



## Practical Theories Make the World Go (Part II) *An Interview with Erol Gelenbe*

### Editor's Introduction

*This interview is the second of two parts of an interview of Professor Erol Gelenbe by Professor Cristian Calude. It appeared in print in the October 2010 issue of the Bulletin of the European Association for Computer Science (<http://www.eatcs.org/index.php/eatcs-bulletin>).*

*Peter J. Denning*  
*Editor*

## Practical Theories Make the World Go (Part II)

### *An Interview with Erol Gelenbe*

**Cristian Calude for Ubiquity:** Tell us about the design of the first random access fibre-optics local area network.

**Gelenbe:** This was a very interesting experience. In the mid 1970's, thanks to Louis Pouzin who is one of the pioneers of the Internet, and an extremely sharp and amusing individual, I was put in contact with the group developing the Arpanet. In particular I met Bob Kahn in Washington. At that time, a new packet communication scheme using satellites had been devised: the ALOHA Network, which was implemented by Norman Abramson at the University of Hawaii. Of course, ALOHA is the “father” of the Ethernet. Abramson and Kahn had published papers that described the scheme and computed its maximum throughput; Leonard Kleinrock and his students were also studying the problem. I felt that the initial models addressed steady-state analysis, in a context where the steady-state might not exist because the system was intrinsically unstable. Together with Guy Fayolle and Jacques Labetoulle, we obtained a strong result, which after some delay managed appeared in the Journal of the ACM [Aloha2] proving that the slotted random access communication channel (i.e. known as “slotted ALOHA”) was intrinsically unstable due to potential simultaneous transmissions between uncoordinated transmitters, and that it could be stabilised and even optimised under a “1/n” policy which was to retransmit previously collided packets at a rate that is inversely proportional to the number of backed-up transmitters. Strong results may sometimes upset your colleagues. But Bob Metcalfe, who invented Ethernet, was very positive about this work, as he wrote a few years ago to Jeff Buzen, his PhD advisor at Harvard.

This work had started while I was at the University of Liège and at INRIA, and then I had moved to Orsay where, with Jean Vuillemin and Jean-Claude Bermond, we founded the LRI, Laboratoire de Recherche en Informatique. At Orsay, I told Wladimir Mercouroff, a senior member of the university, that this work could have practical applications to local area communications. He suggested funding via the DGRST (Délégation Générale à la Recherche Scientifique et Technique) jointly with a company called La Calhène, to build a fiber optics local area communication system for environments with strong electromagnetic perturbations. I would have been happier to use coaxial technology but the funding agency favoured fiber optics. So we ended up building a prototype called Xanthos, which used DEC-LSI11 processors as access nodes, and fiber optics for transport with the random access protocol using our optimal control algorithm with a scheme we had devised to estimate deviation from optimality based on the frequency of the fiber channel's “silent” periods [Aloha3]. Once the system was up and running, I presented it to the French Telecommunications authority for possible commercialisation. They told me that this work was of academic interest, but because we were using random access, we could only guarantee delivery times on average and in probability, rather than with fixed maximum delays; being rather naive at the time, I believed them. So the

project was set aside. A couple of years later Ethernet appeared, and I am sure that some French Telecom people were biting their nails. As I said, Bob Metcalfe knows this story.

As a consolation prize, the French Telecom hired me as a consultant for a few years and I was able to do several other things for them, but they (and I) missed out on a major opportunity. What happened to Louis Pouzin and his team at INRIA, for similar reasons and regarding the Internet as a whole, is a far more tragicomic story.

**Ubiquity:** You patented an admission control technique for ATM networks.

**Gelenbe:** This was an application of my earlier theoretical papers on diffusion approximation models for queueing systems (which was rejected for a prestigious French conference). My results then appeared in the Journal of the ACM [Diff1] and Acta Informatica [Diff2] and were motivated by the need to simplify the calculations of queue lengths, server utilisation and so on, when you have “non-exponential” assumptions. Diffusion approximations and Brownian Motion are well known, and there is even a wonderful little book on the subject by Albert Einstein. This approach had been suggested to approximate road traffic congestion by G.F. Newell (Berkeley), and then by Hisashi Kobayashi (IBM) for computer system performance. My original contribution introduced a mixed discrete-continuous model which could address deal “low traffic” conditions which had been ignored in earlier work. This gave rise to mixed differential and partial differential equations which I was able to solve in “closed form”.

In the mid 1990's, IBM was designing its N-Way switch for ATM (Asynchronous Transfer Mode) Networks. The system design was being carried out at Raleigh (North Carolina) near Duke University where I was Department Head; the hardware was being designed at IBM's La Gaudé Laboratory. The fashionable approach at that time to admission control was to use “large deviations” whose originator, Varadhan from the Courant Institute was, this year, like me, elected to the Hungarian Academy of Sciences. But large deviations only provide “order of magnitude” estimates of packet or cell loss, which is the primary metric of interest in ATM. I was awarded a contract by IBM-Raleigh to look at the problem and we developed an algorithm that used the predictions of my model [Mang1], [Mang2] to decide whether to admit a new flow into the network, based on the model's predictions for packet or cell loss. One of my students, Xiaowen Mang (now at AT&T Labs), performed some simulations. It all seemed to work well, so we patented the technique together with IBM engineers Raif Onvural and Jerry Marin. The US Patent was awarded in 1998 or 1999. Links to these ideas can be found in my recent work on packet travel time in wireless networks [Travel].

**Ubiquity:** What is the “cognitive packet network” routing protocol?

**Gelenbe:** Let us call it CPN [CPN1],[CPN2] to make things simpler. It is an algorithm that can be run on network routers within an IP (Internet Protocol) network or most other networks (including sensor and ad hoc networks). CPN adaptively chooses paths which have certain desirable properties such as better delay or loss characteristics, lower energy consumption,

lower economic cost, greater security, or a combination of such criteria. The criteria are incorporated into a “goal” or objective function, whose instantaneous value is established based on measurements which are collected with the help of “smart packets” that constantly search the network for paths which provide the best possible value of the goal. Since CPN is based on on-line measurements, it responds to observations which are being made and which will in general change with time so that the choices that CPN makes also change. CPN offers the end user the possibility to make the choices, but the end user may delegate the decisions to an agent which manages its access to the network, or which may manage the choices that need to be made for several end users. CPN uses two mechanisms; the first is the “smart packets” which act as scouts and collect and bring back measurements, as already mentioned. The second is the use of recurrent random neural networks (RNNs) which are installed in those routers that take part in CPN's decisions. The RNNs act as oracles; the excitatory and inhibitory connections of these RNNs are updated using the reinforcement learning rule based on the goal function, as a result of the measurements constantly collected by the smart packets. The RNNs are used to route the smart packets (i.e. to inform the ongoing “search”), while all the resulting measurement information concerning the goal is returned to the end-user or to its decision agent. The decision agent then uses the advice it receives, so as to select the best paths in the network. For instance, the decision agent may wish to reduce the frequency with which it makes changes in paths so as to avoid needless oscillations; in that case it may decide to change a path only when the estimated benefit is high. CPN has been implemented in several wired and wireless network test-beds. It has also been considered as a means to direct people in crowded environments, or in emergency situations.

**Ubiquity:** You have collaborated with the telecommunication and computer industry in various capacities. How useful/relevant are theoretical results for this industry?

**Gelenbe:** I think that the value of theory in our field, when it is based on realistic assumptions and sound evaluation, lies in its ability to provide tremendous short-cuts that avoid a lot of tedious work based on experimentation and testing. My first inroads into the telecommunications industry were related to the performance evaluation of the E10 electronic switch in the late 1970s. The E10 was in fact a large scale computer, together with electronic equipment, that was used to connect and automatically manage large numbers of telephone calls. It was to replace the previous “dumb” electronic switching systems. The French Telecom research centre CNET had been involved in trying to evaluate whether the E10 was performing up to specification, and they were relying on simulations which were taking orders of magnitude longer than the time it took the E10 system itself to execute the corresponding task. The team studying this was at the end of their tether, and the team leader finally had a (real) nervous breakdown. Together with my PhD student Jean Vicard we stepped in and within six months we had a mathematical model based on queueing networks which was quite accurate and which could be solved in seconds of computer time, rather than hours or days of simulation time. Thanks to this work, I continued being funded by CNET for nearly twenty years and they hired several of my former PhD students. This also explains that many of my former PhD

students now teach in France at schools such as Institut National des Télécommunications and École Nationale des Télécommunications.

**Ubiquity:** In France in 1982 you designed and implemented a national vocational training programme in computer technology called the “Programme des Volontaires pour la Formation à l'Informatique.”

**Gelenbe:** This was a very interesting experience. At the suggestion of Jacques Gualino, who was one of INRIA's managing staff, I started working with the people who had launched the “Centre Mondial pour l'Informatique” in Paris, Jean-Jacques Servan-Schreiber, Nicholas Negroponte and Seymour Papert. The latter two wanted to help the third world via personal computers, while Jean-Jacques was actually (I think) on a mission to transform the French bureaucracy through a greater use of Information Technology and to attain some form of political power or political role in the process. The year was 1982, soon after the elections that had brought François Mitterand and the Socialist Party to power, so there was an opportunity to make some changes -- but the question was what this Centre Mondial should do. While these people aimed at lofty and global goals, I decided to tackle a relatively small project. I felt that much of vocational education in France was obsolete and essentially taking poorly advised teenagers and turning them into disgruntled unemployed young adults, simply because vocational education was being dispensed by obsolete technical educators who “trained” young people to operate obsolete equipment for jobs that did not exist. On the other hand, both industry and the service sector were looking for people who had some simple computer education that could be used in technical and service jobs. However, instructors who were knowledgeable computer scientists and engineers were just too expensive. Furthermore computer equipment was scarce and expensive. In conversations with Jean-Jacques Servan Schreiber, Pierre Lafitte (then President of the Ecole des Mines in Paris) and myself, we came up with the idea of using newly graduated engineers who could do a “vocational education” service for youngsters instead of their military service where they were often getting very bored. Though conceptually simple, the whole programme had to be “engineered”, which I did, so that several hundred young graduate engineers could do this new form of civilian service instead of going into the military for one year. Several ministries had to be convinced. Several million francs for equipment, wiring and room security (against computer theft) were needed to get started, and the network of training centres for unemployed young people had also to be incorporated into the task. The long and short of it is that we were successful: up to seven hundred young graduating engineers and computer scientists got involved in this programme each year. Personal computers from a variety of sources were purchased and installed in small groups of ten PCs per training centre. Tens of thousands of young unemployed people were trained and many entered the job market successfully. For the first two years I ran the programme and collected detailed statistics about what was happening. Then the programme was taken over by existing social and government bodies. It came to an end towards 1989 after Jacques Chirac became Prime Minister and some things introduced by the Socialist government were abrogated.

**Ubiquity:** You also served as Science and Technology Advisor to the French Minister of Universities, and member of the Executive Board and of the Science and Technology Board of the Data and Information Fusion Defence Technology Centre in the UK and chaired the Technical Advisory Board of the US Army Simulation and Training Command.

**Gelenbe:** The VFI Programme that I discussed before attracted the attention of the then French Minister for Universities, the law professor Roger-Gérard Schwartzenberg. Normally I should never have been in his group of advisers: I was a recent immigrant who had not studied at a Grande École. All of his other advisers, except the three political appointees from his party (who had studied at Sciences Politiques -- Sciences Po in Paris) had either studied at ENA (École Nationale d'Administration) or at École Polytechnique. There was a Professor of Medicine who advised the Minister for the medical side of things, and then this strange individual: me. Some of the other people in his group of advisers and higher civil servants in the ministry obviously thought I was totally out of place. But I was a young professor from the best science campus in France, Orsay and I taught part-time at École Polytechnique, so I could not be all that stupid. I spoke and wrote French well, and the Minister's parents had been immigrants, so he was open minded. I was offered the job because the VFI programme had shown that I was a "go getter", able to handle large projects and deal with the arcane administration.

Sure enough, in this new job I did bring together another large project. The year was 1984, and most students in French Universities and Grandes Écoles did not have an introductory course in Computer Science. The barriers were the usual ones: where to find the lecturers, where to find the computers, where to find the space with appropriate electrical and network connectivity as well as security, and what to teach. The whole funding issue needed to be dealt with because we needed to buy eight hundred PCs and install them in groups of eight or ten, and the installation itself would cost quite a bit of money. We had to open many junior faculty positions for the teaching. We set up working groups to design the course material by group of disciplines: math-physics, biology, medicine, economics, humanities, and so on. I will not dwell on the details, but it worked out well, and also made many of my colleagues unhappy because they thought they could have used the resources much better for Computer Science, without realising that we could not have attracted such a large investment for the benefit of a single discipline, and that the equipment and new faculty were in themselves an investment in research as well.

In my role as advisor, I had several other jobs to do: monitor and expand the different scientific disciplines that I was dealing with, monitor the engineering schools, deal with the transformation of the École Normale Supérieure (a very interesting story to tell separately some day), the transformations of PhD programmes and so on. An initiative which proved very useful and lasting was the "Arrêté" which has allowed some of the Grandes Ecoles to deliver the PhD degree: it has significantly increased the research activities at many of the elite schools in France. At the end of two years of doing this job from 8am to 8pm, I was totally exhausted. I was delighted (!) when in the Spring of 1986 the Socialist Government lost the elections and I

could leave this harassing activity to return to my lab. In those two years I did manage to write one or two decent papers and to graduate some PhD students, but it was tough.

You mentioned the Technical Advisory Board (TAB) of the US Army Simulation and Training Command. I spent 1993 to 1998 at Duke University as Head of Department, and then moved to Orlando, Florida. There I was the Director of the School of Electrical Engineering and Computer Science (SEECs) from 1998 to 2003, which I founded by merging several programmes at the University of Central Florida. I was in contact with the neighbouring modeling and simulation facilities for the US Army where much of the training is heavily computerised. I was asked first to sit on this TAB for a year, and then to chair it for four years until I moved to Imperial College. This offered possibilities for interaction with a sophisticated organisation in areas such as virtual and enhanced reality, games, simulation, networking and distributed computing. At UCF I was also Associate Dean of Engineering and headed an organisation with a total of 2200 students, nearly 100 instructors, three Master's and three PhD programmes, and four distinct undergraduate degree programmes in Computer Science, Computer Engineering, Electrical Engineering and Information Technology. In the five years I was there, we secured \$15 Million for a new building, and it was nice to lead its design team. I visited it in April 2010, and enjoyed seeing my "creation".

In 2003, when I joined Imperial College, the UK Government had decided to transfer much of its Defence related research (except for the top secret stuff) to Universities. I was one of the writers of a proposal to start a research centre around Imperial College joining with other universities and industry, and a budget of £10 Million per year, with half of it from three companies: BT, General Dynamics UK Ltd and QinetiQ, which ran was until 2009. Since I was involved in the inception, I became a Member of the Executive and of the Science Board of this Data and Information Fusion Defence Technology Centre. It meant weekly travel around the UK, and was a good way for me to "immigrate" more rapidly. Now that it has ended, I also appreciate the opportunity to spend more time on own research.

**Ubiquity:** What gives you more pleasure in academic work?

**Gelenbe:** I think that we are so lucky to have such a fun job. That's probably why we are not good at getting a decent salary in relation to the number of hours that we put into it. Exercising my curiosity and learning new things, being able to talk to experts about subjects that are new to me, being a bit of clown when I lecture, and enjoying the interest of young people, these are some of the things I really enjoy in my work.

**Ubiquity:** Did you ever miss a target? How did you cope?

**Gelenbe:** I miss targets all the time, and one reason is that I am rather dispersed as to the subjects that interest me. Right now they are Computer Networks, Gene Regulatory Networks [Gene], Viruses in nature and in computers [Virus], Economics [Auctions], Synthetic Chemistry [Molecule], and a few other things ! Thus I have given up pursuing conference deadlines. I try to



publish papers on my work in the best relevant journals, and to benefit from serious refereeing and criticism: referees are my best teachers these days.

**Ubiquity:** Many thanks!

## References

[Aloha2] G. Fayolle, E. Gelenbe, J. Labetoulle “Stability and optimal control of the packet switching broadcast channel”, *Journal ACM*, 24(3): 375-386 (1977).

[Aloha3] Binh Tri An, E. Gelenbe, “Simulation et amélioration d'un canal de diffusion sous contrôle”, *R.A.I.R.O. (Informatique)*, 11 (3) : 301-321 (1977).

[Mang1] E. Gelenbe, X. Mang, R. Onvural “Diffusion based Call Admission Control in ATM”, *Performance Evaluation*, 27 & 28: 411-436 (1996).

[Mang2] E. Gelenbe, X. Mang, R. Onvural “Bandwidth allocation and call admission control in high-speed networks”, *IEEE Communications Magazine*, 35 ( 5): 122-129, (1997).

[CPN1] E. Gelenbe, “Cognitive Packet Network”, U.S. Patent 6,804,201, Oct. 11, 2004.

[CPN2] E. Gelenbe, M. Gellman, R. Lent, P. Su “Autonomous smart routing for network QoS”, *Proc. First International Conference on Autonomic Computing*, IEEE Computer Society, 232 - 239, Washington D.C., (2004).

[Gene] E. Gelenbe “Steady-state solution of probabilistic gene regulatory networks”, *Physical Review E*, 76(1), 031903 (2007).

[Virus] E. Gelenbe “Dealing with software viruses: a biological paradigm”, *Information Security Technical Reports* 12: 242-250, Elsevier Science, (2007).

[Travel] E. Gelenbe “A Diffusion model for packet travel time in a random multi-Hop medium”, *ACM Transactions on Sensor Networks*, 3 (2), (2007).

[Molecule] E. Gelenbe “Network of interacting synthetic molecules in equilibrium”, *Proc. Royal Society A (Mathematical and Physical Sciences)* 464:2219--2228, (2008).

[Stelios1] E. Gelenbe and S. Timotheou “Random neural networks with synchronized interactions”, *Neural Computation* 20: 2308--2324, (2008).

[Auctions] E. Gelenbe “Analysis of single and networked auctions”, *ACM Trans. Internet Technology*, 9(2), (2009).





[CACM09] E. Gelenbe “Steps toward self-aware networks”, Communications ACM, 52(7):66-75, (2009).