

People that were important during my career

Gregor v. Bochmann (written in May 2016, just before my retirement)

The following pages represent my memories about my research career and the people that were important for me during these activities. I would like to thank, at this point, all my collaborators for their contributions and interesting discussions. I must admit that the following account is not complete, and I want to apologize that many names of collaborators and students are missing in the following text.

My university studies started in 1961 and I was enrolled in the physics program at different universities in Germany. In 1964, I followed a course on nuclear physics at the University of Grenoble with Prof. Depommier. During the following summer holidays, I planned to work in his group at CERN (Geneva), however, I got a hepatitis infection during a trip to Turkey (with the student orchestra of the University of Grenoble, earlier that summer). I met Prof. Depommier later at the University of Montreal where he worked in the physics department.

1967-68, I did my Master project (Diplom-Arbeit) at CERN in the group of the “myon g-2 experiment”. The group leader was Emilio Picasso - he was a kind of “father” figure for the whole group, also for me. During that time, F.J.M. Farley (from Readings, UK) showed me how to program the PDP-8 computer in machine language.

After obtaining my Master degree, I moved to Canada. Prof. B. Margolis was my PhD supervisor at McGill University from January 1969 through summer 1971. I worked on simulation studies of high-energy reactions of particles on nuclei, writing these simulation programs in Fortran and PL/1.

During the year 1970-71, I gave computer science courses at the Maisonneuve College in Montreal. Pierre Cliche was a colleague of mine with whom I had many exchanges. He explained me the organization of the college courses etc. and also explained me the political developments going on in Quebec (from the separatist perspective).

During the year 1971-72 with my post-doctoral scholarship, I started to get familiar with some fields of computer science at the University of Montreal. Prof. Bill Armstrong invited me to collaborate with him on adaptive Boolean networks (nowadays called neural networks) and applications to character recognition. Later I participated with him in discussions about commercialization of a patent he had on this topic. At the end of the year, I became an assistant professor in that department.

From my first years at the University of Montreal, I remember other colleagues with whom we had many discussions and social meetings: Jean Vaucher, Stan Selkow (who later moved to the US), Paul Bratley, Neil Stewart, Edward Cerny, Jan Gecsei, Alex Bacopoulos and others. There were also some visiting professors from France: Jean-Pierre Verjus, Laurent Trilling and Vaucoix.

During my first year as a professor, I gave two courses, one on machine language programming and one on formal languages (I had followed a course on this subject the year before, given by Prof. Vaucoix). I also started to collaborate with Prof. Lecarme (who later moved to the Université de Nice in France) on compiler writing systems. My first Master student developed such a system for LL(1) grammars, and I got interested in semantic attributes which led to my first important paper in the Communications of the ACM in 1976 [1].

During those years, I attended various international conferences on different fields of computer science in order to find a field in which I would concentrate my research efforts. In 1973 at a conference on computer communications (a very new field at that time), I met Louis Pouzin (the director of the French computer communications project CYCLADE) who suggested I should study how one could verify that communication protocols are well designed and well implemented. I took the Alternating Bit Protocol as an example and developed a verification method for protocol specifications given in the form of state machines. This led to my second important paper in 1978 [2].

I continued working in this field until the 1990ies, and it is in this field that I made my name as researcher. Further details are given in my 2010 paper entitled “Some notes on the history of protocol engineering” [9].

During those years, I had research collaborations with several of my colleagues. I remember joint publications with Edward Cerny, Jean Vaucher, and with Michel Raynal and Jean-Pierre Verjus from Rennes (France). With Jan Gecsei, I worked on Videotex systems in the early 1980ies; these were precursors of the current Internet-Web (which started in the 1990ies), including the Canadian Telidon system which supported graphics. With Claude Frasson and Jan Gecsei, I worked on “intelligent tutoring systems”, and with Eduard Cerny I had a company in which we developed such a system, however, without commercial success. During the year 1977-78, I was a visiting professor at the École Polytechnique de Lausanne with professor Giovanni Coray, and my course on the “architecture of distributed computer systems” developed into a book that was published in 1979, and an extended version, “Concepts for distributed systems design” was published in 1983.

Starting in 1976, I did some research projects funded by the Department of Communications (DOC, Canada) in the area of standards for computer communications. This led to my involvement in the development of Formal Description Techniques (FDT) for communication protocols and services which were standardized by the International Telecommunication Union (ITU) and the International Organization for Standardization (ISO). André Dantine (University of Liège, Belgium) organized in 1978 the first international conference on computer communication protocols where I presented my paper on state machine protocol verification. There I met, among others, John Day (from the US) who spearheaded the work on FDT within ISO, Colin West (from IBM Zürich) who was developing an automated tool for protocol verification, and Fritz Vogt (from Hahn-Meitner Institute in Berlin) with whom I collaborated in the subsequent years on example applications of FDTs, and with whom I kept a friendship until today.

In the following years, I was funded by Bill McCrum (DOC) to participate in the international standardization activities for FDT. These developments are described in Section 4 in my “history” paper. I chaired the ISO Subgroup A on describing system architectures, but I was mainly involved in Subgroup B for the development of the Estelle FDT language which became an international standard in 1986, and I chaired the first joint ITU-ISO meeting with the objective of joining the development of Estelle with the SDL language that was being developed within ITU (however, this was not successful). SDL became a language much used with software tools for specifying and implementing communication systems. It was later (2003) integrated into the second version of the standard Unified Modeling Language (UML).

Here are some things I remember from those times: Many vivid discussions with Bob Tenney (chair of Subgroup B), and with Stan Budkowski (INT, Evry, France) who developed an interesting software tool for Estelle. Discussions in Subgroup A with Mr. Bachman (known for his work on databases at Honeywell). Many trips to far-away places (for standardization meetings) such as Melbourne (Australia), Catania (Sicily), Paris, etc. During that time, I also worked with a company from Ottawa (Computer Gateways Inc) for the National Library of Canada on the development of an international standard for the interlibrary loan application.

From my research group at the University of Montreal, I remember people such as Daniel Ouimet, Jean-Marc Serre, Michel Maksud, and others. We mainly concentrated on building tools for different versions of the Estelle language and demonstrating its usefulness for building different prototype applications. One example protocol [3] that was often used was the Transport Protocol Type 2, standardized by ISO and ITU in the context of the OSI standardization effort (at that time, TCP and the other Internet protocols were not yet in general use). We also worked on the ASN.1 (OSI) standard which defines a coding scheme for application-level protocols (a precursor of XML). Some students developed a system for automatically generating code for the coding and decoding of these ASN.1 messages.

Besides protocol verification and implementation, there was also the topic of testing. This topic, how to test protocol implementations for conformance to the given protocol standards, was promoted in the standardization context since 1980 by Dave Rayner (UK). It was of much interest to the industrial community. Rayner organized in 1981 the first conference on Protocol Specification, Testing and Verification, which became a yearly meeting at which I presented much of our research results. My PhD student Behcet Sarikaya was the first person to study the development of tests for protocol implementations from a state-machine description of the protocol [4].

I spend my second sabbatical year 1986-87 (part time) with Siemens in München. I was working in a group that was involved with the development and implementation of protocols for computer communication. During that year, I also did some consulting for Nokia (Finland) and Philips (in Erlangen). This was my first experience within a company. It was interesting.

My first sabbatical year had been very different. I spent that year 1980-81 at Stanford University (south of San Francisco). I was associated with the group of David Luckham and tried to use some tool for theorem proving (but found it much too complicated). I learned about Temporal Logic (which was just invented a few years before) with Susan Owicki and used it in a much cited paper on hardware verification [5]. I also did some consulting with SRI International in the area of protocol specification languages.

In 1986 I stopped my involvement in standardization and concentrated my research effort on building tools and methods for using FDTs for protocol specification, testing and verification. The year before, I had been invited for a 6-week stay with a research group of NTT in Tokyo (Musashino area). My wife Elise joined me there on our first trip to the Far East. Afterwards I had organized a visit through China to different universities in Beijing, Shanghai and Xian (with the help of Prof. Wu from Tsinghua University whom I had met at a conference). Another trip to China, mainly to the Academy of Sciences in Chengdu, followed in 1988. As a follow-up, I had several PhD students from China in my research group, and also several visiting researchers from China and NTT that stayed for a year. The research group became very international. I remember among the visitors: Reinhard Gotzhein (post-doctoral fellow from Germany), Prof. Mme Zhao (from Tsinghua University, Beijing), Mr. He (from the Academy of Sciences of Sichuan), Dr. Masaki Itoh and Mr. Susumu Fujiwara (from NTT, Tokyo), Prof. Higashino (from the University of Osaka).

This was also the time that the Centre de Recherche en Informatique de Montreal (CRIM) was founded and a number of big industry-collaboration projects became available to the university researchers in Montreal. I had been involved setting up the mandate for CRIM, and in the sequel had many collaborative research projects with industry. There were two projects that had a strong impact on my research thinking: A project on network management with BNR (which later was called Nortel, and in the early 2000 became bankrupt) and a project on workflow modeling for business processes with DMR (Montreal). This was the time when “object orientation” was an important buzz word. The project with BNR was interesting through the active participation of the researcher Angelo Bean from BNR. He introduced us to the problems of network management (which was at the same time an important field for standardization) and we developed an object-oriented specification language (called Mondel) and an execution engine for this language. The nature of the collaboration with DMR was similar. We learned about the concepts used for describing business processes and we developed a language and associated tool for evaluating the performance of business process models. A key person in the first project was Michel Barbeau, PhD student (now with Carleton University), and the second project was done in collaboration with my colleague Rudolf Keller (now in Switzerland).

It was probably around 1988 that I got a telephone call from Mr. Decouvreur who was the director of a center for protocol testing recently created in Laval (Montreal) by DOC. He proposed the establishment of an industrial research chair at the University of Montreal, co-funded by NSERC, DOC and the IDACOM company of Vancouver active in the field of protocol testing. The prestigious program of NSERC industrial research chairs was created a few years earlier, and this

was the first such chair at the University of Montreal. The chair was active from 1989 until December 1997 when I left for the University of Ottawa. The organization of the NSERC chair implied a new professor position at the University of Ottawa. During the first year, this position was filled by Pierre Mondain-Monval (who returned afterwards to France). Then Rachida Dssouli joined the chair from the University of Sherbrooke (she was a former PhD student of mine). Supported by the larger funding, my research group grew over the years. Alexandre Petrenko (from Latvia) joined the group as a visiting scientist and also helped with the supervision of our graduate students. Also several post-doctoral fellows joined the group over the years (I remember in particular Gang Luo, who later joined Nortel), and many international visitors came for longer periods, often with their own funding.

The research of the chair was oriented towards the problem of protocol conformance testing, although also other topics were pursued. At the end, I found it a bit tiring to always work on the same topic.

By the way, there were other things going on at the same time. It was the time that federally funded Networks of Excellence were created for the first time. These are associations between research groups at different universities and various Canadian industries that collaborate on a particular research agenda. In 1990, I participated, under the leadership of Maier Blostein (McGill University), in the establishment of the Canadian Institute for Telecommunications Research (CITR). This was a very interesting experience. Initially, I established a collaborative sub-project of CITR with colleagues from the University of Ottawa (Luigi Logrippo and Bob Probert) on FDTs. However, after two years, this sub-project was abandoned due to lack of support from industry. In particular, our partner IBM was pushing for another direction. I had already earlier had some experience with IBM-collaboration through the Center for Advanced Studies (CAS) which was established in Toronto by Jacob Slonim in 1990 (I had participated in the CORDS project with several universities).

Now in 1993, I chose the topic of multi-media applications for a new CITR sub-project and assembled university experts on orthogonal research topics with the common goal of building an advanced application including the results from all these topics: multi-media databases (Tamer Özsu, Edmonton), scalable video encoding (Eric Dubois, INRS, Montreal), multi-media applications (Nicolas Georganas, Ottawa), networking issues (Johnny Wong, Waterloo), and quality of service negotiation (Bochmann). The leadership for this sub-project was later taken over by Johnny Wong (Waterloo), and I looked after the research problems of quality of service negotiation at the application level (with Gecsei, Dssouli, and Brigitte Kerhervé, UQAM). I found this work quite interesting since it was in a new area, different from the concerns of FDTs. It was a very new field of investigation (as the field of protocol verification was when I started to work in it in 1975). With the post-doc Stefan Fischer and the PhD student Abdelhakim Hafid, we wrote a good number of seminal papers [12, 13]. The main issues were related to the quality limitations of the user terminals and the load of the servers for video-on-demand applications.

Then, beginning of 1998, I took advantage of an early retirement plan of the Quebec government (which was followed by the University of Montreal) and took a new job at the University of Ottawa. I still remember that I gave a talk at the engineering faculty in Ottawa and was invited for lunch by (late) Nicolas Georganas (whom I knew from earlier times) and Emil Petriu to an Indian restaurant which now is one block from my home. I got an offer from the Dean by mail when I was on a visit in Japan, and immediately asked for my retirement from the University of Montreal a few days before the deadline.

The University of Ottawa had just decided to create a new department “School of Information Technology and Engineering” (SITE, nowadays called “Electrical Engineering and Computer Science” - EECS) through the fusion of the Computer Science and Electrical Engineering departments - and under the leadership of my colleague Tim Lethbridge an undergraduate Software Engineering program was created which was one of the first in Canada. I was hired by the Dean Gilles Patry who became a year later Vice-President and then President of the university. I remember Tyseer Aboulnasr, who became Dean afterwards, for her thoughtful actions. For instance, she organized a social event, called Faculty Talent Show, where she performed a Karate act, and I played cello solo and in a quartet with another professor and two engineering students.

Concerning my research, I continued the supervision of my PhD students in Montreal and looked around in the Ottawa environment for some new interesting research directions. This was the time of the Internet boom which later led to the tumbling of several telecommunications companies, including Nortel which had been the industrial leader in the Ottawa area. My first PhD student in Ottawa was Khaled El-Fakih whose thesis topic was in the area of state machine testing, for me an old area – but he explored some new avenues. He took a position at a university in the UAE, and we stayed in contact over the years and produced several joint papers on different topics.

I continued work in the field of quality of service negotiation in an outgoing CITR project on “Enabling Technologies for Electronic Commerce”, and I started some new projects with Nortel on teleconferencing systems. At that time, there was also much talk about optical networks in Ottawa. I initiated a collaborative project with colleagues at the Universities of Ottawa and Carleton with funding from NCIT (National Capital Information Technology ??). This was my opportunity for learning about optical networks.

I have learned about many different fields during my career, since I changed my research direction several times over the years. When I started my career in computer science at the University of Montreal, everything except programming was new to me. The first courses I gave were for me the opportunities to learn the particular topic in depth - in order to be able to explain the issues to the students. Later I volunteered to give a database course in order to learn about this topic. This was similar for a course on computer communications at the University of Montreal. A couple of years ago, I studied software requirements analysis for giving a course on this topic.

The creation of the SITE department entailed the planning for a new SITE building. Nicolas Georganas was the instigator for a renewed application to CFI (Canada Foundation for Innovation) for getting matching funds from the federal and Ontario governments for laboratory equipment obtained at reduced prices from various equipment manufacturers. In 2000, I became the leader of a big 5M\$ application, called SAVANT, for three big collaborative labs in the area of “Distributed Collaborative Virtual Environments” (Georganas), “Advanced Software Engineering Research and Training” (Probert) and “Advanced Internet Applications and Systems” (Bochmann). We moved into the new building in 2002.

In 2002 I was invited by David Plant (McGill University) to take part in a funding proposal for an NSERC Research Network called “Agile All-Photonic Networks” (AAPN) with two research themes: “Network architectures” and “Device technologies for transmission and switching”. I became the coordinator for the first theme which included Trevor Hall and Oliver Yang from Ottawa, and Marc Coates, Lorne Mason and Richard Vickers from McGill. We had 5 industrial partners including Nortel. This was a very interesting project. Our work on network architectures was based on a star network technology proposed by researchers from Nortel, including Richard Vickers. During the years of funding until 2008 (and beyond) I spent much time on optical networks. I remember my PhD students Abdelilah Maach, Cheng Peng, Peng He and Hassan Zeineddine, some of whom were co-supervised by Trevor Hall or Hussein Mouftah. When the AAPN project had finished, I worked for a couple of years with James Zhang (a former PhD student of Oliver Yang) in the area of optimization of network routing.

CANARIE is an organization that provides advanced communication infrastructure to Canadian universities and industries for the development of advanced applications. In 2002, the scientific leader of Canarie, Bill St. Arnaud, and Michel Savoie from the Communications Research Center (CRC) in Ottawa and his team (Jing Wu, Hanxi Zhang, Scott Campbell and Mathieu Lemay from Montreal) promoted the idea of “User-controlled light-path provisioning” (UCLP). As a researcher from the university, I became the official leader of a project with CRC on UCLP funded by Canarie. This was also a very interesting project mainly because of the advanced software technologies that were used to implement the UCLP concepts, such as Java Jini, Web Services and BPEL, which were very new at that time. The developed software was used within the Canarie network, and it was also used for certain demonstrations of international inter-connections between high-speed optical networks of different countries and between continents. A second project on this topic was funded in 2005.

Besides these big projects, I also came back to two research themes that have followed me through much of my life, namely “Submodule construction” and “Protocol derivation from service specifications” (for explanation, see Section 3.3.5 in my History paper). I wrote my first paper on submodule construction in 1980 with (late) Philip Merlin (from Haifa, Israel) [6, 7]. Later I came back to this problem with several of my PhD students (Z.P. Tao and J.Drissi). Around 2000, I attended a lecture on stochastic relational databases and noticed that the formula used for solving a particular problem had the same general form as the formula used for submodule construction. In

discussions with Nina Yevtushenko (from Tomsk University, Russia), I had already learned that the formulas for submodule construction look very similar for synchronous and asynchronous state machines. This led me to the idea that the submodule construction is a general concept with different manifestations in different compositional contexts. Finally I found a formulation of this problem in logic and a solution which can be used to derive the different submodule construction formulas that were known in different contexts (paper of 2010) [8]. The correctness proof of the solution in logic is very simple, which simplifies the correctness proofs for the more specialized construction formulas.

During my last sabbatical year in Oldenburg (Germany), I studied an application of submodule construction to controlling cars when overtaking on the highway (in collaboration with Ernst-Rüdiger Olderog, Martin Hilscher und Sven Linker), and I tried to find a simple algorithm to solve this problem with liveness constraints – thus coming back to temporal logic which I had studied first at Stanford in 1980.

The other theme, “Protocol derivation from service specifications”, came into my mind during the research visit at NTT in Tokyo in 1984 and was inspired by some notational approach similar to sequence diagrams, proposed by Haruhisa Ichikawa. Later, back in Montreal, I elaborated, with post-doc Reinhard Gotzhein and PhD student Ferhat Khendek, a framework and algorithm for deriving protocols [10]. Around 6 years later, this was extended by my PhD student Christian Kant in collaboration with Teruo Higashino, and in the early 2000, it was extended in a different direction with my former PhD student El-Fakih and colleagues from Osaka (Teruo Higashino and Hirozumi Yamaguchi). I also spent two months in Prof. Higashino’s group in 2002.

During that time, I got together with several Ottawa colleagues interested in “Requirements-driven development of distributed applications” and prepared an NSERC Strategic funding proposal which we got for three years (with Daniel Amyot, Bob Probert, Luigi Logrippo, Amy Felty and Stéphane Somé). My work in this context was related to protocol derivation. We also had the visit by Rolv Braek (Technical University of Trondheim, Norway) who worked with his PhD student Humberto Castejón on understanding how to avoid race conditions in distributed applications. I visited Trondheim in 2006 for a month, and we came up with a high-level notation, similar to Activity Diagrams, to describe requirements for collaborations in distributed systems which allows for “weak sequencing” [11]. I extended my protocol derivation algorithm to the case of weak sequencing. This was implemented in a tool by my Master student Fedwa Laamarti and applied to Web Services protocols by Mustafa Nasser Mousa Faleh. The NSERC Discovery grant which I received for the coming 5 years also is oriented towards this research theme.

During the last twenty years, I was involved in many different research directions (too many?). I mentioned already protocol and software testing, quality of service negotiation for multimedia applications, resource allocation in optical networks, and the two long-time themes of submodule construction and protocol derivation. Furthermore, during my time in Ottawa, I got involved in several other research areas. In relation with multimedia applications, I got interested in security

protocols for authenticating mobile users, and I worked with Carlisle Adams and Morad Benyoucef (Telfer School of Business) on trust management. I also got involved in the design of peer-to-peer networks together with post-doc Shah Asaduzzaman and PhD student Ying Qiao (addressing geographical information retrieval and load balancing). During the last five year, I participated in a research project funded by NSERC and IBM on understanding Rich Internet Applications (RIA) with my colleague Guy Vincent Jourdan and Vio Onut from IBM. This was a very interesting project due to the participation of Vio Onut in the day-to-day discussions about the research progress and the close interactions with his IBM team. Our project obtained several prizes at the yearly CASCON conferences in Toronto, and many patents came out of our work. Initially, the research topic was recovering the structure of JavaScript applications through crawling in view of security and accessibility testing, while later other related topics came to the forefront of our attention.

Overall, I must say that I do not regret to have worked in so many different fields. I always found it interesting to learn something new and try to apply my intellect to solve problems in these different areas.

A few references:

Note: The copy links do not work. Please look at "List of all publications" on my web site.

- [1] **Semantic evaluation from left to right** (G. v. Bochmann), Comm. ACM 19, pp. 55-62 (1976) -- [copy](#) -- [link](#)
- [2] **Finite State Description of Communication Protocols** (G. v. Bochmann), Computer Networks, Vol. 2 (1978), pp. 361-372 -- [copy](#)
- [3] **Specifications of a simplified Transport protocol using different formal description techniques** (G. v. Bochmann), Computer Networks and ISDN Systems, Vol. 18, no.5, June 1990, pp. 335-377 -- [copy](#) -- [link](#)
- [4] **Synchronization and specification issues in protocol testing** (B. Sarikaya and G. v. Bochmann), IEEE Trans. on Comm., COM-32, No.4 (April 1984), pp. 389-395; russian translation: Express Information (overview of western publications), Information Transfer, 1985, no. 28 -- [copy](#) -- [link](#)
- [5] **Hardware specification with temporal logic - An example** (G. v. Bochmann), IEEE Trans. Computers C-31, No.3 (March 1982), pp.223-231 -- [copy](#) -- [link](#)
- [6] **On the construction of communication protocols** (G. v. Bochmann and P. M. Merlin), ICCS, 1980, pp.371-378, reprinted in "Communication Protocol Modeling", edited by C. Sunshine, Artech House Publ., 1981; russian translation: Problems of Intern. Center for Science and Techn. Information, Moscow, 1981, no. 2, pp. 146-155
- [7] **On the Construction of Submodule Specifications and Communication Protocols** (P. Merlin and G. v. Bochmann), ACM Trans. on Programming Languages and Systems, Vol. 5, No. 1 (Jan. 1983), pp. 1-25 -- [copy](#) -- [link](#)
- [8] **Using logic to solve the submodule construction problem** (G. v. Bochmann), Journal on Discrete Event Dynamic Systems, Vol. 23 (1), Springer, March 2013, pp. 27-59. -- [copy](#) -- [link](#)

[9] **Some notes on the history of protocol engineering** (G. v. Bochmann, D. Rayner and C. H. West), Computer Networks journal, 54 (2010), pp 3197–3209. -- [copy](#) -- [link](#)

[10] **New results on deriving protocol specifications from services specifications** (F. Khendek, G. v. Bochmann and C. Kant), Proc. SIGCOMM'89, July 1989, in Computer Communications Review Vol.19 no.4, pp. 136-145 -- [copy](#) -- [link](#)

[11] **On the realizability of collaborative services** (H. N. Castejòn, G. v. Bochmann and R. Braek), Journal of Software and Systems Modeling, Vol. 10 (12 October 2011), pp. 1-21. -- [copy](#) -- [link](#)

[12] **Cooperative quality of service management for multimedia applications** (S. Fischer, A. Hafid, G. v. Bochmann and H. d. Meer), Proceedings of the 4th IEEE International Conference on Multimedia Computing and Systems, Ottawa, Canada, June 1997, pp. 303-310. -- [copy](#) -- [link](#)

[13] **Some principles for quality of service management** (G. v. Bochmann and A. Hafid), Distributed Systems Engineering Journal 4 (1997), pp. 16-27. -- [copy](#) -- [link](#)