J ANDREW MCCAMMON ORAL HISTORY

COMPUTERWORLD HONORS FOUNDATION INTERNATIONAL ARCHIVES

Transcript of a Video History Interview with
J Andrew McCammon, Ph.D.
Pioneer in Theoretical & Computational Chemistry,
University of California at San Diego

Recipient of the 1995 Cray Research Leadership Award for Breakthrough Science

Interviewer: David K. Allison (DKA)
Curator, Division of Information Technology and Society, National Museum of American History, Smithsonian Institution

Location: San Diego, California

DKA: You said you came from a family with a tradition of engineering, and then you went into science. Maybe you can talk a little bit about that, and your interest in studying science as a college student.

AM: It really goes a long way back. There's more there than you really want to hear, but I grew up as one of the children in a family of a civil engineer. My father was a professor of civil engineering and worked in private industry. He was very interested in bridge design and things like that.

I was a troublesome child early in my life, like many scientists seem to have been, failing years of school and having trouble with mathematics, not learning multiplication tables and so forth. But my parents were always very interested in trying to expose me to scientific things, taking me to museums and seeing exhibits. They wanted me to explore what modern technology and physics and chemistry and so forth could do. And similar to many people in my generation, I think my interest in science really took off in the post-Sputnik years, when there was a great wave of concern in society of trying to bring children into science to compete with our arch rivals in the Soviet Union. I was still at the stage where I was not interested much in what went on in the classroom, my real entertainment was reading comic books, and suddenly the comic books were all about people exploring space, and space stations orbiting the earth. Suddenly I became interested in what one could do with tools, and went from being a failing student to a student who was doing very, very well in classes.

DKA: About what age was that?

AM: That would have been when I was in the fourth and fifth grade, probably in the early 1960s. My parents tried to encourage that further, and confused me in a way that lasted to this day by giving me my first toy chemistry set and my first toy microscope at the same time. So I never was sure whether I was a chemist or a biologist. I would look at leaves, but also make reactions in test tubes. And like many kids, I was really more interested in bizarre colors or explosions I could generate, than deep insights into nature. But that was really the beginning of my interest in science as I remember it, and from the outset it was very inter-disciplinary. Again, my father tried to reinforce that by seizing any glimmer of interest I might show in something - for example, if I was interested in solar mirrors and was trying to construct one to create high temperature to burn pieces of paper - he would get out the encyclopedia and look up the equation for a parabola, and show how the parabolic design really was founded in some sort of mathematical principles. So early on, during elementary school years, sixth, seventh, eighth grade, I began to have an interest in the quantitative side of things, the mathematical underpinnings of the experiments that I was doing. And that really continued through high school and into college.

I think it was really as an undergraduate student at Pomona College that I became interested in computers. There was an organic chemist on the faculty of Pomona College who was one of the early pioneers of trying to use computers to help in the design of new pharmaceuticals and other useful compounds - Corwin Hatch. He used an empirical approach, a statistical approach, in which one looked at a series of potential pharmaceutical compounds, and tried to decide which ones had the optimum effect, and then correlate that into some sort of simple statistical correlation with simple chemical properties of the compounds – their size and other simple features.

DKA: How were you exposed to that as a student? Did he bring you in to see that process at work?

AM: Yes, we were really exposed to that in two ways. One was through the classroom, we had the advantage at Pomona College as one does at certain places of having relatively small classes, where the students were able to work closely with faculty members. We would go through normal textbook-like material but then supplement that with some of the ongoing research. It's not a new point, but one we're emphasizing here is that there is a strong tie between research and teaching. We who do research often try to emphasize that this is important in the educational process. And Dr. Hatch is a good example of it because he was able to bring some of his current research work into the classroom. He would include that in some of the lecture materials. He would show some of the things that were actually going on, and were proving fruitful in society at that given time. Of course that captured my imagination and others'. It made clear that what we were learning in the classroom really did have some real potential for real usefulness in society in medicine and industry and other applications.

In addition to that, Pomona was very fortunate to have a grant from the National Institutes of Health that sponsored summer research activities for the students. Again, one of the unfortunate trends in today's society is, and these things are always cyclical, but there's been a diminishing of federal support for education and research, and a lot of these programs have dried up. But in the mid 1960s, there were programs from the National Institutes of Health and the National Science Foundation to allow opportunities for undergraduate students to spend time in the laboratories actually participating in research activities. So we got a more intensive exposure to some of the uses of computers and other tools through intensive summer research participation.

DKA: What types of computers did you get to use in that program?

AM: The earliest computers that we used at Pomona College came around 1965, and they were pretty primitive by today's standards. They were, in fact, computers where you would program them by using electrical cables that you would plug in to sockets, and you would rearrange these cables if you wanted to change the program of the computer.

Soon thereafter we got an early IBM machine, one that had tubes and so forth, and it had these punch cards that you could use instead of changing the plugs in the sockets. So during the course of my undergraduate work we saw one wave of advance in computing, and then of course later on as I went into more advanced undergraduate and graduate work, we moved on into being able to use some of the super computers of the day.

In my graduate and especially post-doctoral work, we were using a machine called an IBM 36091E, which at the time was the super computer of the era. The machine was actually located at Columbia University, and we would use it over the telephone lines from Harvard where I was doing my graduate work. And occasionally we would have to visit the computer site to pick up magnetic tapes of data. It was an impressive machine to look at, it was fairly large in size, and it generated a lot of heat. So it was in a room that had tremendous air-conditioning capability, and the whole room kind of throbbed with these air-conditioners. You had the sense of enormous power. Of course the same computer power now is something you would have now on a laptop machine. And the calculations we struggled to do in the mid- 1970s, are the things you can do on a laptop computer traveling across the country now.

DKA: What made you decide to go to Harvard?

AM: As I mentioned before, I was very interested in the interface between biology and chemistry and physics. I had pretty much decided by the end of my undergraduate work that I wanted to try to use physical chemistry type methods to study biological behavior at the molecular level. There were several outstanding schools in the country to consider for graduate work. By that time I was married, and my wife was an English major, also from Pomona College. And I think the decision to go to Harvard College was largely based on opportunities for her. There was a very good publishing procedures program that Radcliffe College offered. So she was able to go participate in that and then look for a position in publishing. In fact she worked for MIT's *Technology Review* magazine for a number of years. So it was determined partly by extra-scientific factors, but clearly there were a number of good people on the faculty at Harvard working on this interface area, between chemistry and biology.

DKA: Were they as interested in the computing aspects of it as you were, or was that something that you brought?

AM: That really came somewhat later on. My graduate career was terribly complicated in a number of respects, and it's something that I often try to present as an example to students who are faced with their own difficulties in graduate work. I started at Harvard in 1969, and during my first year I was drafted. It was during the Vietnam War. I was a conscientious objector and so I did one year of graduate school and then I found the nearest hospital, which was Massachusetts General Hospital and went and applied for an alternate service type position.

I simply assumed that I would be doing some sort of degrading and loathsome work for two years, but when they found out I'd just received my master's degree in physics at Harvard and had background in chemistry and so forth, they thought perhaps they could use that to some advantage. So they assigned me to work for two years in a bio-chemistry laboratory. This was far away from the things I had planned to do. I was really working with test tubes and columns, and freezing in cold rooms playing with protein molecules. Yet it was good experience in retrospect, because it taught me to think about the actual physical behavior of these protein molecules in a way that I might not have if I had just dealt with purely theoretical issues. But in the course of the graduate, and in particular this alternate service-type work, I became very interested in antibody molecules - the key component to the immune system. One of the things that came up in the work I was doing at Massachusetts General Hospital was the fact that antibody molecules are in fact quite flexible molecules. They have very distinct globular domains of protein that are linked by rather thin segments. And these globular domains can move around with respect to one another. That was something that even at the time I thought, "Perhaps one could use computers and theoretical methods to study how these parts of the molecule move with respect to one another, and what that might have to do with how these molecules actually work."

When I returned to graduate school, I wasn't able to pursue that idea right away. I wound up more or less by accident working with someone that I hadn't really planned to work with as a thesis supervisor. Originally I had planned to do experimental work in applying physical chemistry to biological systems. And I found one potential PhD advisor who was an assistant professor, who unfortunately did not get tenure. Then I found a second junior faculty member - I'm a slow learner - who also did not get tenure. At that time I was fairly desperate about what to do. I'd lost two years already. I was on my third advisor, and we had a visiting faculty member, John Deutsch from MIT who's an eminent person in theoretical chemistry. Professor Deutsch agreed to take me on as a student and I wound up more or less accidentally learning statistical mechanics, which is a branch of theoretical chemistry that has to do with the behavior of large ensembles, large collections of atoms. More or less by accident I wound up learning material that put me in a position that I could go back and think in a way that I could not have otherwise have thought about the behavior of these protein molecules. So after working with Deutsch and learning statistical mechanics, I was able then to do post-doctoral work with Martin Karplus. He was a Harvard faculty member, again an eminent theoretical chemist, and Martin Karplus was just getting interested himself in the use of computers to study biological molecules. It was not a brand new field, but a relatively new field.

One of the things that I was able to bring to this study, I think because of my accidental work with John Deutsch, was a perspective of how the dynamics of these molecules might be described in very quantitative terms using equations from the "Theory of Diffusion" and equations from the "Theory of Newtonian Dynamics." I really tried to describe the atomic motion in protein molecules, and the overall motion of protein molecules in a fairly rigorous fashion. In order to do that however, we needed to rely on the power of what were then the super computers of the day in order to solve these equations and describe these motions in a quantitatively accurate fashion.

DKA: It was not entirely clear to me from the material that I have when you were able to start doing this. Was that your thesis problem, or did you not really get into this in a significant way until after you finished the thesis?

AM: The actual work on the simulation of the motions of atoms and proteins was something that really started after my thesis work. The thesis work was mostly concerned with how molecules like proteins moved through a fluid environment, and that has turned out to be a very useful thing to consider in much of what we've done recently. But the simulation of the detailed motion of the atoms, and within protein molecules, was something that I really pursued in the post-doctoral work.

DKA: My sense is that this passion that you had to understand the physical behavior, as well as the chemical behavior, and see the bridging of those is somewhat unusual. Did you feel that it was unusual, or were there other people who were interested in this? It just seems pioneering, particularly at that phase.

AM: We did have a sense of doing something new when we tried to take particularly basic equations of motion from physics - very fundamental equations, like Newton's "Equation of Motion" or the "Diffusion Equation," to actually try to take such equations and apply them in a detailed way to individual protein molecules, individual nucleic acid molecules. It was something that had not really been done before. It's something that even at the time, we had a sense of bridging, and in a rather philosophically interesting fashion, two very different areas of science. We had no idea at the time what this might be useful for if anything. It was purely abstract research at the time, and perhaps it's a good example of how something that starts out as pure research, not goal-oriented in any sense whatsoever, turns out to be tremendously useful in practical applications later on. But we saw this work as an opportunity to use the power of computers, which of course is something that's always a moving marker, and we were working with the best machines of the day. But with the power of these machines we were able to bridge the basic concepts of physics on the one side, with the basic elements of life on the other, and say something about how these molecules actually function.

I remember quite vividly when we did the first simulation of the atomic motion in a protein. We did not have nice graphics displays of these things, but we were able to output some of the structures of proteins to pen plotters and rather laboriously draw one snapshot of a protein, and then a snapshot of what it looked like a little time later, and a snapshot of what it looked like a little time later. There was a sense, even at the time, of something truly historic going on, of getting these first glimpses of how an enzyme molecule for example, might undergo internal motions that allow it to function as a biological catalyst.

DKA: Was the response from the scientific community to what you were doing enthusiastic? Was it skeptical? How were you received in the early phases?

AM: When we first tried to report this work, it was really met with a mixture of skepticism and some enthusiasm. The greatest skepticism really came from people who had done the most to determine the structures of protein molecules using experimental methods. This is through the field of X-ray crystallography. There's a very important and now rather venerable field in which people isolate the protein molecules, things like hemoglobin or other proteins. They isolate them and grow crystals of them and by shining X-rays at such crystals, and understanding how the X-rays scatter off the crystals, you can determine what the structure of the underlying molecules is, the actual three-dimensional architecture of these proteins. Now these protein structures were always depicted on the cover of *Scientific American* and journals such as *Science and Nature* as very static images, and so the great majority of the scientific community tended to think of proteins as being fairly rigid structures. If one wasn't looking at a picture of such a structure, then one might be looking at a brass model made of rods and bolts and nuts and again, a very rigid structure. So most people who dealt with proteins, if they thought of their flexibility at all, tended to think of it as very limited, tended to think of these molecules as quite rigid. So because of that, when our first studies of protein dynamics showed rather large motions within these molecules, there was a sense that something was terribly wrong in a part of many people in the community. I should say that some of the very best crystallographers, people who knew the most about the Theory of X-Ray Crystallography and so forth, knew that what they were looking at in these static images was just kind of a time average picture, and they were not surprised that there were fluctuations around that time average structure. Although I think even they were a little surprised at how large some of these motions are.

DKA: And you were beginning to understand this even in your years as a post-doc at Harvard?

AM: Yes, this was part of my work as a post-doc at Harvard, and it's something of course that I continued to develop very vigorously once I moved down to the University of Houston as a faculty member.

DKA: Did you have any question about wanting to go into university life or did you think about maybe working for a pharmaceutical company or did you have a specific career goal?

AM: I don't know that I ever really thought about it in a very focused way, but I was always interested in doing research, and having my hands on research problems and trying to choose directions for the work to go. I think that pretty much dictated trying to pursue a university career if possible, simply because one had the freedom to sort of set directions and go after personal goals somewhat more than one might in an industrial setting or another setting.

At the same time, I had some interest in trying to see if there were not practical applications for the work we were doing. But I didn't want to be constrained to try to have some definite answer to some practical question in six months time. I was more interested in taking a long view what we might be doing in five years, or six years, or ten years, and really unless you're very, very fortunate in an industrial government laboratory, universities are the place where one is able to take that kind of long view.

DKA: When you went to the University of Houston, what research environment were you heading into? Was it an environment that was favorable to the work you were trying to do, or did you realize you were going to have to come in and build a structure that would allow you to continue to work?

AM: When I went to the University of Houston, I recognized that I was moving into an environment where I was going to have to build a lot of the infrastructure to support our research program. I was very fortunate when I looked for a university position to have more than one choice. It was one of those times in the job market cycle when things were very tight. Again, the decision to go to Houston was really based largely on our two career family concerns. My wife had had decided at that time to leave the publishing business. She had gone back to do her pre-medical courses in medical school, so she was interested in finding a place to do her residency program in neurology. Among the two or three places that I could think of to find a faculty position, the great strength of Houston was that there was a very large medical center there and some medical schools where she could continue her training.

So I reluctantly declined offers from some schools that had at that time a stronger reputation and perhaps more resources. It was clear when I went to the University of Houston that it did have some advantages, but it was going to require a lot of work to really get a research program up and running. One of its advantages was that it was a very hungry school. It was a young school, hungry for reputation, it was hungry for achievements of scientists, and so it fostered a kind of entrepreneurial setting, people who were willing to roll up their sleeves and try to get things done would find some support on the part of the university administration.

And that turned out to be very important for me, in that when I first went to the University of Houston they had rather poor computing resources, in particular, I won't name the brand of computer they had, but it was a type of computer that one usually associates with the banking industry rather than a high performance scientific computing environment. So one of the first things that I did was to start working on the administration to upgrade the computing resources at the university. And I'm afraid I sometimes was a little bit devious in trying to influence the administration. After several pleas to the university administrators to try to bring their computing resources up to snuff, and not getting the reaction I wanted to, I composed and sent a telegram to the president of the university, saying, "Computing crisis, have grant from NIH, can not do research, NIH may investigate, situation desperate." And within 48 hours there was an allocation of between one and two million dollars approved to buy a new computer for the university, and I was given significant say in what type of hardware to acquire. So happily within a couple years of being at the University of Houston, we were suddenly near the forefront in terms of computing resources, and able to do calculations that were pushing the envelope a little bit in terms of possible applications.

DKA: What did you buy with your two million dollars?

AM: Well we bought an AS9000, a machine by National Advanced Systems. It's an IBM compatible type machine which was fine with me, because I'd grown up largely with IBM technology. So it was familiar, something we could move on to. And we were able to do quite a bit of work on that machine before we really needed to take the next step, which was to what we think of now as real super computers, the Cray Machines and things like that that really have traditionally set the top end of these machines.

DKA: I want you to talk a little bit about the method that you pursued. See if you can just put it in a fairly simple form that people can understand what a thermodynamic cycle perturbation is. That's a big mouthful, but it's critical to what you did.

AM: There's an interesting story behind this thermodynamic cycle perturbation method. As I mentioned, when I first went to the University of Houston, we were already simulating the dynamic motion of the atoms in protein molecules, and looking at protein molecules basically as any kind of material system that you might look at. There was no real driving societal need, interest or anything. It was pretty much pure research, "How big are the displacements of the atoms? What's the time scale?" and so forth.

My interest in doing something beyond that did stem in large part from a personal situation, which was that when I was a still relatively young faculty member at the University of Houston, the wife of a colleague of mine, a very good friends of ours, was stricken with pancreatic cancer. Fortunately she's recovered and in very good health, but I remember feeling very intensely at the time how powerless I was to do anything to help, and how absurd that seemed that here we were, studying these basic molecular components of living systems, and surely there must be some way we could use this knowledge to help people that were stricken with cancer or other diseases. So I remember thinking quite intensely about it at the time, that there ought to be something we could do in our research to steer it toward a useful end, toward the design of new pharmaceuticals. And this thermodynamic cycle perturbation method really was the outcome of a lot of thinking about such issues.

Many, perhaps most pharmaceuticals, are small molecules that act as what are called enzyme inhibitors. Enzymes - it's a term that almost everyone has heard - enzymes are simply protein molecules, large biological molecules that act as catalysts. A typical enzyme will bind to some molecule, maybe a sugar molecule, and break it up into smaller pieces, or maybe combine it with something else. Enzymes facilitate chemical reactions. Now most pharmaceutical agents act as inhibitors to enzymes. It's like throwing a molecular monkey wrench into the works of an enzyme. If you have an enzyme inhibitor, it will typically bind to the so-called active side of an enzyme and block it so that the enzyme will no longer function as a biological catalyst. Familiar drugs like aspirin and drugs for ulcers and most cancer agents and other things are inhibitors of one enzyme or another.

So an important question in a pharmaceutical industry is if you have one pharmaceutical molecule that works pretty well, are there ways you can tinker with that pharmaceutical to make it bind even more strongly to its target? And that was a question that I immediately began thinking about in connection with this young woman's bout with pancreatic cancer. Was there some new method that we could come up with using computers, and using these simulations of enzymes and other molecules to help guide the design of new pharmaceuticals? We took as a simple model problem, a very common enzyme, one called trypsin, and a very small molecule called benzamidine that binds to trypsin very specifically. And we wanted to see if we could calculate the strength of binding of benzamidine to trypsin, and then compare it with other modified forms of benzamidine that might bind more or less strongly, as a simple model system that we set. And we tried to use our computer simulation methods to calculate these strengths of binding.

It's very hard to do this in a brute force or direct fashion. We had the enzyme in a big box of water molecules, and had the inhibitor molecule, and tried to push the inhibitor molecule into the active side of trypsin, but we were never really able, for technical reasons, to get a very realistic depiction of this binding process.

Basically we were trying to make it go too fast in the computer, and the enzyme was not able to relax in the way that it really would in a binding process. Our computer simulations were looking at, again because of limitations in computer speed, we were looking at picoseconds, or tens or maybe hundreds of picoseconds, very very short times, much, much smaller than the actual binding process occurs in the laboratory. So we were not getting right answers. We were not doing physically meaningful calculations, and I remember spending more than one sleepless night worrying about that, "How can we do this? How can we do this?"

One day I went to a lecture in my other field of statistical mechanics, going back to the John Deutsch type work, and I remember listening to someone talk about ways in which you can make a small cavity in liquid water have different sizes and calculate changes of certain thermodynamic properties. I was sort of halfway dozing through the lecture. It was quite technical, and it was after lunch. I was feeling a little sleepy, and suddenly I remember sitting bolt upright in my chair because I realized that what this distinguished professor was talking about was exactly the missing link that we needed to do the calculations we really wanted to do. What we really want in some sense is not to answer the question how strongly does a potential drug molecule bind to a target, but what binds better, this drug or this other possible drug? We want to put things on a scale of relative binding strengths. And it occurred to me that using this technique that this eminent statistical scholar was describing, we could recast the whole problem we were trying to solve in a way that was much more amenable to a solution - and that is we could start with our drug molecule bound to the active site of the enzyme and simply change it from one drug molecule to another one while it's located in the binding site.

So for example, we could take a hydrogen atom on a drug and slowly change it into a fluorine atom. The name computational alchemy has become popular for these kinds of methods, because you really are in these calculations, changing one kind of element into another type of element. It's not something you can do in the laboratory. But because we're using computer models, because we're using mathematical representations of these systems, it's perfectly possible to take one of these computer simulations and just really change the size of atoms, change the lengths of chemical bonds and actually transform, or in modern parlance, to morph one molecule into another molecule and calculate the thermodynamic changes associated with that. Now, if you use a basic principle of thermodynamics, it's possible to show that the changes in thermodynamic properties that you get by these artificial, or non-physical changes of systems, can be related to these changes in free energies of binding, which is what we're really after. What we were able to do in some sense is to finesse the calculation, to break it down in a different way that was more amenable to solution on computers of current power.

DKA: That was the trick wasn't it, to learn what that mapping was, between the binding energies and the thermodynamic properties? Did that just occur to you, or was it known in the literature?

AM: In some sense we were using a set of ideas which were very, very old ideas, but just combining them in a completely new fashion. We were using an idea that goes back to even high school chemistry classes, of a so-called thermodynamic cycle idea, that if you have a set of states of a chemical or physical system, there are certain functions that have to have no net change when you go all the way around one of these thermodynamic cycles. And what we were interested in was two parts of the cycle: the binding of molecule A and the binding of molecule B, and we were able to relate that to the other two sides of this square where we change molecule A *into* molecule B. The idea that there exists such thermodynamic cycles was very familiar, but it was very unfamiliar to look at non-physical transformations on side of that, and to think about binding processes in that setting. That was completely new, and it was just one of those momentary insights.

DKA: Well that, plus the way that you spoke about using the simulation and seeing and doing almost as if it was the equation, must be something that became familiar to you, but was somewhat unusual at that time period, was it not?

AM: The idea of combining the underlying thermodynamic theory, and combining these computer simulations with experimental data on the structures of these molecules, it was something that we were doing very much in isolation at that time. I think really our group, and the people in our group who were doing this, were the only people in the world who were trying to do this particular type of calculation at that time. It's something that now has become a very, very widespread method, and essentially, every large pharmaceutical company in the world carries out calculations of this kind. It certainly is not a guaranteed way of coming up with ideas for new pharmaceuticals, but it's a way to help people make more educated guesses as to what compounds to try to synthesize and test. It's very widespread now in the pharmaceutical industry, and many universities and medical schools. There must be hundreds of laboratories that carry out these types of calculations. And happily, they are beginning to produce new pharmaceuticals.

As you probably know, there's a very long development time in the pharmaceutical business. If you have a compound that's eventually going to prove something useful in the clinic, there's about a ten year period that that compound has to be tested with cell cultures and laboratory animals and so forth, and then a long time in the clinic before it's actually blessed with the term pharmaceutical or drug and put on the pharmacy shelf.

But there are a number of compounds now for treating influenza, HIV infections, emphysema, a number of diseases, and there are a number of compounds in advanced stages of clinical trials that were really discovered by the surge of computer simulations. So the hope is that with time, some of these will actually be beneficial.

DKA: Maybe you should talk a little bit about how the change in computing power resources allowed you to mature this method, and what it first looked like, and how that grew and changed in your years at Houston.

AM: Much of what we have done has been trying to take advantage of the newest generations of computers. There are people who argue that there's no longer any need for more powerful machines. Small computer workstations are already as powerful as Cray's used to be, so why go any farther? They're partly right. Certainly problems that it took a Cray computer to do five years ago are things that can be done on a powerful workstation, or maybe a cluster of a few workstations today. But there are always going to be problems that require that single-most powerful machine to begin to think about. It's easy to think of problems like that.

Certainly one of the things we've done in our research group, is as we think about what areas of research to pursue. We try to think about what the most powerful computer is going to be like in five years time, or ten years time, and say, let's start working on that problem now, even though we know we won't really be able to solve that problem right away. But when that machine appears, we'll be in a position to really take advantage of that resource and do something new and interesting.

When I mentioned this work on trypsin binding to benzamidine, this initial thermodynamic cycle type work, the only reason we were able to undertake that project was because we had access to the National Science Foundation, to one of the great supercomputers of the day. That in fact was one of the Cyber205 parallel supercomputers that was at Princeton University and the NSF supercomputer center that unfortunately no longer exists. But when that machine was initially set up, we were granted a very generous allocation of time on the machine. And only because we had this extensive block of time on the most powerful machine of its day were we able to think about trying to set up this extremely challenging calculation, where we had a whole enzyme molecule and a box of thousands of water molecules and think about drugs binding to it. It's not a calculation you could have done on any other machine, and there will always be calculations that are waiting for that next machine to do.

DKA: So you were moving into an environment of increased calculations that allowed you to solve these problems faster. What was the visualization side of that as you came in, and how did that develop along with the increase in computing power?

AM: The visualization side of things is very, very important in this kind of work. We're dealing with, as I mentioned, enzyme molecules and their inhibitors. Such molecules are very large. They typically have thousands of atoms in them. They're very irregular in shape, and in order to visualize what's going on, we need state of the art computer graphics capability – not only to visualize the molecules themselves and the motions of these molecules, but also to display data relating to our calculations. If we're interested in for example, the electric field around one of these molecules, which may be very important to explain its properties, then we can use graphical techniques to study such things. At the outset our graphics capabilities were extremely limited. Again, going back to the first molecular dynamics simulation of proteins, we simply used pen plotters to draw just a few snapshots, to give us sort of a hint of what might be going on in the internal motions of a protein. By the time we were doing the simulations at the machines at Princeton we had slightly better computer graphics capability. We were able to make single colorful images of the protein molecule and this inhibitor bound to it. But we were not really in a position where we could easily make fully dynamical images of these molecules, and we've depended really on the development of computer graphics technology alongside the basic computation technology to try to help us get more out of the simulations as we conduct them.

One of the advantages of a place like the University of California in San Diego with the San Diego supercomputer center is you have not only state of the art advanced computing capability here, but you also have people who are experts in the visualization of the data, and you have resources to look at these images and three-dimensional stereo vision. You gather around with a group of colleagues and point to these structures in a three-dimensional representation, and discuss and argue about the results of calculations, plan the next steps and calculations. So computer graphics are very helpful in many, many ways, but among them are understanding the data, discussing it with your colleagues in an interactive fashion and of course, presenting the results to students to use in the classroom to help us as teaching aids, and to present the results to the public to help them understand what's going on in these computer simulations.

DKA: Why don't you describe what we're looking at.

AM: Well one of the enzyme molecules we've been working on most recently is shown here on the screen. This is the enzyme acetyl cholinesterase, and it's an enzyme which is really critical to the operation of the nerve and muscle systems of humans and all other organisms. When one nerve talks to another nerve cell in the brain for example, there typically is a small space between the nerves called a synapse, and when a nerve is excited and it brings its electrochemical pulse down to the synapse, there are neurotransmitter molecules that are released into the synapse. Those diffuse across and excite the next nerve on the line and that's what causes, for example, a thought to develop. Similar synapses operate where nerves come into contact with muscles and are responsible for all neuromuscular activity. Obviously these synapses have to operate very, very quickly. If you touch a hot stove you want to respond instantly. You don't to wait and wait and wait and then respond. So there's been tremendous evolutionary pressure to have the nerve and neuromuscular systems operate at lightning fast speed. Well, this molecule plays one key role in the speed of activity of the neuromuscular and nervous systems.

I've mentioned that in the synapses, there are neurotransmitter molecules that are released, but you need to get rid of those neurotransmitter molecules very quickly. Essentially it's like having an on switch and an off switch. When you release the neurotransmitter molecules, that's an on switch. It turns on the following nerve or muscle, but also you need to be able to turn it off very quickly so you're poised for the next idea or the next response. In order to turn the switch off, you need this molecule which basically destroys the neurotransmitter molecules in the synapse, cleans it up and restores it to its original state so it's ready for another nerve or neuromuscular excitation.

So this enzyme, acetyl cholinesterase, has one very simple but very important job, and that is to break down the neurotransmitter acetylcholine. Acetylcholine is a small molecule, which binds to this enzyme and then it's broken up into two parts and reduced inactive. The acetyl cholinesterase enzyme shown here is actually not from a human form. It's actually from a kind of electric ray, one of these bat-like fish that swim off the coast of southern California, *torpedocalifornica*. And electric fish and electric rays use these same synaptic mechanisms to build up the large voltages associated with their electrical activities so they're good sources for proteins such as this. This acetyl cholinesterase molecule is a dimmer. It has two identical subunits, one on the right-hand side, and one on the left-hand side, but they're facing opposite directions. The one on the left-hand side is facing with the entrance to the active side pointing toward us, and the one on the right-hand side has its active side pointing back away from the screen. The thing I want to call your attention to is that the active side is not really on the surface of the enzyme. It's actually buried very deep in the middle of this large globular structure. And it has a small cave-like entrance to the active side that you can just barely make out where I'm pointing here.

The acetylcholine molecule has to enter this active site and move down a long tunnel in order to get to the catalytic center itself and then be degraded. Well if you actually make a model of the acetylcholine neurotransmitter, and if I were to display it on the same screen, it would be too large to fit into this hole. It would look like the neurotransmitter could not possibly get from the outside of the enzyme deep into this catalytic center. So what's the trick, how does the enzyme mange to handle acetylcholine and handle it so quickly? Well, part of the answer is in these thermal motions of the enzyme. As I mentioned, in our early molecular dynamic simulations, we were able to see that the motion of the atoms inside proteins is of surprisingly large amplitude. And if you do a molecular dynamic simulation in which all of the atoms of this protein are allowed to move due to their thermal energy, you find that this opening breathes, essentially. It opens and closes and occasionally it opens wide enough that neurotransmitter molecules can get into the enzyme and move on down to the active site. And do what I'll show you next will be a set of snapshots that will just illustrate this breathing motion.

What you see in this series of three snapshots are actually frozen images taken from one of these molecular dynamic simulations. We've taken the enzyme molecule, this acetyl cholinesterase molecule, we've put it inside a large box with thousands and thousands of water molecules, and we've allowed all of these atoms to move around according to Newton's Equation of Motion just with the kinds of thermal energy that they would have at room temperature. Then every now and then we'll take a snapshot of the system and display the molecular surface near this entrance to the active site. And what you can see is that as you go from one moment in time to another, you can see how the active site entrance is somewhat opened up in this image. A few moments later it's much smaller in width, and then a little time later on it's beginning to open up again. So protein molecules are not the rigid structures that people thought for many years, they are in some sense, living, breathing entities of their own and these motions are absolutely essential for the functions of enzymes and nucleic acids and really all biological molecules.

I've described how the internal motions within this acetyl cholinesterase molecule play a critical role and a function of the enzyme. Here's a second part to the story, and that has to do with the electric fields around the enzyme molecule. What we've done now is we've taken this whole molecule, we've rotated it, so now the active site of one sub-unit is pointing up this way, and the active site of the other sub-unit is pointing down. We have taken account of the fact that all of the atoms of this protein molecule have small, electrostatic charges on them and have described here the electrostatic potential around this enzyme molecule that arises from these charges. In simple terms, many of the atoms around the entrance to the active sites have negative charges on them, and so there's this large region of negative electrostatic potential. Why this is important is that acetylcholine, the substrate, the neurotransmitter that gets destroyed by this enzyme, has a positive charge on it, and it's familiar to everyone that opposite charges attract one another.

So the positively charged acetylcholine molecule is pulled into the active site, is actually steered toward the active site by the electric field around the enzyme molecule. This enzyme gets its tremendous speed really from a combination of two factors. One is this initial electrostatic steering of the diffusional encounter of the neurotransmitter with the enzyme. Then the next part is the breathing motion of the enzyme that allows the neurotransmitter to move down into the active site and allows the subsequent reactions to proceed so rapidly. This electrostatic steering contributes something like a factor of 10 increase to the speed of the enzyme.

Again, this is an enzyme that has operated under tremendous evolutionary pressure to work as quickly as possible, and nature has achieved a tremendous speed of this enzyme by this initial electrostatic steering, and then this dynamic motion that facilitates the actual reaction. So in the studies of acetyl cholinesterase, we've provided at least part of the answer to a very basic biological problem, as how do nerves and muscles work so quickly? Why are reflexes seemingly instantaneous? But there's more to this story, because this enzyme turns out to be a potential important target for pharmaceutical work.

I've mentioned that these simulations are of great value and helping us to deepen our understanding of biology, but there's a practical side of this, too. There are a number of diseases in which this enzyme plays a role in one way or another, one that's familiar to almost all Americans, is Alzheimer's disease. We all have relatives or know of people in advanced years who begin to lose their memories of recent events, and it's a progressive neurological degeneration. The actual cause of Alzheimer's disease is not well understood at all. But one of the aspects of Alzheimer's that is known is that the level of this neurotransmitter, acetylcholine, drops to dangerously low levels in the central nervous system, in the brain. So if there's any way you can bring the level of neurotransmitter back up you should be able to restore some cognitive function. In fact, one way to do that is to slow down the action of this enzyme, so you're not destroying the neurotransmitter so quickly, that allows the acetylcholine level to come back up. The only drug that's on the market now for the management of Alzheimer's disease is actually an inhibitor of this enzyme. It's a drug called Tacheryn or Cognix, it's widely prescribed. Not all patients tolerate it because it's quite a toxic, rather experimental compound, but it is widely used. It's a small molecule that binds in these active sites and is this kind of molecular monkey wrench that blocks the action of the enzyme to a limited degree, and allows the neurotransmitter level to come up. For patients who can tolerate this drug, it's a small miracle drug. A patient that one day is confused and disoriented and can't recognize her husband or children, after a few weeks of therapy with this drug, is often times able to recognize relatives again and to participate in some of human activities in a recognizable fashion. But it is a very experimental drug still, it's got serious problems with liver toxicity and serious side effects, so there's a great interest in trying to re-engineer such drugs to try to develop more specific, more effective versions of them.

So we're beginning work using these thermodynamic cycle methods to see if we can take the same target compound, the same enzyme, and hopefully develop a better generation of Alzheimer's agents; the same molecules involved as a potential target for drugs for glaucoma and many other diseases.

DKA: I'm just curious whether you're finding in other enzymes the same molecular properties, or whether this is unusual.

AM: That's probably a good point to cut to another molecule. What I'm showing here is yet another enzyme molecule that operates with tremendous speed. This is an enzyme called superoxide dismutase. Its job is basically to clear your body of a very toxic compound called superoxide. Superoxide is just like normal molecular oxygen that we breathe and depend on, but it has an extra electron attached to it, and becomes a very toxic, highly reactive molecule. Superoxide develops in the normal course of metabolism and nature's developed this enzyme to get rid of this toxic material, superoxide. Again, since it has an extra electron on it, the superoxide molecule has a negative charge to it, and this enzyme takes advantage of simple static electricity to speed its rate of clearance of superoxide from the body. Here you can see the molecular surface of the enzyme; again we have a dimeric enzyme, one with active sites on different sides of the molecule. And again the electric fields operate to steer the substrate molecule into the active site to speed its rate of action.

What's nice about this story is that we had made these predictions a number of years ago that the electric fields around the enzyme might speed the diffusion of the substrates into the active site, and this stood as just a theoretical statement for two or three years, until in the early 1990s, a team of investigators headed by Elizabeth Getzoff here in La Hoya, at the Scripps Research Institute, essentially took up this as a challenge. They said, "Well if this enzyme operates with tremendous speed because of the electrostatic steering, and if McCammon and his coworkers were right, that the electrostatic steering really is responsible for the speed of the enzyme, we ought to be able to make even a faster version of the enzyme by what's called site-specific mutagenesis." Site-specific mutagenesis is simply a way of taking the gene that contains the blueprints for this enzyme and making certain changes in the gene so that chemical groups in the enzyme are replaced by other groups. They were able to use these mutagenesis methods to create variations on this enzyme that had increased or decreased electrical charge around the active site, and the calculations would predict that the rates would be increased by a factor of 2, increased for other changes by a factor of 4.

Libby Getzoff and her group did what we as computationalists can not do. She and her coworkers went into the laboratory and actually made these alternate structures of the enzyme and measured their kinetic properties, and showed that in fact, they do work twice as fast, they do work four times as fast.

And I think this is interesting because it really the first example of engineering at the molecular level of biology. People have talked for years about protein engineering, or genetic engineering. As I mentioned earlier, I come from an engineer's family, so I take this engineering stuff very seriously. One of the people on my dad's PhD committee was a person who designed the main span of the Golden Gate Bridge in San Francisco. One of my dad's most prized possessions is a book on the design of suspension bridges. It's an old book and in the middle there are about thirty pages that are worn and yellowed with repeated use, and if you open up those pages, what you see are equations – pages and pages of equations. Mathematical principles that connect physics to some structure that is not yet built, but when built will behave in a certain way, that's engineering to my mind. To have a deep enough physical understanding and the supporting mathematical structure to design something that will behave in a predicted way. And this is the first time that that's been done in biology. Libby Getzoff and her group, I think, deserve tremendous credit for carrying out this first real feat of bio-molecular engineering, of taking physical theory and mathematical computational principles, and using them to actually build a new enzyme that works as specified.

This is another enzyme molecule that we've been studying in our laboratory. There are perhaps two things to say, one is that this is a little bit different representation of the enzyme than you may see in some other images. Some of the other images that you'll see today have a kind of a white sheet wrapped around them, and those other images sometimes I call Cristo renditions. There's this popular modern artist who likes to wrap islands and wrap German Parliament buildings and things like that. Here we really pull the wrapping off, and so you're looking underneath the molecular surface to see the actual bones of the enzyme molecule, if you will. This is the polypeptide chain Alpha helix that were predicted by Linus Pauling before they were actually discovered in experimental structure, alpha helix and other elements of the structure of the protein.

The enzyme that you're looking at here is one called adenosine deaminase, and it's again an enzyme of very great medical importance. There are unfortunately certain children who are born with certain defects of this enzyme and they suffer from something called severe combined immunodeficiency syndrome. Those children will normally die at a very, very young age and infancy unless dramatic measures are taken to try to restore the presence of this enzyme or to reactivate it somehow. That's one of the first targets of modern use of genetic engineering in medicine, is to try to restore the gene for this enzyme into children that have this severe combined immunodeficiency syndrome. This enzyme is also involved in a number of diseases, among them, certain leukemias and lymphomas.

What you see here in addition to the enzyme shown in green is a small pharmaceutical molecule, deoxycoformycin. Its one of the few compounds that's in clinical use for the treatment of what's called Hairy Celled Leukemia. And now that the experimental structure for this enzyme and this drug molecule are known, we're engaged in calculations in which we're trying to change the structure of this drug in certain ways to see if we can increase the binding of the drug to the enzyme. This is work that's being done by a Tammy Marone, a young woman who's doing post-doctoral work in our group, and certainly shows some promise I think for the engineering for better anti-cancer agents.

DKA: You're not only a researcher in the laboratory but also a teacher, and a part of your work has been to reform not only what teaching and chemistry means, but bringing these bridges across. How have you worked in that part of your career?

AM: One of the fun things in some sense about being a university professor is that there is this opportunity to have interplay between work in the laboratory, the research side of one's work, and teaching, the educational side of our work. We do try to continually re-examine how we teach science in view of what we're learning at the research frontiers of our field.

One good example of that perhaps is I will be teaching in a few months an undergrad physical chemistry course. It's an undergraduate physical chemistry course that will have a lot of biologists and future doctors and other people in the course who are interested in biological molecules. One way that we will try to liven this course up, it can be awfully dry just learning equations and pages of dreary text, we will try to bring some of what we've learned about our research work into the classroom, and actually set up examples where students will be able to get their hands on the computer, and carry out at least simple versions of some of these simulations as part of their undergraduate curriculum, their undergraduate studies. For example, we will try to make it possible for students to calculate the strength of binding of different drug molecules to an enzyme, at least in a simplified way, to underscore the importance of thermodynamics, and help them turn this rather dry material into something that's clearly and compellingly important for their future work. And we'll try to study the diffusional encounter of substrates with enzymes to underlie the importance of again what can be dry kinetic theory.

One of the attractions of the University of California at San Diego, one of the reasons that I moved here recently, is that there are tremendous opportunities, not only on the research side, but also on the teaching side. The San Diego supercomputer center plays a critical role, not only in supporting the nation's research efforts and computer simulations, not just in biology but in environmental modeling and studies of global warming and so forth, but it also serves as a tremendous resource for teaching, for the development of video tapes, showing examples such as this.

It can be dispersed to classrooms throughout the country and to generate model simulation exercises that students certainly here on campus, but eventually students everywhere throughout the country and even around the world can plug into the internet and come to the San Diego supercomputer center through kind of a virtual experiment and be able to do some of the kinds of manipulations that I've described here. This will have, hopefully, two very valuable results. One is it will help students really learn the basic theory in a more solid fashion, but also it will keep them excited about the science. It's that sense of excitement that is so critical to a person's being successful in their career.

DKA: You talked in a number of points about the role that teachers had for you in guiding you along your career. And I know that you've spent a lot of your own personal effort in repaying that favor to your students. Talk about some of the students who have not only studied under you but have had you as their mentor and how they've gone on to change the field of this science.

AM: Well certainly one of the great satisfactions that any university teacher will tell you about is seeing promising students really flourish in their careers, to get a good solid grounding, and certain fundamentals to learn whatever you were able as a teacher to convey to the students, and then to go on and do things that really can change the whole shape of the field. Often times they do things that you know ideally are far beyond what I might have been able to do myself. I'm very, very fortunate to have a number of students that have gone on to positions as university faculty members at the University of North Carolina and the University of Washington and many other institutions, also some smaller institutions that really concentrate on undergraduate teaching.

These former students of mine are themselves advancing the frontiers of this field, developing new applications and teaching yet new ranks of scientists who will be entering the field. Other former students have gone on into federal laboratories, and into pharmaceutical companies where they're putting these methods to work, to do the actual discovery of new pharmaceuticals, to do the discovery of new enzymes and really put these tools to work.

DKA: What is the secret of teaching them and leading them well? How do you do that?

AM: Well the secrets to good teaching are something that's really quite nebulous. I don't really have any magic formula. I have tried to develop some outlines of suggestions to students. But one thing that I always encourage them to do is to find a problem to think about as far as learning about research is concerned. Try to find a problem that is an important problem to work on as much as possible, something that they themselves are very interested in. If a student is not deeply interested in what he or she is working on, there's really no hope of progress.

So I encourage students to try to find a problem that they're really obsessed by, that will keep them up at night, worrying and tossing and turning. And then beyond that to try to learn tools that have a very wide utility - things like molecular dynamic simulations, things like theories of how molecules diffuse. These are very general concepts that are fluid enough that a student will not run out of gas just on finishing one particular research project, but the student will be equipped with tools that can be used again and again in different settings. So those are at least two key ideas.

DKA: But I sense in you also a very personal quality of leadership, in caring about your students. It seems that you're more than just someone who gives them information. Do you have that kind of relationship with them?

AM: I do care about each and every one of my students in a very, very personal fashion. I very much look at my group as a family in some sense. We try to get together on a very regular basis, and we talk largely about scientific things, but also about more personal things, too; people's concerns about the job market, people's concerns about teaching. We oftentimes have undergraduate students who will come and spend time working in the group, and we talk about ways that we can make that experience as rewarding as possible during the time that they visit. So we do try to create a strong family atmosphere in the group. I do encourage the students in our group and students in my classes to interact with one another, as much as they can, not to feel that they should advance by climbing over other people, but I do try to encourage a sense of cooperation among students. It is true that there is inevitably a degree of competitiveness in science, and that that's an essential ingredient, in fact, for people to feel challenged to do the best they can. But it's possible to be competitive in science without stabbing other people in the back and that's certainly something that I try to encourage my group to do – to work together on things, but for each person to try to put his or her best talents to use in moving the whole enterprise forward.

DKA: You've described a research environment, first at Houston, and maybe even more fully now at San Diego that is dependent on a number of factors. It's dependent on a growing and developing world of supercomputing, a research center supported by a university climate. Do you see that continuing in its present form; do you see that structure changing, what does the future look like?

AM: The universities are always evolving in the setting of society at large. Clearly we're in a situation right now where we're undergoing a large change in our general climate with the end of the Cold War. There's been a tremendous decrease in the willingness of the federal government to support basic research and science in general. There's this loss of this concern that we had in the post-Sputnik era, and so we're facing a difficult time certainly in terms of federal support, that's compounded to some extent by the current changes in the industrial sector.

Obviously the work that we do, much of it involves the pharmaceutical industry, and there are these tremendous pressures for cost control in the pharmaceutical industry that led to a reduction of support and a reduction of opportunities to some extent, for people who are interested in drug discovery. So we're facing difficult times.

One can only try to make the best of the situation by educating society as to the ultimate benefit of continuing this enterprise, to point out the value of advanced computing tools as something more than just a way to beat an international competitor in the defense arena, but to point out the potential for advanced computers as tools for conquering diseases, for controlling diseases if you can't conquer them, and for dealing with a myriad of other problems that society faces, environmental problems and many, many others.

So I think one of the changing features perhaps in our time right now is university faculties, that we need to be a bit more advocates of the university enterprise. We need to remind people of the enduring value of what we're doing, and point out that a dollar invested in a university may benefit society to the tunes of hundreds of dollars later on in terms of reduced impact of disease, in terms of reduced impact of pollution, in terms of ability to improve society in many, many ways.

DKA: In a more technical sense, as you look for the tools that you developed for now a number of years, what does the future look like? Is this a blossoming period? Are you coming to a close on an era? What lies ahead for the scientific tools that you've spent so much time developing?

AM: I think the future of the kind of field that I've been so involved in in recent years, that future looks extremely bright. In fact, I feel a little bit as often been described of Newton and other great scientists, of just having one little grain of sand in your hand and you realize that there's this whole beach and this whole ocean out there. We obviously have started to look at a few molecules, a few enzymes, and learned a few things about how they function, there's much more to do in that arena. But there's much more to be done in extending this kind of work in several directions. One is looking at the fine details of the reactions that go on in enzymes, and what's the nature of those and how those can be influenced in other ways.

My personal interest, I think, for the next few years is going to be going up the scale a little bit. I've mentioned how acetyl cholinesterase is a critical enzyme in synapses, and how one nerve talks to another nerve. There's much now that is known about the detailed structure of entire synapses, and so what I would like to do in the next few years with my research coworkers is to try to move the next level up, and to go from computational molecular biology, if you will, to computational cellular biology.

I think especially if we're able to continue to develop more and more powerful computers, if the envelop of supercomputing continues to move the way that it principally can, we ought to be able to think about integrating these molecular components into models of entire synapses, and integrate the molecular picture into a computational cell biology. I think computational cell biology founded on truly molecular understanding, could be one of the tremendous exciting areas of science in the early 21st century. It will depend on federal support. It will depend on the development of the advanced computing capability. We do need high-performance machines to really move in this direction, but given those resources, I think the next decade or two looks very exciting indeed.

DKA: I was just wondering if you had anything else to add.

AM: I can't, I think I've finished.

DKA: Thank you so much for sharing your thoughts a research with us. It has been a pleasure.

AM: Thank you.