An Interview with Erol Gelenbe

Interviewed by Cristian Calude

Editor's Introduction

This is Part I of an interview with Professor Erol Gelenbe, conducted by Professor Cristian Calude. Gelenbe holds the Dennis Gabor Chair Professorship in the Electrical and Electronic Engineering Department at Imperial College London and is an associate editor for this publication. This interview also appeared in the October 2010 issue of the [Bulletin of the European Association for Computer Science](#) and is printed here with permission.

—Editor

An Interview with Erol Gelenbe

Interviewed by Cristian Calude

[Erol Gelenbe](#) holds the Dennis Gabor Chair Professorship in the Electrical and Electronic Engineering Department at Imperial College London. He is an alumnus of the Middle East Technical University, Ankara, Turkey, and received a PhD from the Polytechnic Institute of New York (Brooklyn Poly) and the Docteur ès Sciences degree from the Université Pierre et Marie Curie (Paris VI).

As one of the founders of the field of computer system and network performance analysis, he is also well known for inventing random neural networks and G-networks. Two of his four books were published in English, French, Korean, and Japanese. Elected to the Turkish Academy of Sciences, the Hungarian Academy of Sciences, the Académie des Technologies (France) and Academia Europaea, he is a Fellow of IEEE (1986), an ACM Fellow (2001), and has received honorary degrees from the Universities of Liège, Bogaziçi (Istanbul), and Rome II. He has graduated more than 50 PhDs including many women computer scientists and engineers, and was awarded the ACM SIGMETRICS Life-Time Achievement Award (2008) and other technical awards.

Cristian Calude: Tell us about your experience of studying and working in so many countries.

Erol Gelenbe: No two countries and no two environments, and indeed no two institutions, are identical. There is no “best” way to do things and each institution has evolved or adapted according to the personalities of its leaders and the specific context within which it is operating.

Moving from one institution to another is interesting and challenging, especially from one country to another, and can be a source of fun and learning (for me, at least). However, being a foreigner almost everywhere, I can group countries and institutions into two broad categories: those that are open to “allogens” and are willing to be inclusive, and those which have (sometimes in subtle ways) set up significant barriers to “foreign” penetration.

Erol Gelenbe (right) holding his diploma from the Middle East Technical University, Ankara, Turkey, in 1966.

It is quite different if you are a visiting professor: You are there temporarily and do not constitute a threat to established ways. If you arrive as a potential permanent addition, matters are different, more challenging and more interesting. I have held chairs in Belgium, France, the U.S. and U.K. In some countries foreigners are discriminated against illegally, including in matters of promotions, awards, etc., and have to bear with derogatory comments. Such things may continue even after a “foreigner” actually acquires the country’s citizenship, or is from the EU and works in another EU country. One hears about illegal immigrants, but seldom does one talk about illegal barriers imposed on lawful immigrants.

My strangest experience was with my candidacy to become Institute Director at CNR (Italy) in 2000 at the request of Italian colleagues. I am fluent in Italian. My application was eliminated from consideration because they officially said that they could not read my signature: “*firma non leggibile*.” Yet my credit card payments have always been accepted in Italy without problems regarding my signature! In that instance, a large number of CNR institute directorships were being filled, and they deftly managed to avoid appointing even a single foreign candidate.

Recently, Brice Hortefeux, the French Interior Affairs Minister, was condemned by a French court for “racial insults” to a person of North African descent who belongs to the same political party. Hortefeux has not resigned despite the court decision. Some 20 years ago, I suggested that the French Computer Society AFCET co-sponsor a conference attended by French colleagues that was held in Turkey. The then president of AFCET said that he did not want to join a Turkish quagmire (“*bourbier turc*”). Since then AFCET has disappeared, while the Turkish mud hole (a.k.a., the International Symposium on Computer and Information Sciences) is holding its 25th anniversary conference in 2010 at the Royal Society in London! Even more strangely, a couple of years ago the lady at a *Metro guichet* in Paris refused to sell discount day tickets to two of my Greek and Greek Cypriot PhD students on the grounds that she would not sell discount tickets to foreigners. Well, such stories are endless and unfortunately common.

CC: As a computer scientist, engineer, and applied mathematician, you have done a lot of theory, part of which was used in commercial products. For example, your early work was incorporated into the QNAP2 software package.

EG: I have indeed attended to theory, but except for my early work on stochastic automata [Automata], my work is motivated by “practical” problems or direct physical inspiration. For instance, I got involved in performance modelling and then in queuing theory because of two practical drivers. Shortly after defending my PhD, I spent two summers at the Phillips Research Laboratories in Eindhoven, where I was asked to work on memory management algorithms for stack oriented re-entrant programmes [Phillips1],[Phillips2]. I knew nothing about the subject but was annoyed by the “ad-hoc” nature of the design choices that were being made. So I felt that some theory was needed, for instance in the choice of the page and memory segment sizes, so as to optimise the overhead.

Similarly, in my first job at the University of Michigan (Ann Arbor) as an assistant professor, they asked me to teach computer architecture. Everyone already there had taken over the “nice” theoretical courses on automata theory, formal languages, etc. so I (the newcomer) was “stuck” teaching the subjects that others did not wish to teach. Well, there again, I became involved in developing a more quantitative and seemingly rational (at least to me) approach to computer architecture and operating systems, which has given rise to the field of computer and network performance analysis and evaluation. For instance, I was able to prove results on paging algorithm performance [Paging] which attracted the attention of some Hungarian and Russian mathematicians and physicists in the mid 1970s, as well as on memory space optimisation [Lifetime] which drew theoretical conclusions from Laszlo Belady’s earlier measurements at IBM on

“life-time functions.” My development of new “product form” networks [PF,CRAS], which are also linked to statistical mechanics and theoretical chemistry, was motivated by listening to presentations from neuroscientists while visiting the RIACS/NASA Research Centre, but that is yet another story to be told below.

Erol Gelenbe (right) summarizing “diffusion approximations for queueing networks” [Diff1], [Diff2] at IRIA in 1973 or 1974 to Jacques-Louis Lions (left) and Alain Bensoussan (center).

So yes, much of my work has had a theoretical bent, but it has almost always been driven by a strong link with engineering requirements or by observations from nature.

Another example inspired by engineering is the research I did on optimum checkpointing [CP] in databases, which appeared in the *Journal of the ACM*, but was motivated by a practical issue that was recounted to me by Claude Delobel in relation to the automatic storage in a database of “hits” during some fencing championships that were taking place in Grenoble, and the computer being used had intermittent failures! This work gave rise to a few PhD theses around me and to much more work around the world.

Another quite theoretical result on an equivalence property of servers with vacations [Iasno] was motivated by a property I had observed in a simpler Poisson-arrival context while writing a book, and on the intuition that it actually was true in a more general framework. A probability models of uncertainty in databases [Hebrail] was motivated by a problem that was posed to me by Jean-Louis Bacquet-Grammont, a historian of the Ottoman Empire, concerning the interpretation of historic data.

The QNAP2 (then Modline) software tool for performance evaluation was developed by our group at IRIA (now INRIA), and the specific technique that I personally contributed was on diffusion approximations, first published in the *Journal of the ACM* [Diff1],[Diff2]. This software tool has generated some €200 million over 20 years for the companies that commercialized it (initially SIMULOG an INRIA spin-off company). The developers and inventors themselves hardly got anything. We were naive about such things. Throughout my career I have been involved with industry, via patents, via tools such as QNAP2, via consultancies or short-term assignments inside industry, and also via contracts to my university that are directly funded by industry. Of course, many of my PhD students have gone on to work for industry, most recently in the financial sector.

CC: You are a pioneer in the adaptive control of computer systems.

EG: I am bit like the elephant in the dark room. Someone comes into the room, touches the leg of the elephant and thinks it is a tree, another person feels the tail and thinks it's a rope, and yet another catches the elephant's nose and thinks it's a hose! Some people think that I am a pioneer of computer system performance evaluation (at least that is what they say on my ACM Sigmetrics Award, and my IEEE and ACM Fellowship Awards). The French Academy in 1996 gave me the France-Telecom Prize for developing mathematical models of communication networks. The Hungarian Academy of Sciences, in its recent election, mentions both my work on system and network performance and on neural networks and learning.

As you indicate, in the last six or seven years I have been involved in developing ideas on Adaptive Computer Systems and Networks [Users] such as the "Cognitive Packet Network," and this has helped generate research projects in "Autonomic Communications" in Europe. I was asked to [write a paper about this work](#) for the [July 2009 issue](#) of *Communications of the ACM*. The Internet is largely a legacy system based on principles that find their origin in the computers and data communication systems of the 1970s, and it works pretty well. Thus it is hard to introduce new concepts and methods into the Internet. Much of the current network research only addresses tweaks to well understood aspects.

CC: Tell us about random neural networks. You developed their theory—mathematics and learning algorithms—as well as some applications to engineering and biology.

EG: Let me first tell you what the random neural network (RNN) is [RNN], and then I will get around to telling you how it came about.

Consider a system composed of N counters, and let them be numbered $i, j=1, \dots, N$. Each counter can have a value which is a natural number. At any instant t , only one of the following events can occur: the i -th counter increases by one (external arrival of an excitatory spike to neuron i), or if a counter has a positive value it may decrease by one (neuron i fires and the spike is sent out of the network), or a counter i may decrease by 1 and simultaneously some other counter increases by 1 (neuron i fires an excitatory spike which arrives instantaneously to neuron j), or i decreases by 1 and so does j if both start in a positive value (neuron i fires an inhibitory spike to j), or finally at time t nothing happens. What this is modelling is a network of N neurons which are receiving and exchanging excitatory or inhibitory spikes. The system operates in continuous time so that t is a real number, making it a continuous time network of counters. It is quite extraordinary that this very simple model has some very powerful properties including the ability to learn [RNN-Learn1], and the ability to approximate continuous and bounded functions [Approx1], and also some very neat mathematical properties such as “product form.”

It all started when I was visiting the RIACS Research Centre at NASA Ames, in Moffett Field, California in the summers of 1987 and 1988. This was a lot of fun because at lunch time, going out from the back of the laboratory one ended up directly on an airfield where the U2 spy aircraft was taking off. The body of this airplane is very small and thin, with just enough place for one pilot and his instruments and commands, but the wings are very long. In fact, the wings have small wheels at the ends which support them at take-off. If they did not have the wheels, they would be scraping the runway because the wings are long and too heavy to stay in a horizontal position. From Moffett Field, the job of the U2s was, officially, to fly along the U.S. Pacific Coast to try to spot and track Russian submarines. The Director of RIACS at that time was my friend Peter Denning who had just recently left Purdue University where he had been the department head, and had invited me. My official job at RIACS was to work on the performance of parallel processing systems, since I had just published my small monograph on multiprocessor performance.

NASA Ames had some of the largest supercomputers at that time since they were supposed to eventually replace wind tunnels (another specialty at NASA Ames) for the testing of aircraft and

rockets. It is amusing to note that both supercomputers and wind tunnels are “energy-voracious.” NASA Ames, and these were the days before September 11, was that it was supposed to be a very secure facility, and I was a non-resident alien ... what a funny name for a non-U.S. citizen without a Green Card! Therefore I was not allowed to enter the airbase officially and had to work in an external building just at the border of the base.

Erol Gelenbe delivering a keynote speech about the analysis of auction algorithms at the IFIP World Congress 2008 in Milan.

But ... the building’s back door was unlocked, so I could walk onto the tarmac and observe the U2s. I could also just walk over through the back door to RIACS. Anyway, this is just to set the tone about my working environment.

Peter Denning had recruited an interesting man called Pentti Kanerva from Pat Suppes’ entourage at Stanford. Pentti was originally a forestry engineer from Finland who had done a PhD in Philosophy with Suppes. He had invented a model of computation called sparse distributed memory (SDM), which is not too different from the Kohonen maps that you may

know about. SDMs offer a nice adaptation or learning algorithm both for numerical and non-numerical data based on slowly modifying a data representation based on new data that is presented. As a result Penti was also interested in natural neuronal networks, as were some other people who worked at RIACS, so they had organized a series of seminars by prominent neuroscientists. My work on the RNN started in those circumstances. In the meanwhile I also published a paper on the learning properties of SDM, which has remained rather obscure.

As many people of my generation, I was familiar with the McCulloch-Pitts model of neurons, and about the Minsky-Papert controversy concerning non-linear perceptrons. I knew of John Hopfield's model and his results concerning optimization through relaxation, and of the work of the PDP Research Group at San Diego, and the contributions of Dave Rummelhart and Terry Sejnowsky, and about the backpropagation algorithm. At that time, Françoise Fogelman, who joined the Ecole des Hautes Etudes en Informatique in Paris that I had founded, was a strong proponent of these techniques. My former student Andreas Stafylopatis from Athens was also quite interested in these things and we tried our hand at some collective stochastic models for large numbers of neurons [Stafy]. But I felt, after listening to several presentations by neuroscientists, that none of these models actually captured the spiking activity of natural neuronal ensembles, and furthermore (except for John Hopfield's work) the PDP Group's work did not address the important issue of feedback in natural neuronal systems, or "recurrence" as people say in that area.

Filled with all these interesting neuroscience lectures, I set to work upon my return to Paris in September of 1987. Also I had the good luck of being hired by ONERA (French Aerospace Research Organization) as a consultant in AI which was not my area, and I felt obliged to produce something significant.

In six months, I had developed the spiked random neural network model, and obtained its analytical solution, but the people at ONERA could not understand what I was trying to do. The following summer I was back at RIACS, and met Dave Rummelhart who had moved to the Psychology Department at Stanford. I dropped abruptly one day into his office without knowing him personally, and told him what I had done. He was very friendly and interested, and invited me to give a seminar the following week. After the seminar he told me to submit my work to the journal that Terry Sejnowski, Dave and others had recently started, *Neural Computation*, and my first paper was rapidly accepted and then published in 1989.

Several papers have followed over the years [GelFour2],[Stelios1]. Since the journal indicates the name of the handling editors after the papers are published, I owe a debt of gratitude to Dave Cowan and Haim Sompolsky who acted as handling editors, and neither of whom I know

personally to this day. My learning algorithm [RNN-Learn1] came in 1993. It was the first algorithm that established that learning for an N neuron recurrent network is of time complexity $O(N^3)$, while it was well known that the back-propagation algorithm for a feed-forward network is of time complexity $O(N^2)$. In the course of this work, there were applications to imaging [Atalay1], [Atalay2], [Compress1], [Compress2] and adventures and complications related to non-linear mathematics. There have been other applications and extensions, and a return to biology while I was at Duke [Oscill], but that would lead to an even longer story.

CC: Please describe the famous G-networks.

EG: Queueing theory has been around for at least as long as telephone systems have existed. The literature contains many tens of thousands of papers which appear either in publications related to the application domain (e.g., telephony, manufacturing systems, computer systems, the Internet, road traffic, wireless communications, etc.), or in more mathematical journals related to probability theory or operations research. It is a theory based on mathematical probability that considers a dynamical system composed of service centers and customers.

The latter move around the service centres according to a designated probabilistic or deterministic behaviour, and at each service centre a customer waits in line and is then served according to a service discipline, e.g., first-in-first-out, round-robin, last-in-first-out as in a push-down stack, or according to some priority scheme and so on. The service time of each customer in a service centre is typically given in the form of a probability density function or probability distribution, and this will typically differ in each service centre.

In addition, customers may belong to “classes” so that the service time distributions and the routing of customers among different centers, may both depend on the class to which the customer belongs. Often independence assumptions are made about the service times and the routing of different customers even though they may belong to the same class. This is a very useful theory in that it is widely used in industry to design telecommunication systems, manufacturing systems, transportation, parts handling and assembly, etc. When such systems have a steady-state or long term solution in which the joint probability distribution of the number of customers in each of the queues can be expressed as the product of the marginal distributions in each of the queues, despite the fact that the distinct service centre queues are in fact coupled, then we say that the queueing network has “product form.” Examples include the Jackson Networks that Len Kleinrock used in his early work to model packet networks, and the Baskett-Chandy-Muntz-Palacios (BCMP) networks. Product form is a remarkable property that reduces dramatically the time and space computational complexity of using queueing

networks, removing the need to enumerate all possible states, and bringing the computational cost down to a polynomial time and space complexity. Jeff Buzen has made the key contribution to the complexity issue.

Erol Gelenbe enjoying the Bosphorous in Istanbul, from Dolmabahçe Palace in 2009.

G-Networks [G-Net1] extend a network of queues, to include certain new types of customers that can modify the behaviour of others. Thus “negative” customers destroy other customers. For instance they can represent external decisions that are made to reduce traffic because of congestion, or to remove packets in a network that may contain viruses. Triggers are yet another type of customer which can move other customers from one queue to another. Multiple class G-Nets are discussed in [GelFour1],[Gel-Labeled]. Resets [G-Net3] are customers that replenish queues when they are empty, to represent (for instance) a situation where we wish to keep some part of the system busy, or when queue length represents the degree of

reliability so that “replenishment” corresponds to repairing a component. Thus you can think of G-Networks as queuing networks that also incorporate some useful control functions. For instance, while ordinary customers can be packets in a network, these special customers can represent control packets that travel through the network and affect the ordinary packets at certain specific nodes. The link between G-Nets and neural networks is discussed in [G-Net2].

G-Network models, and related aspects discussed by my colleagues Jean-Michel Fourneau and Peter Harrison, lead to product forms. However the solutions obtained, starting with [G-Net1], differ from the earlier Jackson and BCMP networks in that they rely on non-linear traffic equations that describe the flow of customers of different types and classes throughout the network. Because of this non-linearity, one also has to address questions of how and when the solutions one may obtain actually exist and are unique. My first paper on G-Networks was turned down at an ACM-SIGMETRICS conference because the reviewers did not quite believe that new models in this area could be found and also solved analytically. Thus I turned to journals dealing with applied probability, and some of my most cited papers are in this strange new area which has attracted much attention over the last 20 years.

References

- [Automata] E. Gelenbe “On languages defined by linear probabilistic automata,” *Information and Control*, 16: 487-501, (1970).
- [Paging] E. Gelenbe “An unified approach to the evaluation of a class of replacement algorithms,” *IEEE Trans. Computers*, Vol. C-22 (6): 611-618, (1973).
- [Life-time] E. Gelenbe “The distribution of a program in primary and fast buffer storage,” *Comm. ACM*, 16 (7): 431-434 (1973).
- [Diff1] E. Gelenbe “On approximate computer system models,” *Journal ACM*, 22(2): 261-269 (1975).
- [Phillips1] E. Gelenbe, J. Boekhorst, J. Kessels, “Minimizing wasted space in partitioned segmentation,” *Comm. ACM*, 16 (6): 343-349, (1973).
- [Phillips2] E. Gelenbe, P. Tiberio. J. Boekhorst “Page size in demand paging systems,” *Acta Informatica*, 3 (1): 1-23 (1974).
- [PF] E. Gelenbe, R. Muntz “Probabilistic models of computer systems - Part I,” *Acta Informatica*, 7: 35-60 (1976).
- [Diff2] E. Gelenbe “Diffusion approximations: waiting times and batch arrivals,” *Acta Informatica*, 12: 285-303, (1979).
- [CP] E. Gelenbe “On the optimum check-point interval,” *Journal ACM*, 26 (2): 259-270, (1979).

- [Iasno] E. Gelenbe, R. Iasnogorodski "A queue with server of walking type (autonomous service)," *Annales de l'Institut Henry Poincaré, Série B, XVI* (1): 63-73, (1980).
- [Hebrail] E. Gelenbe, G. Hebrail "A probability model of uncertainty in data bases," *Proc. Second International Conference on Data Engineering*, pp. 328-333, Los Angeles, CA, Feb. 5-7, IEEE Computer Society, (1986), ISBN: 0-8186-0655-X.
- [CRAS] E. Gelenbe « Réseaux stochastiques ouverts avec clients négatifs et positifs, et réseaux neuronaux," *C.R. Acad. Sci. Paris, t. 309, Série II*: 979-982, (1989).
- [RNN] E. Gelenbe "Random neural networks with positive and negative signals and product form solution," *Neural Computation*, 1 (4): 502-510 (1989).
- [Stafy] E. Gelenbe and A. Stafylopatis "Global behavior of homogeneous random neural systems," *Applied Mathematical Modelling* 15 (10): 534-541 (1991).
- [G-Net1] E. Gelenbe "Product form queueing networks with negative and positive customers," *Journal of Applied Probability*, 28: 656-663 (1991).
- [Atalay1] V. Atalay, E. Gelenbe, N. Yalabik "The random neural network model for texture generation," *International Journal of Pattern Recognition and Artificial Intelligence*, 6 (1): 131-141 (1992).
- [Atalay2] V. Atalay, E. Gelenbe "Parallel algorithm for colour texture generation using the random neural network model," *International Journal of Pattern Recognition and Artificial Intelligence*, 6 (2 & 3): 437-446, (1992).
- [RNN-Learn1] E. Gelenbe "Learning in the recurrent random network," *Neural Computation*, 5: 154-164 (1993).
- [G-Net2] E. Gelenbe "G-networks: An unifying model for queueing networks and neural networks," *Annals of Operations Research*, 48 (1-4): 433-461 (1994).
- [GelFour1] J.M. Fourneau, E. Gelenbe, R. Suros "G-networks with multiple classes of negative and positive customers," *Theoretical Computer Science*, 155: 141-156 (1996).
- [Compress1] E. Gelenbe, M. Sungur, C. Cramer P. Gelenbe "Traffic and video quality in adaptive neural compression," *Multimedia Systems*, 4: 357-369 (1996).
- [Gel-Labed] E. Gelenbe, A. Labed "G-networks with multiple classes of signals and positive customers," *European Journal of Operations Research*, 108 (2): 293-305, (1998).
- [Oscill] E. Gelenbe, C. Cramer "Oscillatory corticothalamic response to somatosensory input," *Biosystems*, 48 (1-3): 67-75 (1998).
- [GelFour2] E. Gelenbe, J.M. Fourneau "Random neural networks with multiple classes of signals," *Neural Computation*, 11 (4): 953-963 (1999).
- [Approx1] E. Gelenbe, Z.-H. Mao, Y.-D. Li "Function approximation with spiked random networks," *IEEE Trans. on Neural Networks*, 10 (1): 3-9, 1999.
- [Compress2] C. Cramer, E. Gelenbe "Video quality and traffic QoS in learning-based sub-sampled and receiver-interpolated video sequences," *IEEE Journal on Selected Areas in Communications*, 18 (2): 150--167, (2000).
- [G-Net3] E. Gelenbe, J.M. Fourneau "G-Networks with resets," *Performance Evaluation*, 49: 179-192, (2002).

[Users] E. Gelenbe “Users and services in intelligent networks,” *Proceedings IEE (ITS)*, 153 (3): 213-220, (2006).

DOI: 10.1145/ 1895419.1899474