(\delta)$  for  $\theta = 1/2$  closely resembles that of Figure 5 whereas the logarithmic form is very similar to Figure 3. As a comparison, a large data sample taken from a national enterprise network is shown in Figure 10. Again, real data deviates from any theoretical model but the general shape is apparent and the reduced initial value is clearly visible. Bearing in mind the fluid nature of data communications traffic, it would be presumptuous to claim to have found the 'correct' distribution but we do appear to have something that works!

## Conclusion

This is a simple exercise, not unimportant, but simple. If we think of the study of networks as being like a network itself – with

Table 1.  $f(\Delta)$  for various  $\theta$

	$\theta = 0.5$	0.4	0.3	0.2	0.1	0.05	0.01
$\Delta = 0$	0.500	0.400	0.300	0.200	0.100	0.050	0.010
$\Delta = 1$	0.303	0.268	0.222	0.164	0.090	0.048	0.010
$\Delta = 2$	0.184	0.180	0.165	0.134	0.082	0.045	0.010
$\Delta = 3$	0.112	0.120	0.122	0.110	0.074	0.043	0.010
$\Delta = 4$	0.068	0.081	0.090	0.090	0.067	0.041	0.010
$\Delta = 5$	0.041	0.054	0.067	0.074	0.061	0.039	0.010
$\Delta = 6$	0.025	0.036	0.050	0.060	0.055	0.037	0.009
$\Delta = 7$	0.015	0.024	0.037	0.049	0.050	0.035	0.009
$\Delta = 8$	0.009	0.016	0.027	0.040	0.045	0.034	0.009
$\Delta = 9$	0.006	0.011	0.020	0.033	0.041	0.032	0.009
$\Delta = 10$	0.003	0.007	0.015	0.027	0.037	0.030	0.009
$\Delta = 11$	0.002	0.005	0.011	0.022	0.033	0.029	0.009
$\Delta = 12$	0.001	0.003	0.008	0.018	0.030	0.027	0.009
$\Delta = 13$	0.001	0.002	0.006	0.015	0.027	0.026	0.009
$\Delta = 14$	0.000	0.001	0.004	0.012	0.025	0.025	0.009
$\Delta = 15$	0.000	0.001	0.003	0.010	0.022	0.024	0.009
$\Delta = 16$	0.000	0.001	0.002	0.008	0.020	0.022	0.009
$\Delta = 17$	0.000	0.000	0.002	0.007	0.018	0.021	0.008
$\Delta = 18$	0.000	0.000	0.001	0.005	0.017	0.020	0.008
$\Delta = 19$	0.000	0.000	0.001	0.004	0.015	0.019	0.008
$\Delta = 20$	0.000	0.000	0.001	0.004	0.014	0.018	0.008

Table 2.  $g(\delta)$  for various  $\theta$

	$\theta = 0.5$	0.4	0.3	0.2	0.1	0.05	0.01
$\delta = 0$	0.213	0.176	0.136	0.094	0.048	0.025	0.005
$\delta = 1$	0.310	0.272	0.222	0.164	0.091	0.048	0.010
$\delta = 2$	0.188	0.182	0.166	0.135	0.082	0.045	0.010
$\delta = 3$	0.114	0.122	0.123	0.110	0.074	0.043	0.010
$\delta = 4$	0.069	0.082	0.091	0.090	0.067	0.041	0.010
$\delta = 5$	0.042	0.055	0.067	0.074	0.061	0.039	0.010
$\delta = 6$	0.025	0.037	0.050	0.060	0.055	0.037	0.009
$\delta = 7$	0.015	0.025	0.037	0.049	0.050	0.035	0.009
$\delta = 8$	0.009	0.017	0.027	0.041	0.045	0.034	0.009
$\delta = 9$	0.006	0.011	0.020	0.033	0.041	0.032	0.009
$\delta = 10$	0.003	0.007	0.015	0.027	0.037	0.030	0.009
$\delta = 11$	0.002	0.005	0.011	0.022	0.033	0.029	0.009
$\delta = 12$	0.001	0.003	0.008	0.018	0.030	0.027	0.009
$\delta = 13$	0.001	0.002	0.006	0.015	0.027	0.026	0.009
$\delta = 14$	0.000	0.001	0.005	0.012	0.025	0.025	0.009
$\delta = 15$	0.000	0.001	0.003	0.010	0.022	0.024	0.009
$\delta = 16$	0.000	0.001	0.002	0.008	0.020	0.022	0.009
$\delta = 17$	0.000	0.000	0.002	0.007	0.018	0.021	0.008
$\delta = 18$	0.000	0.000	0.001	0.005	0.017	0.020	0.008
$\delta = 19$	0.000	0.000	0.001	0.004	0.015	0.019	0.008
$\delta = 20$	0.000	0.000	0.001	0.004	0.014	0.018	0.008

numerous paths connecting many different points – then we have merely taken a single journey from a source to a destination, one of many on that route alone and amongst many others. However, in so doing, we have highlighted some general points, applicable in the wider field.

Networks today are large and complex. This is true of many systems of course, including other types of network (transport, physical, chemical, economic, etc.) but the sheer speeds involved in telecommunications tend to compound this complexity. Whilst cars on a motorway, for example, can be counted manually or with fairly primitive equipment, and traffic patterns adequately described using time slots of so many minutes or hours, network data traffic requires a far greater precision. The effects of minute changes in behaviour can be both considerable and quick to take effect. Congestion can build up in data networks in a fraction of a second, congestion patterns can be widespread but complex and interdependent; and in the (possibly) short time it takes the network to respond, huge quantities of data can be lost. This data may then be retransmitted by higher-level network protocols, adding to the original congestion. The effect of motorway traffic bunching too tightly together is well understood nowadays. One vehicle slows down, the car behind brakes harder, the car behind harder still. Eventually one vehicle stops altogether and a traffic jam appears from nowhere. Imagine a similar situation happening in a data network but millions of times a second and at electron speed!

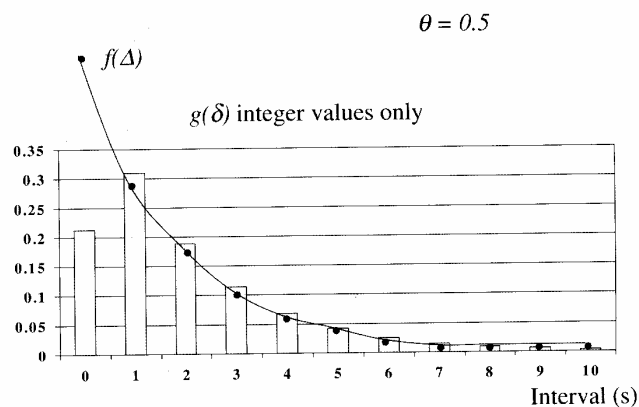


Figure 9.  $f(\Delta)$  and  $g(\delta)$




Figure 10. Observed data (national network)

It's not, in principle, difficult to measure data communications traffic, even extremely accurately. However there are a variety of methods at our disposal and choosing effectively is vital. We need to ensure that we know what we intend to measure, that our approach will in fact measure what we intend to, that data taken in this way is recorded appropriately and that it can be delivered without biasing its own source. We then have to take care in interpreting this data: it's only too easy to overlook a small point of great importance or to overplay the relevance of an incorrectly presented or misunderstood feature.

Traffic analysis is an important area of study in networking. An article such as this can do little more than open the door and suggest what's there. Essentially, the limited discussion here illustrates how a closer look at a system can reveal the underlying truth. In a high-speed/large-throughput field such as data communications (and many other areas of applied mathematics), understanding such subtleties can make a world of difference.

## Acknowledgement

The advice from the referees was invaluable in preparing this piece for publication and was gratefully received. Thanks also to John Worden for proof-reading the final version. □

## END NOTES

1. It's probably worth pointing out that there are many ways of measuring network traffic using what we might informally call *high-level* or *low-level* approaches or alternatively, working on a *large* or *small* scale. The method discussed here is a *medium* – possibly about a size 14!
2. A *circuit-switched* network is being described here. A circuit is established between end-points for the duration of the call. An alternative is *packet-switching* in which data is transmitted in independent blocks. There are different flavours of packet switching but their arrival characteristics are similar in nature to circuit establishment. However packet lengths tend to be more uniform, due partly to the presence of a number of control packets as well as data packets having an upper bound.
3. A obvious case in point is the Internet itself. This evolved fairly rapidly from the ARPANET, a US defence project network. The ARPANET was extremely experimental in a number of areas to the extent that network scientists were still discussing lessons learned from it long after the Internet was a global fixture!
4. These figures were actually taken from a laboratory test network – possibly about as idealised a 'real-world' network as we can get!
5. It is of course possible to record with a much greater accuracy but to the nearest second is typical. Call logging is not a 'natural' function of a network in that it does not contribute to its real-time operation. In fact, used carelessly, it can add substantially to the network traffic we are trying to measure. If data recorded across the network is transmitted to a central point then it may bias its own figures: in a worst case scenario it may even cause unnecessary congestion!

**Vic Grout** graduated from Exeter University in Mathematics and Computing in 1984 and was awarded a Ph.D. for his algorithmic design work in communication engineering at Plymouth Polytechnic in 1988. His research has been concerned with the simulation and optimisation of large systems including the design and management of communication networks. He is currently investigating the new problems in network administration emerging from recent advances in network technology.

## Instructions to Authors for *Mathematics Today*

*Mathematics Today* is primarily a general interest publication for mathematics graduates and is not a research journal. It aims to provide members of the Institute with news and other informative items on mathematics and related topics. It includes articles of wide mathematical interest, overviews of recent developments or applications, mathematics in education, news items, news of members, book reviews, puzzles and letters. Contributions from Institute members and non-members will be considered.

### Submissions

Manuscripts may be submitted as hard copy, by email (in Word 97, Word Perfect up to version 6, or Latex), disc or fax and must be entirely the work of the named author(s). If the article has been published before or if it is being considered for another publication this must be clearly stated and full details given

- **Titles:** Authors should choose short and interesting titles for their articles and the articles should not normally exceed 3000 words in length.
- **Abstracts:** A short summary of articles should be provided.
- **Author Details:** For each author please provide a full name, address and email address/daytime telephone and a photograph and short biography.
- **Photographs, illustrations:** Please ensure that copyright approval has been obtained for all illustrations used.

### Editorial and refereeing policy

It is Institute policy to send all articles, letters and puzzles, that have mathematical or scientific content to referees, as a matter of routine, in order to ensure the accuracy of the material. At the same time confirmation is sought from the referees that the article is suitable, both in subject matter and style, for *Mathematics Today*. A balance of topics, and items, is maintained within issues and over several issues. Non-acceptance does not imply that an article does not meet the academic standard required. The Institute reserves the right to edit manuscripts for clarity of expression and to adjust the length of articles.

Permission must be obtained before items from *Mathematics Today* are reproduced or published in other publications. Submissions should be sent to Gayna Leggott, Editorial Officer, Email: [gayna@ima.org.uk](mailto:gayna@ima.org.uk) Tel: 01702 356111