Algorithms & Power Systems Architecture:

Using Historical Analysis to Envision a Sustainable Future

Oral History Project

Interviewee: John Undrill

Interview Date: July 16, 2025

Place: Online

Interviewer(s): Julie Cohn, Daniel Molzahn, Sairaj Dhople

Transcriber: Temi

Keywords: John Undrill, University of Canterbury, New Zealand, General Electric Company (GE), Power Technology Incorporated (PTI), Electric Power Consultants, hydroelectric power plant (hydro plant), radial transmission, power system dynamics, Leon Kirchmayer, Ontario Hydro, Manitoba Hydro, Itaipu Dam, Paraña River, voltage control, voltage regulation, analog computer, digital computer, Ōhakuri Dam, Waikato River, Woodward governor, stability limit, hydro-governor, laptop computer, IBM 7094 computer, IBM 1620 computer, Monroe calculator, Lyapunov theory, Nyquist method, Bode method, model, EAI TR10 analog computer, transfer function, Prime 400 computer, Power System Simulator [for Engineering] (PSS-2, PSSE), electromagnetic transients [program] (EMT, EMTP), Positive Sequence Load Flow (PSLF), Power System Computer Aided Design (PSCAD), load flow, GE 400 computer, Hewlett Packard 2120 computer, IBM PC AT computer, UNIX, computer memory, dynamic network analyzer, Newton-Raphson, visualization, Roger Clayton, John Burns, Western Electricity Coordinating Council (WECC), Charles Concordia, Paul de Mello, Gordon Carter, Len Garver, Gabriel Kron, Harvey Happ.

***Abstract*:**

In this interview, Dr. Undrill offers reflections on his academic and work experiences in the electric power industry. He begins with a description of his education at the University of Canterbury in New Zealand, describing his early and continuing interest in the dynamics of power system. He mentions his positions at several private companies and describes relationships with academic institutions

throughout those years. He discusses development of a dynamic model of the Ōhakuri hydroelectric plant using a vacuum-tube analog computer. He mentions directly testing models on the New Zealand hydro plants for accuracy. He recalls good fortune in being able to learn mathematical analysis on the university side and have hands on experience on the industry side. He discusses the risks involved in running tests directly on hydro plants. He mentions working with the woman whom he eventually married. He talks about tuning a voltage regulator and running hydroelectric generators at close to their stability limits. He zeroes in on the difficult conversations with plant operators about increasing the gain to achieve greater efficiency. He mentions a 1995 blackout in the western United States that renewed interest in generator testing following a 30-year hiatus. He discusses the transition from analog to digital computing and moving from hand-cranked calculators to high-powered electronics on which he completed frequency domain computations. He mentions competing control techniques, and the various computers available to the industry across the decades. He notes that a program he wrote to study a single hydro plant in New Zealand became embedded in the PSS-2 software and was subsequently treated as a definitive hydro governor model. He underscored his own practice of understanding, and often producing, the software within every programs he uses for his own work. He describes the early years of attempting to complete load follow computations on a dial-up phone line to a computer and an even slower teletype machine that produced the results. He charts the evolution of the various iterations of PSS products and their commercialization. He notes that running PSSE and PSLF still takes longer than it should but anticipates that computation times will come down soon. He reflects on working with Leon Kirchmayer, Charles Concordia and other leading lights of the power industry of the mid-twentieth century. He closes by describing the supportive and inspirational working environment of General Electric during his years there.

Note regarding Transcript and Timestamps:

Timestamps appear in the transcript but may not align with the recording. Pauses and interruptions may have been removed, and the overall length of the recording may be reduced. Factual corrections provided by the interviewee are indicated in italicized text within brackets throughout the transcript.

Transcript:

Julie Cohn (00:00:02):

This is Julie Cohn. I'm conducting an interview with John Undrill on Wednesday, July 16th, 2025, online. Daniel Molzahn and Sairaj Dhople will assist with the interview. Dr. Undrill, during today's interview, we are especially interested in your early work on dynamic system analysis, power flow computations, and the development of software that has been widely used by industry to plan and operate power systems. Understanding that your career has spanned 60 years, we are also interested in your perspective on key events and significant changes in technologies that affected your work. But first, we'd like to ask you to describe your educational background, and your professional positions over these several decades.

John Undrill (00:01:09):

Okay. In some respects, that challenges my memory, but here we go. I remember walking around a hydroelectric plant in New Zealand when I was 10 or 11 years old. And from that age until about now, I've never been very far from power plants and particularly hydro plants. It came about because a neighbor of our family home in New Zealand was assistant general manager of the state Hydroelectric Department of New Zealand, which was a government-owned utility. And he arranged for my father and myself and my brother to visit the plants and see how they worked and get a feel for them. It's turned out to be an advantage. I got my academic education at the University of Canterbury in New Zealand, B.E. In 1963, and Ph.D. in 1965.

John Undrill (00:02:29):

I regard my professional education as having been my first four years with General Electric Company. And I've been involved in pretty much the same thing from—I won't say from the age of 10 or 11—but from the start of undergraduate education at University of Canterbury until now. And in between has been a long and continuous education in the, I'm going to say, dynamics of power systems. But—and this might sound strange to people coming into the power systems industry now—but won't to the old timers, the dynamics of power systems and the understanding of power systems is the understanding of the dynamics of large moving things. The large moving things that we've dealt with have been rotating generators, but it could equally well be ships tied together by refueling hoses and appropriate lines, or vehicles riding on a rough road. It's the dynamics of things where the masses of the things involved are significant, or large, compared with the strength of the things that are holding them together. Is that sufficient on the education question?

Julie Cohn (00:04:06):

Yes. Thank you. Would you also take a moment to describe the positions you've held throughout those years?

John Undrill (00:04:15):

Okay. Well, I joined GE in 1966, basically straight out of university. I was a postdoc fellow at University of Toronto for a year. And while I was at an IEEE meeting, during that postdoc year, I met Dr. Kirchmayer from GE. I was young and cheeky at the time, and having met him and exchanged politenesses, I asked him if he'd give me a job. He gave me a job application form on the spot, and 10 days later, he offered me a job. So, it was an excitement to be able to join GE. It was the kind of thing that a student like me would dream of, and led to a long, happy professional life. I left GE and joined PTI shortly after PTI spun off from GE.

John Undrill (00:05:19):

I joined PTI as one of a group of 13 engineers in running a spinoff technical company. I left PTI in 1986. I was vice president at the time. I founded Electric Power Consultants with two other gentlemen that left PTI with me. That was operated from 1986 to 1994. In 1994, for reasons that we'll talk about in a few minutes, GE bought Electric Power Consultants, so brought me back from having been a first-year engineer in GE in 1966 to an old engineer sitting in the seat of the guy that hired me in 1966. And I stayed in GE until 2006. I've been retired from the profession, but still answering the bell, from 2006 until now.

Julie Cohn (00:06:39):

So, you are very firmly on the private sector, industrial consulting side of the industry. Did you do anything in academic positions at all?

John Undrill (00:06:52):

Well, I had a year as a postdoc fellow at University of Toronto, which was both an introduction to university life and an introduction to the relationship that existed between the university and Ontario Hydro. Ontario Hydro Research Lab. And then I spent a year back on the faculty of University of Canterbury between first stint at GE and going to PTI. That was intended to be the, the beginning of the long association with the university in New Zealand. But PTI offered me an offer I couldn't refuse and probably took me a microsecond to decide to come back to New York. And PTI—much less easy decision for my wife and daughter to come back to the US, but we're glad we did.

Julie Cohn (00:08:03):

Thank you. So, how would you describe the relationship between academic institutions and the companies you worked for over those years?

John Undrill (00:08:18):

Oh, that turns out to be easy to do because it's a happy answer. I joined GE at the height of the large transmission expansion projects that were going on in the USA, in the Midwest and in the far west, at the time. And a spell of maybe 10 or 15 years when large hydro plants were being developed and worked into the electric system of Brazil, Argentina, Venezuela, where the issues that were challenging the people that were interested in what I was involved in, were how to build, how to design, and then how to operate large radial transmission systems. And a part of the work that was required was understanding the dynamic behavior of these machines in situations where they were only economic, if they were operate operated close to their stability limits.

John Undrill (00:09:55):

So, the engineering work that was being done by GE Electric Utility Engineering operation, which is where I was, by PTI after it spun off from GE, by Canadian GE in Dominion Engineering, who had major interests in electric system in Brazil, the Italian CESI Organization, and English large system research labs. The world was concentrating on basically the dynamics of these very large systems. It was the same dynamics as had been of interest in the very large systems that were developed, let's say between 1915 and 1925. They were systems with capacity, radial transmission capacity of, oh, perhaps a hundred megawatts, and were operated close to their stability limits for economic reasons. The difference in 1965, '66 to 1972 was these large systems now were operating at the stability limits, where those limits were transfers of 5- and 8,000 megawatts.

John Undrill (00:11:21):

The hydro plants that were being built had nameplate ratings of 4- to 8,000 megawatts, and then ultimately the Itaipu power plant between Brazil and Paraguay, whose nameplate capacity was, I think, 18,000 megawatts. The questions involved with the dynamics of these large radial systems, with a large mass attached by a highly stressed electrical interconnection to the legacy electric system, that was going to receive these large inputs of power. The questions at the time were the behavior of the new radial extensions that were attached onto these relatively compact, quite large, but relatively compact urban systems—Caracas in Venezuela, San Paulo, and Rio, and Bella Horizonte in Brazil and Buenos Aires in Argentina, where there were electric systems that were quite large, thousands of megawatts, tens of thousands of megawatts, but with an injection of power coming in radially on double circuit 400-KV transmission, or, in the case of Itaipu project, in the end, five parallel, 500-KV circuits and putting this large injection of power into a large, but relatively not huge, receiving electric system.

John Undrill (00:13:09):

So there were the questions relating to holding the thing in synchronism and transient stability problems that showed up in the textbooks from 1925, but also very real problems as to how you were going to control the voltage on this 500-KV transmission system, or 400-KV transmission system, when at night and on the weekends the receiving system was a load, but not a big enough load to hold the voltage down. The radial transmission lines, once they get to the 500 kilometers long, or so are, at light load, a large capacitor, the voltage goes up. And so the problems of transient instability, which is the first thing everybody thought of, were joined by problems of voltage control, both steady state voltage control—you were involved in switching shunt reactors and shunt capacitors at 500 KV and 400 KV level—and voltage in the dynamic sense that the voltage regulators and the voltage controls on the synchronous condensers at the receiving end of the transmission had to be tuned and adjusted so that the whole thing was stable and damped.

John Undrill (00:14:41):

And now that's a long answer, a non-answer to your question, Julie. The question was, what was the relationship between GE and the investor-owned industry side of the discussion to the university side? The answer was, in my case, a very happy one. What I did in New Zealand, during my education, was set up models of the elements of what turned out to be the big thing of the late 1960s, 1970s. I was given access to an analog computer at University of Canterbury. It was a home-built, very ingenious, hundred-volt vacuum tube analog computer. And with that, I built up a dynamic model of the Ōhakuri hydro plant, very conventional hydro plant, four by 28-megawatt generators, about a hundred-foot head, 128 megawatts nameplate capacity, I believe.

John Undrill (00:16:11):

And I started from an analog computer model of a hydro plant constructed by the engineering people at Manitoba Hydro. A gentleman named Hovey. Extended that a little bit and got, by virtue of playing with an analog computer, which was equivalent to playing with a hundred-megawatt hydro plant, understanding how the thing worked. I was given a gift there that would be hard to imagine now but was possible in an island in the Pacific Ocean in 1965. I asked one of the engineers in New Zealand Electricity Department, would it be possible to do some tests on one of your hydro plants to see if this model that I've made in the analog computer is any good and is reasonably accurate? The expectation, the reasonable expectation was that the answer to the question would be, "no, we don't let graduate students play around with our electric power plants.

John Undrill (00:17:29):

We have to have our businesses supplying electricity to the nation." That wasn't the answer I got. The answer I got was, "yes, we're curious. Write us a test plan, and we'll see if we can arrange for you to do some test work at one of our plants." Well, that was a simple statement, but it illustrated the difference, very quickly, between university life and industrial life. "Yes, we're interested, write us a test plan." And writing a test plan put me into closer contact than I was comfortable with the professor at the head of the electrical engineering department in the university, who, turned out, had a fairly strong opinion as to how a test plan should be written. And there were many drafts. And the test plan was accepted, and we were permitted to take our test gear from the university to the Ōhakuri plant on the Waikato River, New Zealand, and run tests on the governor.

John Undrill (00:18:43):

The tests that we ran on the governor were a follow-on and follow-up on suggestions made by Manitoba Hydro, who had done some test work themselves, and had developed their analog computer simulation in association with the test work. We did the test work. It was successful. Resulted in a paper that went to the IEEE in 1966. I believe that Dr. Kirchmayer, who hired me at GE, had read the paper, and that may have had something to do with why he would offer me a job. The thing that showed up was, in the university environment, the thing that mattered—and this was because of being at the entry level of the profession and being in the university—

John Undrill (00:19:53):

The thing that was interesting was the technique that was used to address a question or a problem. The problems that were created were somewhat academic. I spent quite a lot of time while I was in the university on trying to develop Lyapunov functions to assess stability of dynamic systems. Turned out to be not necessary, and not very fruitful. But it gave me the connection with the analog computer, and then subsequently with the digital computer of the university, to simulate this hydro plant. And then, the ultimate thing that lasted with me for 50 years in the profession was not the technique that I was studying in the university. It was the exposure to the real thing—hydro plant in New Zealand—until I was finished there. And then, when I joined GE, my goodness, I could go at lunchtime and walk around the factory and look at the barrel of the steam turbine on the milling machine being fabricated. And I didn't, at that stage, get to run around hydro steam plants and smell the coal dust, but I was certainly able to walk around the turbine factory and see the size and the shape and the style of the steam turbines.

John Undrill (00:21:33):

The thing that I say to people all the time, when that question, your question, is asked is, had I not had the technique-oriented, mathematical, physical education of the university, I could never have understood the industrial side of the industry, as I became able to. Had I not had the connection with the industry I would've had a completely different connection over the 50 years with the technique and the mathematics and the analysis of the academic side. I think the good luck that I've experienced was by circumstances and geography—being thrown into both the academic side of system dynamics, and at the same time and immediately in parallel with it, have my hands on the knobs.

Julie Cohn (00:22:51):

That sounds like—

John Undrill (00:22:53):

If you ask me to choose one or the other, the academic aspect of the industry or the industrial aspect of the industry, in isolation from one another, I wouldn't be able to make a choice. I think the big privilege was being able to be involved in both sides of the coin.

Julie Cohn (00:23:17):

I have a couple of follow-up questions. First, briefly, was there, when you were preparing and performing a test on an active hydroelectric plant, were there any risks to the customers that you considered going into it?

John Undrill (00:23:34):

Oh, we can spend four hours or 40 hours on that question, if you like, Julie. Yes, there were. And when we hooked up test equipment to the governor of one of these units at Ōhakuri, no risk involved there to the machines. And then we disturbed the machines. Now it turned out that the Woodward mechanical governors on those machines had a characteristic that had been built into the Woodward design of governor for the evolution of the things from the 1920s to 1960s. By a well-calibrated movement of a lever in the governor linkage, you could make the machine think that the electric load on the generator had suddenly increased by 5% or 8% or 10% of the nameplate, and that caused the speed of the generator and the power produced by the turbine to change in a predictable manner. We used the analog computer to predict the change.

John Undrill (00:25:00):

The owners, of course, demanded that we show how big a change we needed to be able to get a useful result. And the discussion that took place then has taken place every time I've gone to test a power plant in the 50 years. "Good morning, sir. I'm here to do a test on your plant. I'm going to have to disturb the turbine." "Okay. We know that you've done your homework. You can make a 1% disturbance of our turbine." "Well, sir, I won't be able to see any useful information from a 1% test of the turbine, and I need a 15% test." The air turns blue with the intensity of the discussion for a few seconds. "Well, you know, we really don't want you disturbing our turbine by a large amount."

John Undrill (00:26:01):

And the discussion proceeds. And the kind of disturbances that are settled on, were settled on in the 1960s, by me and by others who were doing a similar thing, Manitoba Hydro, notably, were fairly small disturbances for reasons that we'll get to in a minute. Now one can go to a large coal-burning steam plant and inform the management that you're going to do testing, and ask, "would they please exert a 20% change in power production as quickly as you can so that we can observe what happens?" And the answer from the owners of the plants has become, "yes, we can do that." And that's the benefit that the industry has had from being able to test and stimulate the responses of the equipment that we've had; and see what happens and learn by doing it what's safe and carries acceptable risk, and what carries unacceptable risk. We've never, that I'm aware of, had an incident of not being safe. Now, there, of course, was, in the middle of this evolution of what we were doing, a very major evidence of "be careful when you go test a power plant because you might be doing something unsafe." That was the incident in Chernobyl.

John Undrill (00:27:50):

Well, we weren't unsafe at that level. And the risks that were involved were real, but not, of course, of the Chernobyl level. The risk was, and it proved itself, what we worried about happened—the risk was that we would take a plant that had been operating steadily and happily for 10 years, hardly moving from its standard operating condition—the river runs, the water goes through the turbine, the turbine makes 90% of nameplate rating 24/7, 365 days-a-year. Its joints get stiff, things wear, and nobody's aware of what's happened because nothing's been done, nothing's been exercised. It's never been taken to the operating point that we're asking for in a test. Here we come into a plant. We take a plant that's operating in a standard condition that's well understood by its operators, and we disturb it from that known, steady condition to a different condition that they've never seen before. Something happens that was not expected. It turns out that an electrical connection in the control somewhere had been deteriorating slowly, and in fact, had been failed for the last couple of years. Nobody knew, because that aspect of the control was never exercised. We come in and do a test, the thing that was failed reveals that it's failed. And the unit trips. Well, the one thing you don't want to hear when you're doing testing in a power plant is sudden silence.

John Undrill (00:29:43):

When you trip a generator, things get very quiet very quickly. And we've had, in testing—what—300 units over 50 years—we've had two or three incidents where something was not as expected. It was not known until we came to do a test. And when we did a test, we caused the unit to trip. Now, the question that the owners asked, and the first instance when I was still in the university, and they still ask now when I go to go to a plant, is if we trip, how long will it take you to get back in service? You trip, you'll be off the line. You may go all the way down to standstill. You'll have to run the turbine and generator back up, synchronize back on the grid, and pick up load.

John Undrill (00:30:47):

If everything's in good shape with a hydro plant, you can do that in 10 or 15 minutes. Add a half an hour for everybody to walk around and make sure that everything's lined up the way you want it, because this stranger's been in the plant doing test work. If you're talking about a coal-burning steam plant in 1973—if you tripped and went down to turning gear, it might take you two days to come back. Might have taken you two days to come back. If you went to the same plant today, modern electronic digital turbine controls, boiler controls, still burning coal from a mine mouth pit, you might go down to turning gear quickly, if everything stays hot, get back in four hours. So, the risk involved in this testing in the plant was at the beginning.

John Undrill (00:31:56):

And there's still now, but to a lesser extent now, the risk that something was abnormal, [*but*] was not known to be abnormal until you do a test, and going back into service will require rectification of the thing that was not normal. And that may take some time. While the unit's not producing power, the owner of the unit probably has to buy power from somewhere else, because of its obligation to keep the lights on. And so, the economic impact, the economic risk is substantial. It's several hours in an easy case, and a couple of days in a difficult case of lost revenue, or differential cost of power, which can be large.

John Undrill (00:32:51):

The risk is going down because of two things. People testing generators was very unusual in 1965. Now, it's required by NERC standards, and it's done as a business. There are people who are able to come into a plant, wire up their test gear, do a test, and be gone in one day. It's a very well understood thing. Doing it has revealed what works and what doesn't work, and what's easy and what's not. But remember, when we started this, measuring frequency deviations of the grid, with a bandwidth of five hertz or three hertz, it was a huge undertaking. To measure speed variations at that first testing that we did at Ōhakuri, we had to basically run 35-millimeter film across the front of an oscilloscope, recording the 420-hertz wave of a permanent magnet generator on 35-millimeter film.

John Undrill (00:34:08):

Then after the film was developed and run across a light table, measure the length of film required for 15 periods of the 420-hertz wave. It was then writing down by punching the numbers into punch cards for the IBM 1620 computer. It was a laborious task. It was boring. Had to be done. We didn't know what test results looked like—that aspect of the test result—until days after the test, because the film had to be developed and the cycles counted. I had a very beautiful young lady from the art school in the university, volunteered to help me write the numbers and count the cycles. We did that night after night for quite a long time—seemed reasonable, after she had done that for me, that I should ask her to marry me.

John Undrill (00:35:08):

I did, and she did. And she said, "I'll no longer count cycles in the test work." Now the equipment that we're asked to test to satisfy the NERC standards and for other reasons—commissioning testing in particular—is done with equipment that's built into the control systems. A digital excitation system has a digital sampling and data storage capability built into it. That's been the case for the last maybe eight years. So that you can do a test and see the result as the test goes on in real time, and you're basically [*re*]gaining the advantages that we had 50 years ago with the analog computer

[REMOVED FROM THE TRANSCRIPT]

John Undrill (00:36:18):

The risk is that you'll trip it, and it won't go back quickly. So, you would do with an analog computer, things that you wouldn't do with the real turbine, but you—you can't, or yes, you can see me—<gestures> you can grab the knobs and say, "I wonder what would happen if I turned that knob that way, and that knob that way." Well, we did it on the analog computer and saw oscillographic traces responding. When we did it in the power plant, we heard the noise change and saw the meter needles swinging. Now of course, you don't see any meter needles swinging because the meters have been replaced by CRT tubes [*digital displays*].

Julie Cohn (00:37:04):

Thank you. I noticed that Sairaj was about to jump in with a question, so I'll invite Sairaj to go ahead and ask that now.

Sairaj Dhople (00:37:13):

Yeah. John, thank you for sharing these recollections. There's several questions that I could ask, but maybe I'll just limit it to one for now. You, maybe towards the very beginning of your remarks, you mentioned a problem in which generators had to be operated close to their stability limits to fully unlock the economic benefits of running the system. Can you say a little bit more about that?

John Undrill (00:37:40):

Yeah. If you're building a 3,000-megawatt hydro plant, 600 kilometers from the load, you've got to get 3,000 megawatts across 600 kilometers. And that's probably going to involve two, or three, or maybe even four transmission circuits, which are going to be radial to a relatively weak receiving system. But nevertheless, mother nature does mother nature's thing, and she hits transmission lines with lightning, and circuits go out of service, and the system's expected to operate continuously. Now, you can always, or not always, but, for practical purposes, you can, if you wanted to build a transmission system and say, "this system shall be stable and shall never, never fail to deliver its 3,000 megawatts from the hydro plant to the load." You could do that by building a sufficient number of transmission circuits.

John Undrill (00:38:53):

Well, the Itaipu plant, 18,000 megawatts, does have five parallel, 500 kV—I think there's actually 765 parallel AC transmission circuits and a parallel DC project. [*The Itaipu plant, 18,000 megawatts, has two 765 kV AC circuits and two 600 kV DC circuits.*] So, if you build enough transmission, then certainly you can achieve a very high, well, let's put it the other way. You can achieve a very low probability that an act of God event, like a lightning strike, will result in a disconnection to the plant. But remember the transmission part of a large development project, a dam, a hydro generating plant, transformation stations, switching stations, and the receiving substation, has the transmission cost be a fairly significant part of the cost of the project. I don't have the number on my head, but my guess is that the dam and the plant might represent 70% of the cost of the project, and the transmission and receiving substations might represent 15% or 20% of the cost of the project.

John Undrill (00:40:16):

Well, if you double the cost of the transmission by insisting that we have a very low probability of an outage—a lightning strike taking out the plant—then you've increased the capital investment of the plant by, let's say, double the transmission, increase the cost of the project by 15%. That's 15% of $3 billion. That shows up on the balance sheet of a nation like Venezuela or Brazil or Argentina, which is where these projects were built—or Northern States Power [*or other utilities*], if you're talking about transmission across the Dakotas from the coal fields in North Dakota to Minneapolis. By the economic requirements, you don't build transmission that you don't need. The transmission projects were built so that they could reasonably be expected to withstand a lightning hit, a known probable event, and remain stable with respect to synchronous stability, stable with respect to the damping of oscillations, responsive and stable with respect to the response to the voltage controls. You could surely get a system and say we're operating far from our stability limit by building more transmission.

John Undrill (00:41:59):

But you know, a well-tuned voltage regulator doesn't cost anything like the cost of a 500 kilometer, 500 KV transmission line. You're talking tens of thousands of dollars to get the voltage control voltage—just picking an example—voltage regulator controls of a generating unit, right? You're talking million dollars a mile to build 500 KV transmission.

Sairaj Dhople (00:42:31):

Well, thank you for sharing that, John. Just one more question, if I may. You know, you referenced, you know, Woodward governors and, you know, a lot of what you're mentioning also suggests that there was some level of standardization that was probably required for parameterizing control settings and so on. Can you say a little bit about what was the ecosystem back at the time and whether the relationships that system operators had with the manufacturers and the plant owners permitted. Standardization and conversations that were not too abrasive?

John Undrill (00:43:13):

Yes, I think I can. I got curious a few years ago, as to when the literature of hydro-governor adjustment, hydro-governor tuning, started. When I was doing my graduate work, the papers were coming from Manitoba Hydro—this was 1965, '64—Manitoba Hydro, Ontario Hydro. There were papers published in France somewhat earlier than that in a publication named La Huile Blanche, the White Oil. And these were dated [*shortly before*] 1940, '44, '45. There was nothing between 1939 and 1945, of course. And I thought, well, you know, these guys were pretty good. They were adjusting, building these governors, and adjusting them, and getting pretty acceptable behavior. But we didn't really know what the behavior was, because there was not much measuring equipment to record it. The good behavior meant the operators and the plants were able to operate them. There had to have been something known about governors earlier than that. And I went looking for publications, and I found one dated 1826. The author of the paper was Mr. James Clark Maxwell.

John Undrill (00:44:55):

Well, I thought that was rather cool because I could go to morning tea with the guys from the radio frequency and the plasma physics and the internet technology and say, "yeah, but you guys deal with Maxwell's equations, but we deal with Maxwell's equations relating to mechanical governance." Maxwell's paper on governing is well worth looking at. The notation, of course, is somewhat different from what we use now. And he was a hundred years before Bode invented the Bode Plot and Nyquist, had done his thing, and he understood the differential equations that—order of four or three or four or five, depending on his assumptions—just as well as we have for the last hundred years.

John Undrill (00:46:03):

Back to your question, the governors from Maxwell, from James Watt to, let's say, 1955 or 1960, were mechanical. They were flyball speed detectors with hydraulic power amplification to move the hydraulic jacks that move the wicket gates of the turbines. Whether they were operating satisfactorily or not, was by trial and error, where the trial and error was done very cautiously because—to the discussion we just had about the risks involved, when you're maneuvering something and you don't know how it's going to respond, you have to be very gentle and very cautious.

John Undrill (00:46:57):

What I found when I got to testing governors and operating plants, starting, let's say, 1974 through now—because that's when the equipment became available—we found that when we could test hydro plants to see how they responded, they were adjusted extremely conservatively. It was not a question of an oscillatory response. It was a question of a very overdamped response. Well, the reason was, the only way that you could find out what the unit was going to do was to build it and set it up and see what happened. And everybody that's tried to adjust a controller, when he doesn't know what the characteristics of the controller are, learns very quickly the same thing. If it oscillates the gain's too high, if you want to kill an oscillation, lower the gain.

John Undrill (00:48:07):

So, as you lower the gain, the thing does become less oscillatory, and then overdamped, and then very heavily damped. And we found that the grid was operating with many hydro plants whose response could have been: this plant can pick up—increase its output by 10% in 10 seconds. And in fact, what it was capable of doing, because of the way its governor was adjusted, was it could pick up 10% of its output in—now, hold your seat, and don't laugh when I say it—10 minutes. Because the people needing to adjust the governor—not having laptop computers, not having IBM 7094 computers that took up a warehouse, and not having even power detectors or speed detectors that could react with the kind of response they wanted—their visualization of what—or not their visualization, their information as to what they'd achieved when they applied a test disturbance was what they saw on the indicating meters on the panel.

John Undrill (00:49:41):

You know, the panel in the power plant. It's black marble, it's eight feet high, three feet wide, and has great big electric megawatt and megavar meters that are brass castings with pointers that moved back and forth across the quadrant of the meter. Well, they're all gone now, but that was what was available to the guy who had to come into a plant and decide how to adjust the governor. So, the governors were typically adjusted extremely conservatively, and responses were not very impressive at all. In fact, they were frighteningly conservative.

John Undrill (00:50:25):

The standard discussion that I had to have with plant people, when I got to a plant whose controls were clearly adjusted with excessive caution, was, "look, we need to raise the gain. We need to raise the gain from whatever it is now to a lot higher." And we were adjusting the gain, let's say, by adjusting a slide linkage on a mechanical linkage. So, you loosen the thumb screw, move a little slider along, and tighten it up, get an Allen key and cinch it tight, and then try the test again. And, I would say to the people that were making the adjustments for me let, "let's move this from there," <gestures out of view of camera> and now I'm pointing with the fingers in front of you. "Let's move it so that instead of being like that, it's like that we need a lot more gain." And the plant owners would say, "oh my, you can't do that. We've never done more than that. And in fact, we've actually gone the other way. You're asking us to go the way that's instinctively wrong to us and by a large amount." [*Interviewee noted after reviewing transcript that this type of conversation was typical for anyone doing the testing*.]

Julie Cohn (00:51:45):

John, can you just hold up your hands a little higher to show us those differences?

John Undrill (00:51:48):

Oh, yes. So, you know, the thing [*adjustment*] was set with a slide linkage—on a slide adjuster on a linkage like this. <gestures> And I'm saying, "we need more gain, please adjust it from there to there." And, their reaction was, "my goodness, we've never adjusted at anything like that much, and we've always gone from there to there. You're asking us to go in the direction that we think is dangerous, and by a very large amount." And so, you can see that this test session, which started out at seven o'clock in the morning, is still going at seven o'clock at night, because we're inching very slowly to get to where we think we need to go. How did I know—here I am, I'm a graduate student in the university. I'm not an engineer with 50 years' experience.

John Undrill (00:52:42):

How did I know that I wanted the adjustment to go from there to there, as distinct from their instinct, which was to go from there to there? I'd been simulating this thing for the last three months on an analog computer, and I had no idea, simulating the thing on the analog computer, whether that starting point of their adjuster represented a gain of two or a gain of 10 or a gain of 0.2. But what I did know, from the analog computer work was, regardless of what the gain is, if I double the gain, the behavior of the thing will go from a response that's characterized by that <gestures> to a response that's characterized by that. Or if I go from whatever it is to a half of whatever it is, it'll go from that, which is a bit bouncy, to that, which is very conservative.

John Undrill (00:53:47):

So, by having the analog computer simulation, and having twiddled the same knobs, the same parameters, with the knobs that I twiddled on the analog computer, that I was now asking him to twiddle in the test, I knew which way it was going to go. And that was the role that testing played in the early days of generator testing [*commissioning*]. Now I'm skipping 30 years here. My list of generator testing that I've done, starts in 1964, I believe, at Ōhakuri, and then it's silent from 1964 to, let's say, 1983—a little bit longer—1987, because testing was just not done and certainly was not done by people from outside the industry, outside the plants, other than manufacturers doing commissioning testing. So, when we came to realize we had to do testing, the industry came to realize at 4:40 PM on a nice November day—9th of November, and what was it, 1995, that testing had to be done. Because the western system went black.

John Undrill (00:55:23):

And the simulation work that was done to explain what happened when the western system went black showed very clearly that the simulations that were done were not matching the reality. And it was a very logical conclusion that the industry made, quite, quite firmly and quite quickly, that testing needed to be done. And then a lot of testing was done, and now it's required by NERC standards. But the adjustment of the governors, up until about mid-1960s was done by the manufacturers and the commissioning people, proceeding very carefully and very conservatively. The processes that we do now, where you can have a debate as to whether you want to use a root locus technique or a frequency response technique or a try-it-and-see technique, and set a performance expectation that's an overshoot like this, <gestures> or an undershoot like this, is done now with laptop computers and running simulations on the site.

John Undrill (00:56:52):

And "hey, we'll, let's try this." "No, don't try it yet." "They give me five minutes to run the simulation." Run the simulation, bzzzt, "yeah, that will look okay, we'll do it." But that's now, and 1960, '65 to 1975 was then. You have to remember that, you know, I'm sitting in front of a 16-gigabyte laptop computer. When this career started, there wasn't 16 gigabytes of memory in the entire computer industry. When I was doing this analog computer work in the university and doing it in parallel digitally on the IBM 1620, the university had one computer. It was in the engineering school. But the university had one computer, and the computer would be operated by a guy from the geography department doing statistical analysis of the population question from 9 to 10, by somebody from the chemical engineering department doing simulations of the freezing of meat carcasses in the freezing works from 10 to 11. By me, from 11 to 12 doing governor adjustment. The next one would be a civil engineer looking at earthquake-responsive building. The computer was running 24 hours a day.

John Undrill (00:58:26):

You booked a computer three days in advance to get your one hour slot, and then you always showed up early for the start of your one hour slot on the computer, because the chances were that the guy in front of you had trouble with his logic, couldn't get his program to run, and the computer was then going to be idle for a half an hour. It was going to be taken by the next guy in the queue.

Julie Cohn (00:58:54):

I am guessing, Dan, that you have some follow-up questions now.

Daniel Molzahn (00:59:00):

Yeah. So first I want to say thank you, John. This has been super interesting and really enlightening. I've really enjoyed hearing the stories, especially going back to maybe the 1800s in some cases. So, one, you've touched on, on this question I'll ask in a little bit, but maybe you can expand it—so, you know, at different stages of doing all the testing and control settings, you alluded to using some theory. But it sounded like you also did a lot of experimentation and actual testing on the real hardware. So how much—was the theory that you were drawing on, and maybe the models as well, was that out there and you were applying it, or you were developing the theory, developing the models as you went along? Did you talk with control theorists, and they told you what to do? Or this was you guys developing it as you went?

John Undrill (00:59:56):

The fine print of the project that I undertook for my PhD was to try to apply a Lyapunov theory to the adjustment of hydro governance. Try to develop a Lyapunov function that would show that the governor was properly adjusted and show how the parameters of the governor should be changed to achieve that. I made some progress. It turned out that it was a huge computation, a computational effort that wasn't worth it because a practiced engineer doing pencil and paper calculations on the back of an envelope could get an equivalent result much more quickly. When I got to GE, with my head full of Lyapunov theory and so on, I found that GE approached a governor-tuning question or a controller-tuning question in just about any industry in a way that had come to life during the Second World war. The controls of interest in GE, in the part of GE that I was in when I arrived there, were people whose experience during the second War World War was gun turret controls on B-17s.

John Undrill (01:01:33):

And the control theory of the time that had been used at GE, graphically, by hand calculations on the back of envelopes, and by rooms full of young ladies operating Monroe calculators—that is, you know, crank the handle to multiply, crank the handle seven times to multiply by seven—calculators—had been based on basically Nyquist and Bode methods. And there was a lot of discussion as to how much gain margin and phase margin you should have in your controller. When I got to GE, the analog computer was still at its zenith. And so, the understanding of what parameters you would want to adjust and what amount you'd want to adjust them by was based on: how much do you want to shift this Bode plot to the left or right. Bode plot is a plot of amplitude of response versus frequency of response.

John Undrill (01:02:48):

And the guys who'd grown up with this theory, which is, after all, the theory of complex numbers—it's as theoretical as you can get—had gotten very comfortable with how you can adjust a gain in phase and phase lag to get the response that you want in terms of aiming for an adequate phase margin or adequate gain margin. When you're sitting in front of an analog computer, it's just so logical, so natural to be thinking in terms of the Bode plot that you get on the desk on one side and on the adjustment panel of the analog computer on the other side—adjust the time constant of the OPAMP, of the integrating blocks and summing blocks to move the break points of the Bode plots left and right as you want to get.

John Undrill (01:04:01):

So, I learned very quickly, when I got to meet my mentors at GE, Dr. Kirchmayer, de Mello, the integral—the Lyapunov theories, theories based on the energy functions and so on, certainly didn't do me any harm to investigate that and learn something about it in the university. But when the question came, not how did you get the answer, but what is the answer? Things moved very quickly from the energy function gain to the gain and phase margin theory of frequency domain control theory. That persisted. When I arrived in GE, the other young guys in the department were all ahead of me in terms of the digital computer work that had been in universities like Northwestern, MIT, RPI. They'd had access to, had more background in digital simulation.

John Undrill (01:05:34):

And I had the digital version of my hydro plant simulation. But I'd not had access to the bigger computers and able to do the matrix operations that you can do when you've got the memory. What I found was that—I found it very easy to talk, much easier than they did perhaps, to talk to the old guys in GE who were my mentors, because I'd had several years of hands-on experience with the analog computer and the digital computer. And I'd not used frequency domain methods in the university, but when I saw them used in GE, it was kind of like I was a water dog on the pier. "Oh, look, there's water!

John Undrill (01:06:42):

I can go jump in it." And I jumped in, and I've been swimming in the frequency domain analysis of the systems ever since. Things have changed now, of course, because the stuff we're adjusting controls for now is electronic. If you look at the controls that actually work, that people say, you know, "here's a controller, I've got to adjust it. I'm interested in, not in the method I use to adjust it, but in the result that I get when I have adjusted it," the controllers tend to be proportional plus integral. They tend to be proportional plus integral controllers with neat things done to make them less sensitive to noise at the high frequency end of the spectrum and so on. But the adjustment of controls is—really, you use the analytical method of your choice. Some people prefer to work in root locus domain. Some people prefer to work in Bode plot domain. Some people don't want to use either. They want to go direct to simulation. As long as you're getting the right answer, which is provable by running a simulation, that's the objective is get the answer,

Daniel Molzahn (01:08:08):

That's really interesting. And I appreciate the perspective on the different control techniques. One other question along these lines is—it relates to models. So, you know, when you were starting this in building both, I guess, the analog and digital models of these systems, was there well-established models that sort of, everyone agreed on, this is the right physical representation at the right level of fidelity to capture all the behavior you wanted? Or was that, you know, something you had to develop, or were there arguments or discussions about how to best do that?

John Undrill (01:08:40):

As I started out on hydro governors, the accepted model at the time had been published in France in, I think, the late 1930s. And it was a differential equation model. It was not expressed in transfer function terms. It was expressed in terms similar to what Maxwell had used in terminology in his paper in 1800-and-something. The Manitoba Hydro paper had taken the differential equation description of the governor and implemented it on a little EAI TR10 analog computer. And that was the model. Generator model, of course, was well understood. That was Park's and Kron's and the theory of synchronous machines. Governors—steam turbine governors—were fairly well understood in terms of the mechanical construction and ability to move.

John Undrill (01:10:03):

But they were not asked to move anything like within their dynamic capabilities. So, they were not particularly well understood. Transfer function started to show up to model the controls. About the time I was joining, I was learning the ropes at GE, the gun turret people from the airplane industry and the radar antenna people, and then, in the later stages of rockets, the Apollo rocket people were all describing their systems in transfer function terminology. But when I started out to model hydro plant control in 1960, '64, there was no, "hey, here's the model of a hydro plant"—there were papers on what people had done with varying success, but it was: write your differential equations, understand the physics of the thing, and then get the thing into a form that you can integrate by the integration process of your choice.

John Undrill (01:11:25):

There were debates at the time as to whether it should be second order Euler integration or a fourth order Runge-Kutta integration or implicit or explicit. And you know the story. But remember you were, prior to about 1976 or '77, you were always trying to stuff a large simulation into a small computer, small digital computer. If you were on an analog computer, you needed 25 operational amplifiers, but there were only 23 available. You were always trying to stuff a simulation into a limited computing resource. In—call it 1966 or 1967 [*1976 or 1977*]—Prime Computer produced the 400, which gave access to virtual memory access, and it increased the memory range that you could access hugely. And then models started to get developed and become a part of these programs. The programs started to become used by people saying, "instead of going to Westinghouse in Pittsburgh or General Electric in Schenectady to do a stability study, buy a computer program, or write a computer program, and I'll do the study myself." That mode of operation was in the industry from, call it 1960 to 1970. From about 1970 on, the electric companies didn't develop their own programs.

John Undrill (01:13:15):

They bought programs. The first program they bought in any volume was PSS-2, which they bought from us at PTI. Then came the phenomenon: well, PTI has a hydro plant model. That must be the hydro plant model. Well, no, that's not true. PSS-2 had the HYGOV model in it, because I needed a hydro plant model to do a little study that I was getting ready to do regarding a new hydro plant in New Zealand. Question was asked, "what shall we set as the operating rules for this plant in relation to the limited capacity of its surge tank?" I needed the model. Well, I put the model in PSS-2 because I was the author of PSS-2, and I had control over it. I put the model in and used it.

John Undrill (01:14:22):

Then when the program was sold to <inaudible>, it was delivered with generator models: GENROU and GENSAL, a turbine governor model which still exists in the programs T-gov1 and HYGOV, the hydro governor model. But that should never have been, and should never be interpreted as, The hydro governor model. It's a hydro governor model that some joker who happened to be writing a program in Schenectady in the mid-1970s happened to leave in the program when it was delivered to the industry.

John Undrill (01:15:09):

So, the professional engineering question that arises here is this: is the person using a computer program responsible for the modeling that happens to be in the program when he takes delivery of the program? My answer is, the only person responsible for the modeling, in an engineering sense, is the person using the model. That makes me nervous. I still do calculations. I still do dynamic simulations. I still do them using programs that I wrote myself—from the early days on analog computer to now on my 16-gigabyte laptop. I'm not comfortable if I don't know exactly what the program's doing. Exactly what the program's doing means knowing the C code, or the FORTRAN code, or the Python code, perhaps, but it's C code in my case. I want to know the code that runs the calculation.

John Undrill (01:16:26):

And that makes me old fashioned. But it also makes me a little bit wary of a lot of the engineering that's being done at present using models that are in computer programs. Because that's what was convenient to put in the program 15 or 20 years ago. 15 or 20 years ago, the computers that were available to run these programs were less than a handheld cell phone is now. And so, the models that we're using in these production programs don't represent the state of understanding and the state of knowledge of the way the physical systems work. And that's going to change. It's going to change because, as the power system becomes electronic, the modeling of the electronic components—the batteries, and battery chargers, and statcoms, and electronic equipment—are being developed in the computers that are available now, using the knowledge of the systems that are available now. And in my view, they make me less nervous, perhaps, than the models that have been in the programs for a long time. And probably in many cases, not probably, but definitely obsolete and should be replaced.

Daniel Molzahn (01:18:18):

I suspect I might be stealing Sairaj or Julie's question here. But how much do you think, and you are kind of alluding to this, but I just am wondering if you can expand that a bit further. Like the models that you put in the system or in the software tools, for possibly very specific purposes, that were not intended to be generalized but yet because they were there, did people broadly assume these were the defaults, and sort of take them as standards?

John Undrill (01:18:48):

Aha.

Daniel Molzahn (01:18:49):

I mean, how much of this was intended and how much were people actually customizing the models? Do you have a sense of that over the years?

John Undrill (01:19:00):

I have a sense of it, and I have something to say on it. I have in front of me, my page of notes that I made when we talked a couple of weeks ago, and in capital letters across the top is "the dilemma of detail." The point that that comes up again and again and again in modeling is, "I need to model this thing," be it the governor of a turbine or a controller of an electronic element of a plant. "I need to do a study of this thing." And the person that's advocating the equipment—that is the vendor that's going g to sell it, sell the equipment, or the planning engineer that's got to choose between alternatives to buy equipment—the people raising the question asking for simulations.

John Undrill (01:19:56):

Now that's going on today. In fact, the latest IEEE monthly journal came out in my letter box yesterday, with a whole issue devoted to the virtues and difficulties of electromagnetic transient modeling. Well, that issue is of the dual purpose of explaining to people what is going to be possible as these EMT and analysis programs mature and will come to production grade. And it's also advertising the ability of the various consulting companies to develop models that are superior representations of the equipment that needs to be studied. And it lets me jump up on the soapbox and point out that adding detail to a model does not necessarily improve the accuracy of the model with respect to the question that's been asked. It's very hard to—it's hard to find somebody who's in the modeling business who will say, "you know, that's a nice model of electronic controller, but you don't need that much detail. You get a just as good an answer to your question, if you would use this very simple model." It's a very hard thing. Particularly for a person developing models who's very much involved in getting the model to work, as distinct from using the model to get a piece of equipment to work. It's very hard to say, "yeah, that's an interesting detail, but it's not necessary for what you're trying to do."

John Undrill (01:21:53):

It's worse than that. Many of the people who are the intended recipients of modeling work that gets done are people who understand the need for the model, are asking the question that the model's addressing, but are not going to understand the model. They're going to be people at the management and executive levels of electric companies, people in the political sphere who are going to approve the purchase of pieces of equipment, the development of transmission lines on public right-of-ways, and all of the things that are characteristic of industry. They're not going to be able to look at a model written by consultant A and a slightly different model written by consultant B and say, "this simulation superior to that simulation because he has modeled this thing versus, he has modeled that thing.

John Undrill (01:23:00):

Furthermore, if you put a detail into a model, that requires that you have detailed knowledge of the equipment that you're modeling. And that needs to be at two levels: A, the transfer function description of the thing, or the differential equation description of the thing, or the discontinuous curve description of the thing is this, and it has this form, and—you can see I'm moving my hands <gestures>—a graph with a straight line section and a jump and a straight line section and a jump; or a transfer function segment or a structure of a model. If you're going to put more detail into a model, into modeling something, you've got to understand the structure of the thing, and that's going to affect the structure of the model. And then you've got to have access to the numeric description of the parameter values that make the model meaningful.

John Undrill (01:24:04):

And that's usually impossible to get. In the situation we have now, is people are putting battery systems and battery chargers and battery discharges onto the grid, and they're needing to do studies determining whether these electronic controls will operate comfortably with one another when they're connected together in a grid. And they need to be modeled. They're being modeled in PSSE, PSLF, EMT, EMTP, PS-Cad in various levels of detail, all with basically no detailed knowledge of the numeric parameter values that are going to need to be put into these models to make them accurate descriptions of what really happens. The proper use of simulation, the proper use of the analysis in the early stages of use of these models, is to use them for experimentation with what would happen if this gain is doubled or halved or multiplied by 10 or halved.

John Undrill (01:25:28):

What would happen if this transport delay or this phase angle lag in a transducer is doubled or halved or multiplied by 10? If you can't take a—basically a study of a piece of equipment and say, as I said with the moving link on the hydro governor, "I want you to take it from there and go to there" as distinct from "take it from there and go to there." If you can't say, "what will happen if I go that way, and what will happen if I go that way? This and this," you haven't done the engineering. So, back to: is this the model of the thing? It's a model that you, as a professional engineer with the appropriate time, interest, and computing resources can use to determine what would be the likely behavior of this thing in the anticipated environment, or how would this thing best be adjusted to meet a performance requirement in this particular environment?

John Undrill (01:26:59):

There's not enough of that going on, in my opinion. There's, there's too much, "I've run thousands of simulations using a multiprocessor computer, and I've done literally thousands of simulations, and I understand something." But too often it's based on, "I used the blank-X or blank-Y or blank-Z model of a thing." And, "did you adjust the parameters and see how the parameters of this controller affect the behavior of the system?" And the answer is very often, "no. I used blank-X model, which came out of the X, Y, Z computer program."

Julie Cohn (01:27:45):

Sairaj?

Sairaj Dhople (01:27:53):

John, I think we would miss a big opportunity if we didn't ask you, and I think this is a good lead-up from, from where we are jumping off anyways, about your efforts in developing PSSE and, and PSLF. You are credited with having come up with both of those pieces of software. Can you say a little bit about how that all started, people that you worked with, you know, roadblocks that you ran along, and when you were building it, did you have an idea that at some point of time in the future, they would be running grids all over the world?

John Undrill (01:28:34):

Sort of, sort of. PTI got off the ground in transmission planning business in Brazil. Brazil had huge loans from the World Bank and was developing hydro plants on the Paraña River, which were remote from Rio and San Paulo. So, they were building this radial transmission from large hydro plants to relatively smaller, weak loads. And we, in PTI, had realized we were going to have to do a lot of stability simulation, dynamic simulation, a lot of load flow work to support our intended role in the planning of these projects. We had left GE recently, and the company was—the book value of the company was less than the book value of the computer that we could afford to buy.

John Undrill (01:29:42):

We tried to do load flow work using a timesharing system that was available in the Albany, New York area. It was a GE 400 computer. We talked to the thing by a dial-up phone line, 10 characters a second. The teletype printed out results. The computer was fast enough to do the calculations we wanted, but the teletype would produce the results at 10 characters a second—clunk, clunk, clunk, clunk, clunk, clunk, clunk, clunk, clunk, clunk. That's one line of text. And you need a hundred lines of text to see whether your load flow calculation was useful or not. That wasn't going to work. The company bought a Hewlett Packard 2120 computer that had 64 kilobytes of memory. It actually had 32K 16-bit words. [*16K of the 64K bytes*] The operating system took 8,000 words of the 24,000 words.

John Undrill (01:30:45):

So, we had 16,000 16-bit words. 32K, sorry, yeah, 32K less 8 K. We had 24K bytes to put a load flow program in. We got a 256-bus Newton Raphson load flow into PSS-2 on the 2120 computer. The little computer ran 24 hours a day. I used to go home and sleep from seven till two in the morning, get up at two in the morning and go and meet with John Burns in the computer room and take over the computer, and run load flows and stability calculations. It worked. It took that little 2120 computer about 10 minutes to solve a single 256-bus load flow case. But that 10 minutes was it. You knew that you were getting a result, and you knew that it was going to be good probably five minutes into the 10 minutes.

John Undrill (01:31:53):

And if you knew it was going to be a bad result, you just stopped the program and didn't run any further. The productivity was very, very good. And we were producing results for our Brazilian clients at a very effective rate. Those clients, not surprisingly, started to ask, "well, if you can do it with a 2120 computer there, why don't we get a 2120 computer here, and do the calculations in Brazil?" Sounded like a good idea. The information technology people in the client organizations said, "no, you can't do that. That's a computer. Computers belong in the IT department." What had to be done was to relabel this HP 2120 computer as a "dynamic network analyzer." The paperwork was all done, not mentioning the word "computer," so that these computers could be delivered in Brazil, and used as we were using them in New York, but they were in Rio de Janeiro.

John Undrill (01:33:09):

By the time that business undertaking ended, I think there were probably five or six 2100 computers in the office in Rio, running 24 hours a day. They were not running 24 hours a day, because in the summer, the thunderstorms would come into Rio, and there'd be a lightning hit, and the city would go black. You'd have the computer running, power would fail, the computer would shut down. Thank God they shut down gently. The discs had failsafe heads so that the disc didn't wipe when they shut down. But you'd lose the run that you were doing. Half an hour later, you'd see the lights across the bay coming on, and the lights would pop, pop, pop, pop, pop, pop, and the city would be back to full voltage again, and we'd plug the computers in, and off they'd go for another nearly 24 hours. That program, PSS-2, 256-bus Newton-Raphson load flow,128-bus dynamic simulation, was adequate to do basically the planning, by the standards of the late 1960s and early 1970s, of that EHV network.

John Undrill (01:34:39):

Well, the EHV networks in South America. About the same time as South American projects were reaching their end, the Midwest started to develop large transmission projects from the mine-mouth coal-burning plants in North Dakota and South Dakota to the Minneapolis area, and also transmission from Manitoba Hydro down to Duluth and Minneapolis. And the dynamic simulation models that had turned out well in PSS-2 became of interest to Minnesota Power and Light, Northern States Power, the Basin Electric Co-op, and companies in North Dakota. And here came a market for PSS-2, and then PSSE in the upper Midwest. PSSE was PSS-2 reprogrammed in a language called Flex, a variant of FORTRAN, and running on a Prime 300 computer. That evolved quickly into PSSE on the Prime 400 computer, which was the first computer that we had access to commercially that had virtual memory access.

John Undrill (01:36:20):

It gave us access, I think, to 64 kilobytes. No, it gave us access to a megabyte of memory, where the HP had given us 24 kilobytes [*48 kilobytes*]. That let PSSE expand very rapidly to 4,000 buses. My goodness, we got 4,000 buses! We can plan anything with this. Well, now it's, what is it? 40,000 buses? Or some very large number. But PSS, PSSE, PSS-2 was basically sold to our sister company in Brazil. And then a few US companies adopted PSS-2 on the Hewlett Packard 2100. Northern States Power among others.

When PSS-2 was in its limited sphere, it was basically, what the computer industry would call these days, a "sandbox." It was a very tricky program because it was a “quart” pushed into—a “two-liter” program, pushed into a “one-liter” computer.

John Undrill (01:37:42):

But in engineering terms, it was a very simple, very straightforward protocol where you could put a new dynamic model together very quickly and easily, because you didn't have to fight with a complicated data structure. The data structure was simple and small. 24 kilobytes for everything. So, we didn't really develop PSS-2 with the intention that it would be sold as a revenue-generating part of the business. By the time PSSE was developed, particularly on the virtual memory computers, it had become a significant part of the revenue stream of PTI. And so, the development of the program now was no longer an engineering tool, where simplicity of putting an engineering model together was the primary reason. And the productivity was inherently a simple productive thing.

John Undrill (01:38:57):

The Prime 400 was a timesharing system, and it had the emerging good display capabilities that came with the pen plotters and then the electrostatic plotters and laser printers that we use to put results out on. Now that pushed PSSE in three directions. It still was the program available to me and a couple of others of us who wanted to use the program and whose interest in the program was primarily: we want to simulate power plant stuff. We want to be able to understand the way this equipment works, how we're going to adjust it, how we're going to tune it, and put it into service. The transmission planning people who were interested in buying the program, buying our engineering services for transmission line design work, buying our planning services for economic planning and production costing, buying our PSSE programs so they could do their load flow work, were interested in, "well, why would we argue about making an equivalent of the company that's next door to us so that we can have an extra hundred buses to model our company, when we can just have one model that's going to be big enough to have our company modeled in detail and our neighbor companies modeled in the same detail?"

John Undrill (01:40:30):

And the pressure came on the program: handle more buses, handle more generators. Instead of having a plant with seven identical generators be represented as one generator of seven times the size, put in the seven generators, then you don't have to have the data management issue of how many generators are online for this particular case. So, PSSE was drawn and still had its roots in the power plant dynamics. Was drawn in the large system model direction, which forced the development of data management capabilities, that PSS-2 had never had.

John Undrill (01:41:21):

Because if you're going to run a 40,000-bus program, you're not going to sit down with a manual key punch and key punch the data to put into the program. And so, the thing acquired a personality of data management, which is, I think, still developing. And then, as the expectations of the people buying software evolved, and this still precedes the original PC, or precedes the Apple, people buying software started to see—because the famous picture of the space shuttle coming through the atmosphere in the, what was that? That was the early 1980s. You couldn't open a computer industry magazine or a technology magazine without seeing an advertisement which showed the red-hot glowing belly of the space shuttle flying in to land in California. And visual display of results became absolutely vital.

John Undrill (01:42:38):

PSSE had developed a one-line diagram capability. It contributed very little, in my opinion, to the ability of the program to deal with the engineering questions at hand. But it was absolutely essential for the marketing of the program. For the program to be able to be seen in the same light as other programs that had advanced graphics. The PSSE is still pulled, I think, in the same three directions, and it's now—let's see—it was first marketed, I think, in 1976, so it'll be 50 years old next year sometime. And it's still got its fundamental, original structure. Well, PSSE and I parted company in 1986, and I went off with Roger Clayton and John Burns to start Electric Power Consultants. On the first day, we announced to ourselves, "we're not going to spend time and effort at marketing a computer program."

John Undrill (01:43:57):

We're going to develop computer programs, because we're going to need them. We knew that. We're going to develop computer programs on this wonderful IBM PC AT that we have. This is a desktop computer with 32 kilobytes of memory. It's the first 32-bit minicomputer, desktop computer. We'll develop a program for it. We developed a load flow—it was 512 buses in 64 kilobytes of memory. We did it, used it, Roger Clayton used it to do work for a couple of clients in the New England area. It was a success. The clients, perfectly logically, said, "why should we pay you guys to run this program for us? Sell us the program. And we'll run the program for ourselves." And PSLF developed similarly to PSSE, but with the difference that PSSE's roots were in the Hewlett Packard 2100 computer, with extremely limited memory.

John Undrill (01:45:07):

PSSE [*PSLF*], by the time it was in the market, was running on a desktop computer called the IBM RT. It was a very obscure computer. It never got a lot of public recognition, but it was an enormous success. PSLF was written in UNIX, because that was the operating system of the IBM RT computer. And the combination of UNIX and the RT just gave us magic in terms of the ability to get that program to do what we wanted it to do. And it was adopted by Green Mountain Power. The godparents of the program were Green Mountain Power Company in Vermont, Southern California Edison, Northern California Power Agency. The companies that took it to use for themselves. It was then sold, subsequently sold, to Western System Coordinating Council and adopted by WECC.

John Undrill (01:46:20):

WECC got interested in the program. We tried to sell the program to WECC. They didn't want to put that many eggs in such a small basket. We were still only a five-man company, and they had very correct anticipations of what it would take for us to support PSLF from the small company that we had. Being small gives you the advantage of a very efficient communication among yourselves, ability to turn on a dime, and a very strong disadvantage when you have to take on a 24/7 support obligation and a 52-week-a-year support obligation. We, at the time, were approached by GE, who had heard about PSLF and gotten interested in it. And what happened in the end was that GE decided to sell PSLF to Western System Coordinating Council.

John Undrill (01:47:36):

To do that, they just bought our company. And I wound up sitting in the seat of the gentleman that had hired me in GE 40 years earlier. PSLF is alive and kicking now. It's still one of the mainstays of the grid simulation work that's done in the WECC system. I don't use PSLF any longer. I use a program that I started developing a couple of years ago with the intention of going back to the virtues of the little old Hewlett Packard 2100 computer. But instead of being limited to 32 kilobytes of memory on a single computer, I'm sitting in front of 16 gigabytes of memory running the thing in a similar style. No graphics, no fancy data processing, no elaborate data structures, no pushing of a large program into a small computer.

John Undrill (01:48:52):

I'm using the memory spaces as I need it. If I need to invert 32,000 square matrix these days, I just do it. It takes long enough to get a cup of coffee, but I can live with that. That's going to change. The phasor representation of the transmission system as done in PSS-2 and PSSE and PSLF is already approaching. It may already be at the limit of its credibility. By the time the next generation of software development people are talking to a group like you about the subject we're talking about now, they'll be talking about electromagnetic transient simulations. They'll be talking about time steps of microseconds, where we're still talking about time steps of milliseconds, though they probably will not be going to 60,000 nodes at the EMTP level.

John Undrill (01:50:02):

We've probably reached a point now where the studies of the electric system that involve the spread of the phenomenon—how far will a unit trip of a thousand-megawatt nuclear unit in the Carolinas spread, and how quickly will it spread to northern Minnesota and North Dakota—will probably still be answered by programs that are the descendants of PSLF and PSSE. But those programs will no longer be the forefront of understanding of how the equipment works. There'll be specialist programs for a specialist application. The fundamental engineering question from my point of view—which is, "how does this equipment work? Why did that equipment do that particular thing in that particular test?" Where the details of the equipment are important—is going to have to happen in the EMT domain. And the programs that are being used as—not quite science project level, but certainly specialist study level now—EMT, PSCAD, EMTDC, a couple of others— are going to have to grow and

John Undrill (01:51:19):

mature into production type programs of the style of PSSE and PSLF, where they have extensive data handling and data management capabilities tightly linked with them, tightly linked to the load flow capability. They're starting to grow those linkages now, but I don't think they're mature and they're going to have to mature. And then the question of running time will disappear, because nobody expects that the computers that they'll be using 10 years from now will be of the speeds of the computers we're using now. They'll be correspondingly faster. And the computation times, which are still too long today, will come down to be usable.

Julie Cohn (01:52:22):

John, you have brought us full circle from the beginning of your education profession to your projections for the future. And I think after keeping you online for two hours, maybe we have reached a good moment to pause for now. We really appreciate hearing everything you've shared with us about your fascinating, challenging, fun, and illustrious career. And of course, we look forward to the next phase of it.

John Undrill (01:53:00):

No, I just realized I don't envy you the editing task that you—.

Julie Cohn (01:53:07):

Well, the hardest part will be figuring out the acronyms and making sure I got them right. But maybe you'll help me with that.

John Undrill (01:53:13):

Oh, yeah, let me know.

Julie Cohn (01:53:15):

Sairaj, Dan, anything else you wanted to add at this point?

Daniel Molzahn (01:53:21):

If have time?

Sairaj Dhople (01:53:23):

Yeah, we can keep going around for a whole day, but Dan, I think maybe you have a question.

Daniel Molzahn (01:53:30):

Yeah, so, one other question that I was curious about—as you were describing this—non-technical—you had mentioned several times working with Dr. Kirchmayer, and the respect you had for him, and sort of the role that he played in your career, and what it was like at GE. So, I'm just curious if you can maybe say a bit about that, you know, we haven't talked with anyone else that, as far as I know, had worked with him or had been in those roles. What was it like being in that environment? What was the culture like? How did you know, how did work get done? How did people interact, these sort of things?

John Undrill (01:54:07):

Oh, this is going to take five minutes, but Julie, stop me at five minutes. I walked into GE with my head full of what I could read in the IEEE—Transactions of the AIEE, it was still AIEE. IEEE came along, a year later, I walked into GE, and in the first hour I was in the building, I met Dr. Kirchmayer, Charlie Concordia, de Mello, Gordon Carter, Len Garver. So, it was impossible to realize that "here I am walking around on the sixth floor of Building 2, and the same floor of the building of, basically, the seven giants of the business. Kirchmayer was the manager of the department I was in. He was my professional father for the first couple of years I was there. I had two professional fathers at GE, Kirchmayer and Paul de Mello.

John Undrill (01:55:29):

And the atmosphere of the place for a starter-upper like me—I was young, it was early days—was what you would see in a really good university environment. It was just nothing but encouragement to go try what you thought might work. If you've got a question—go down to the library in the basement and read up on it. The place was—you can't call it a lab because it was very definitely a business running for, you know, to be profitable, make revenue and be profitable. But the style of the place was—those guys were the technical encouragement department.

John Undrill (01:56:30):

I don't think I ever went and asked either de Mello or Kirchmayer, can I do this thing? They would always say, "what are you going to be doing now?" And I'd tell them. "Alright, go ahead and do it." The unspoken piece of it was—and they never needed to say it because it was in the way they spoke to you—"go ahead and do it, but if you find that it's not working, stop! And turn your attention to something else."

John Undrill (01:57:21):

And it was just, Dan, the best way I can describe it is, it was just an encouraging environment. It was built into the GE system.

[REMOVED FROM THE TRANSCRIPT]

John Undrill (01:58:37):

If you go on the technical track, you're not expected to become a manager. You're expected to have a technical lead but not be burdened with management responsibility. I was never asked which one I wanted to go on. Never told them which one I wanted to go on. But when I finally realized which one I was on, I found I was on the one I wanted to be on, which of course was the technical track. So, you asked who did I meet and what was it like? Well, that was what it was like. I met and had as my immediate manager Dr. Kirchmayer, the most important technical lead I had was Paul de Mello, who was my technical father.

John Undrill (01:59:51):

If I got frustrated and couldn't get something to work and I wanted to go and see what somebody else thought, I could walk down the corridor and knock on the door and walk into Charlie Concordia's office and change my point of view pretty quickly. When he would explain things to me, Gabriel Kron went, and Harvey Happ would show up and get in an argument. I shared an office with Happ, Kron, and Kron and Happ would argue about network solution technique in great detail and quite loudly, and I learned a lot from them. You couldn't turn around without bumping into somebody who could show you things that you otherwise wouldn't have been able to see.

Julie Cohn (02:00:57):

It sounds like a marvelous experience and a fantastic way to enter a profession.

John Undrill (02:01:07):

It was, yeah. Agreed. It sounds, it sounds like a really, really exciting place to be. Yeah.

Julie Cohn (02:01:15):

Thank you so much for spending this time with us. We are grateful for your anecdotes, your reflections, your projections for the future.

[End of interview]