Oral History of Ken Thompson

Interviewed by:
John Mashey

Recorded: February 8, 2005
Mountain View, California

CHM Reference number: X3091.2005

Q: I am John Mashey; I'm a trustee at the Computer History Museum. It's my pleasure to interview Ken Thompson about his involvement with computer chess over the years. So Ken, welcome to the Museum.

Ken Thompson: Oh thank you very much, thanks for inviting me.

Q: Oh sure. So let's start at the beginning. Where are you from, where did you go to school, and how much game playing did you do before you got to college?

Ken Thompson: Oh, going way back.

Q: Yes, all the way.

Ken Thompson: My father was military and the technical term for that is Navy brat, so I was a Navy brat for the first 20 years of my life. And we moved all over, I never spent more than a year or two any one place. So that was high school and then I graduated high school in Chula Vista, California and went to school in Berkley and that's essentially my educational career

Q: Okay. So in high school or before, or in college, what sort of involvement with games did you have, whether chess or other games?

Ken Thompson: I was always interested in games. I read game books and crosswords, things like that. Chess in particular, I took up kind of the year that Bobby Fisher was mopping up in the U.S. when he was – he's about one month different age than I. And you know you get this feeling that, you know, why aren't you any good? You know he's on the cover of Look Magazine at 14 and I'm going to school in Texas. So my parents played bridge and – the whole family actually played bridge, I played bridge a lot. Never liked the game -- and they would take me to this hotel once a week where they played duplicate bridge, in this little town in Texas. And there were a couple of old men there that played chess in the lobby. And when I didn't play bridge I'd go and watch them play chess. And eventually I started playing and then I read a couple books on it with all of the Bobby Fisher hullabaloo going on. Then I played for about six months, on the team in school. This is – I believe it was 7th grade; it might have been 6th grade. And I guess nobody in 6th grade ever read a chess book before because as soon as you read a chess book you're better than everybody else.

Q: Everybody else.

Ken Thompson: And then I quit it and then essentially never played after that. I was always a fan of chess, always followed it, always read the games, read the books, followed the tournaments. But after that 6th or 7th grade I never played personally.

Q: Okay. So how about other games? Any other ones that you were interested in before college?

Ken Thompson: I was interested in college in three dimensional tic-tac-toe, four by four, by four. And I had a gamer for a teacher, Berlekamp; very nice man. And I used to get midnight computer time. Basically, when it was idle I'd go in and take over the machine.

Q: What machine was this?

Ken Thompson: 7094.

Q: A 7094, okay.

Ken Thompson: And I wrote some game playing programs, one of them being 3D tic-tac-toe. And Berlekamp is still the strongest person I've ever known at that game. And I used to wake him up in the middle of the night, like two in the morning, and he'd come out of the computer room and play this game with me and beat it. I'd go lick my wounds and go back and try to do better, and try to do better. But that's – I did that. Also, there was a gamed called BridgIt which is an N by N array where N can be somewhere around 19, 20 something like that. And one player connects red lines on the screen and the other player connects blue lines. And the red player tries to connect the bottom row, wiggling up to the top row. And the blue player tries to block and connect left to right or right to left. And the first person to make a connection wins. And again, there were some people at school that were very, very good at this, and this was my first game at cheating. Basically, the screens in those days were kind of round and they kind of bulged. And at one point I had an idea I couldn't beat this guy so I made it 19 x 18 [laughs], which is a real advantage, especially if you go first. And he didn't notice and he actually got beat a couple times before he figured out that the screen wasn't square.

Q: [laughs] Yeah, I think that's a little harder to get away with that in competitive chess, yeah?

Ken Thompson: It is.

Q: Okay, so –

Ken Thompson: _____.

Q: Yes, yeah. So then how did you end up at Bell Labs?

Ken Thompson: My teachers recommended me. I wasn't looking for a job I was just hanging around. I was having a good time. I kind of almost ran the school as far as computers go. I was into every nook and corner, and they sicked a Bell Labs recruiter on me, who tried to get me down for an interview the entire week he was there. And finally on the last night before he was going home he called me at home and came over -- actually came over to my house to interview me. And disappeared, then I got a letter inviting me for an interview on the East Coast. And I had a lot of old friends that were spread all over the U.S. from the Navy brat kind of days and so I told him that I wasn't interested in a job but I'd take their interview just to get a free ride to the East Coast to visit some of my friends. And they said “anything you want, you know, that's fine”; and went to the East Coast, I interviewed two days out there; one day at Whippany, one day at Murray Hill. Whippany was grim, and at Murray Hill there were actually two departments that looked very good. And when I got home there was a letter saying, if you want as job you can have it, just rank the departments in order. And I thought about it for a week or so and then ranked them, sent them in and I got offers from the top two that I'd ranked, and thought about that for another couple weeks and said, “okay”.

Q: Okay. Now an interesting testament to the persistence of that recruiter, yes. Okay, so anyway, you moved to Bell Labs and when was that?

Ken Thompson: '66.

Q: '66, okay. And what did you start doing when you were there?

Ken Thompson: I was hired to work on Multix.

Q: On Multix, yeah.

Ken Thompson: And you know it was just a big monstrous project with GE, Bell Labs and MIT to build the time sharing system that was going to end all time sharing systems. And it was hugely over-designed and wildly uneconomical. It had very unique hardware. It was a research project on about four fronts, when, you know, a research project should only have one front. Bell Labs quit the project when it decided that it wasn't going to satisfy their needs. And then I was almost like out of a job at that point and I kind of had these old Multix computers lying around before they disappeared. And no real thing to do so I just did what I wanted to do from them on.

Q: Which was?

Ken Thompson: Games and operating systems and – I don't know -- anything that crossed my mind. I did some positional astronomy, some audio, some – I don't know, anything; absolutely anything that I wanted to do.

Q: So we were looking at the Dec PDP-1 in space war. So let's talk about that. I guess that was one of your first applications on mini computers as I recall.

Ken Thompson: Well PDP-1… – The only PDP-1 I ever programmed or had access to was at Stanford when I was at school in Berkley. Stanford and Berkley had co-op arrangements, and at some point we were moving out our 7090 for a 7094 and we had no computer for an extended period. And we got Stanford to loan us their 7090 and 7094 at night. So I'd come down to Stanford at night with jobs to run and stuff like that and take the stuff back in the morning. It was like midnight to 8:00 and in the computer center there was their PDP-1. And I'm sure I played a lot. So that was my only actual contact with the PDP-1 and the real space war. After that Dennis Ritchie and I found a PDP-7 with a 340 display, the same -- roughly the same -- display.

Q: What kind of display is that?

Ken Thompson: It's a vector display. It'll move to and draw lines and it runs a vector list. The display runs an instruction list that says “draw this, draw this, draw this”, and at the bottom it'll stop and then you wait a 60 of a cycle and then start it again at the top. And one of the games we wrote for that was Space War. We pulled up various copies of it and then copied it over and made it run on the PDP-7. And got it running, you know, it was always very popular. And then I did a bunch of other programs. Like there was a three dimensional space where this machine, this PDP-7 was remote job entry station for circuit analysis. And it had a real high speed dataset; it had a 2000 baud dataset on it that was used to send the real computation task to a big central computer. Well there were several of these satellite computers around and we'd have them call each other up and exchange coordinates -- game coordinates. And you could play somebody else on some other computer. And this computer had a big hood that you could put over the screen that was partitioned so that one eye saw half the screen and one eye saw the other half of the screen. And so I wrote a three dimensional space game where you'd travel around in a cube and find the other guy who was on the other machine doing the same thing, and shoot him and kill him. You know, all the games are shoot and kill.

Q: Yeah, yeah.

Ken Thompson: So I did some of that, and then there was a simulation of the solar system that I did.

Q: The space travel game?

Ken Thompson: It was the space travel.

Q: Yeah, okay.

Ken Thompson: That you would take off and go to someplace else. You'd go into orbit around something, dock with something, those kinds of things. It had all the solar system objects that -- it was two dimensional it wasn't three -- that we knew and loved then. There are more now of course. But anyway that was a fun game. And that got us into the PDP-7 and to really running it. And then on the PDP-7 we wrote the first version of UNIX.

Q: Sure, of course, yeah. Now were there – was there any chess on that or was that later?

Ken Thompson: There was an MIT chess program that you could just load stand-alone and then play. And I played maybe one or two games of it. I think it was a variety of the Kotok program.

Q: Okay.

Ken Thompson: But I'm not that sure what its origination was. It was – but it did come from MIT for the PDP-7.

Q: Okay. So then there was a move over to the PDP-11, right, someplace in there?

Ken Thompson: Yes. But we bought a PDP-11 to…. the excuse was text processing but the real thing was to play more, you know, to play.

Q: Right. As I recall this – once upon a time weren't you trying to get a PDP-10 or something like that for the lab?

Ken Thompson: Yes, we were arguing that the Multix machine should be replaced with a PDP-10. And there was such a huge backlash from Multix that it was pretty soundly turned down. It was probably a good idea, the -10 is kind of a trashy machine with 36 bits – the future just left it behind.

Q: Yeah.

Ken Thompson: It –

Q: The world would have been a little different perhaps if that happened. Okay. But in any case you got off onto PDP-11's and so when did the sort of chess emphasis get restarted in your mind there? Or when did you start getting serious writing chess software on that?

Ken Thompson: Oh there was a display on the first PDP-7 and I wrote a set of display programs and then that moved onto the -11, and they were just kind of parlor games to increase interest in it. I wrote a pool program where you'd line up and shoot pool balls and they'd strike and break and bounce over the screen. And I wrote a chess program. The chess program was actually – I took it out and I actually competed it a couple times. I took it the New Jersey Open in about '72 or '73 and won one game and drew a couple games and lost the rest. It was a very, very weak program and it was then distributed on the first versions of the UNIX as CHESS.

Q: CHESS, yes.

Ken Thompson: I did a better job of it in software and took it to the big national chess tournament. There were two sets of chess tournaments going on in those days. One was American, which was run by the ACM and it was routinely called either the ACM Chess Championships, the U.S. Chess Championships (Computer Chess Championships) or the North American Computer Chess Championships, as it got broader, and it ran every year. And it was typically at the ACM conference. And then there was one every three years which was called the World Championship, and that was actually weaker. It had a bigger title but it was weaker because they would try to not keep it dominated by the U.S. teams. And they'd throw out the bottom of the U.S. teams and invite essentially everybody else in the world, and that actually weakened the pot. And then that was run every three years.

Q: So, in that era, could you sort of talk about the competition and the general caliber of play? What sorts of things could the programs do at that point?

Ken Thompson: I think they played well. They were probably – when I started, my program was probably 1200.

Q: So you might want to explain that a little bit.

Ken Thompson: Okay.

Q: Talk about USCF ratings and –

Ken Thompson: The ratings are where one sigma is 200 points. And the classes are every 200 points; they draw a line and they say the class. And if you and I play and I beat you a lot, I steal some of your points, and if you beat me, you steal some of my points. So this is kind of a big point pool, called the ELO Rating System, of all the chess players in the United States, and also other places too, not just the United States. It was originally designed that 1500 is the average tournament chess player. And then if you're good you go to like 1800, that's Class A, 2000 is expert, 2200 is Master, 2400 Senior Master, 2600 is Grand Master, roughly. So that's every 200 points. And then below – 1800 and below -- is like 1816 is A, 1614 is B, 1412 is C, so this was low C; quite a bit below average for chess players. If you play at home with your parents or your friends quite a lot and you go to a real tournament, you'll probably end up around 1200 or 1300. If you play tournaments for a lot and are kind of average -- you don't actually study it, but you play a lot, you go once a week to a tournament -- then you'll be 1500. And then if you show real talent and start winning, then you're in the 2000 to 2200. And then if you're World Class you're finally up around 2600, 2700.

Q: Okay. Yeah I thought that was worth getting into give people context for these – if they're not – if those numbers are not ones that they happen to know.

Ken Thompson: That's an exponential scale too. That's a logarithmic scale. It represents exponential growth, that each one is a sigma better than the other. And a sigma means that you're going to beat the other person 70% of the time.

Q: Okay.

Ken Thompson: And so somebody 1200 is going to be beat by somebody 1400, 70% of the time and then 1400 to 1600 is another 70%.

Q: Seventy percent then, yeah.

Ken Thompson: Yeah. And so when you go out and you play this,– and you're rated and all of this, you kind of think of it – well numbers that you can just kind of increase. And you don't realize the vast

difference between playing skill that the difference between 1200 and 2600 means. You know, how much better Bobby Fisher is than me. And you just work for that next notch, and the next notch, and you don't realize that you're, you know, that you're climbing mole hill before Everest.

Q: Yeah. There certainly is quite a difference. I actually played Bobby Fisher once and it was simultaneous, I didn't last that long. So I know that feeling, yeah. So I guess your take is that the programs of that era were 1200, 1300 kinds of things? Is that –

Ken Thompson: They were probably slightly better than that. Computers were expanding quickly there and the Northwestern program was the perennial champion all during that era. And they broke 1800 towards the end of that era, towards the late 70's, probably middle 70's. They won a Class B tournament, which is limited at 1800. They actually came in first at -- somewhere in the valley here.

Q: So okay, you were running around with basically a pure software solution at this point, running what, PDP-11's?

Ken Thompson: PDP-11, yes.

Q: Okay.

Ken Thompson: Yes.

Q: Okay. So then how did you get going on doing Belle, and getting into hardware as well as software?

Ken Thompson: It just came naturally; that you find out that -- you know, you profile trying to make this thing go better. The algorithm is where you start off with a current board position, you try all the moves and then you turn the board around and let the opponent play; for each of those moves he plays all of his moves, and then for each of those you play. And so this grows exponentially and you plot -- all this stuff is in the literature -- you plot how many ply you can look ahead.

Q: A ply is?

Ken Thompson: A half a move.

Q: Half a move, yeah.

Ken Thompson: Where you turn the board around and play yourself, turn the board around and play yourself for each of those. So there's 20 or so moves on a board so two ply will be 400 and three ply will be, you now, bigger and bigger and bigger. That… You profile what you're doing and you find out all your time is in the mechanics of chess, it's not in the algorithms or the thinking, or the evaluation. It's all mechanics, you know, finding the legal moves. And so the very first thing you think of is an accelerator to make those legal moves a lot faster. And you say, ah ha, one more ply. And so I made a little tiny chess move accelerator that essentially did what software does is – it finds a piece and then walks the – looks for blank squares, you know, wherever the piece would move, radiating out from that piece, and lays down those moves. And then what happens is it presents these moves to the software and the software has to pull it out of the hardware. And then -- it's kind of hard to go into -- but you have to sort those moves before you actually play them. You can't play them in the order, in some random order, that they come out of the hardware, you must sort them. So the hardware generates these moves fairly fast and then you pull them out meticulously and laboriously into software. And then you sort them and by the time you do that the hardware is swamped by the sort and the – So it did speed things up, but not by what you'd normally think of as a hardware solution. It sped it up by two or three, something like that, instead of 100 which you'd expect out of hardware. So that was the first solution. I took some –

Q: So now – sorry. Quick question here, so how many plies deep would your original software solution go and then– did this give you, what, another ply?

Ken Thompson: No, it probably did nothing more. It was probably four ply with software, and then four ply still with hardware. And that competed in exactly one tournament which was the 1977 World Championships, which was in Toronto I believe. And I came in roughly at the center of the field, maybe a little higher. I was a point higher than the center of the field. So it was a wash, you know, but it was fun. These tournaments in those days especially, you'd sit across the table from somebody who knows chess just like you do. And the computers would sit there and play and you're bored. You know, you watch the game but it doesn't take all that much time. And you'd talk about algorithms and you'd talk about ideas on what you're going to do about chess. And then that would actually fuel you. You'd play six or eight of these people over the course of four or five days. Then you go back home, kind of absorb everything you learned, and try it for the for the next year for the next tournament and then come back and it all happens again. So that's how it was sustaining.

Q: Are there any particular people that come to mind in terms of having interesting exchanges with?

Ken Thompson: Oh yeah, essentially all of them. Atkins and Slate were the perennial champions; they were running on the fastest machines. They were running on Control Data 66- and 7600's, and they were sponsored by Control Data. And then Bob Hyatt –

Q: Which program was that?

Ken Thompson: They were CHESS X.Y for 2, 3.0, 4.0, 4.5, 4.6. Hyatt and Blower from Hattiesburg, Mississippi were sponsored by Cray when Cray came out, and they were the champions for a while. Berliner is an amazing man; he won the human -- the individual -- World Correspondence Chess Championship, just an amazing feat. And he sat down and at one of these meetings told me what he went through personally to win this championship; and against whom, and against these… This is the epitome of chess -- the Correspondence Championship -- and he was the best by far.

Q: You might want to explain that, because not everyone knows what that is.

Ken Thompson: Correspondence Chess is where you are given days instead of minutes to make your move. And you do it at home with open books and open chess boards and you move the pieces. And when you make your move you write it on a postcard and you send it to your opponent and your opponent logs it when he gets it and starts his clock. I mean a calendar instead of a clock, and sends it back. And so you get so many moves in so many days, and you carry on simultaneous games. You're playing the whole field at the same time. And then there's the – a lot of the strategy is that where you're playing – if you want to stagger your clock, you play one move and send it, come back. But if there's obvious replies and you want to dump your opponent without the ability to think during the postage time, what you do is you lay out a big – it's like a contract. If you move here I move here, if you move here I move here, if you move here. So you lay down six kind of forced moves or the best moves and dump him right in a position that is coming, and he knows is coming and instead he gets two or three days per move to do it rather than nearly double or triple that time if you count in the postage time, because there's just one postal exchange. So that kind of strategy, playing people's best lines, their own lines, the research involved, it was just a massive toll on his personal life, his professional life it took to do this, he just – it became a full job, more than full job. You know, 20 hour a day kind of job for the duration of this tournament, which is several years.

Q: Wow that is hard, yes. So perhaps before we get into the further hardware work that happened, talk about the structure of the software that you were using at that point in terms of how things like opening book and the middle game and the end game. You know, what were its abilities in these areas?

Ken Thompson: It was really run of the mill. All of the computer programs were roughly the same. There were a couple of outliers, a couple of people who tried to do real planning kind of things. But the other ones were progressively – became strong enough that they never made a bad mistake, and all of the planning type of programs maybe played better and looked better, but they always made some horrible mistake, somewhere along the game. You just have to make 20 moves in a row with no mistakes and if you make one mistake you get your head handed to you. So they can't… They can play chess better probably, but they can't win endgames and win the tournaments better. So they… Darwin

pretty much culled those guys out. All of the other ones were brute force or semi-brute force. And what they would do is start at the opening position, play all the moves for each of those moves, play all the moves for each of those moves, play all the moves. And most of them – at first they played to a fixed depth and then evaluated. And then they started playing to varying depths where they would play to some minimum depth and then extend, based on what they found along those depths. Like if their king was in trouble, they'd move again, or if their – but basically they would never skip to very shallow depths any move. They looked at everything and therefore they'd avoid the really bad mistakes. And then at the bottom they'd evaluate. And an evaluation involved an exchange analysis of what happened and then turn back. And then there's this marvelous algorithm called alpha beta that is based upon the idea that if it's your – if you find a move that beats – if… It's you and I. You make a bad move and I show its bad by taking your queen. You put the queen on free[?], I take it. I don't have to find the best way to take it -- I don't have to find the best move against you -- all I have to do is wipe that out to the point where it looks worse than some other alternative that you have. And so as soon as I refute one of your moves I don't have to find the best refutation, I can quit at that point. And if you sort the moves -- this is what I was getting at earlier -- if you sort the moves, you want to make the best move first so the refutation comes statistically higher on the list and then you throw – you don't look at the rest of them. And so typically if there's a – 35 moves on the board, very typical, one side makes 35 moves, it's on the offense, and it's trying to find the best move, and the one on the defense has to just prove that's a bad move, he makes one move. And then the offense makes 35 moves to try to count that. So it's – the tree looks like 35:1, 35:1, 35:1 and so it grows as the exponential of the depth over two rather than the exponential of the depth. Means you look twice as far with this algorithm as you can without it. And it's just immensely powerful.

Q: So when did that start to get widely used?

Ken Thompson: It was used before I got into chess.

Q: Okay.

Ken Thompson: In history, if you look at the books, there are three or four people who “invented” it. Samuels says he invented it in his checkers program but if you look at the records of the games he played, I doubt that very seriously. Simon and Shaw say they invented it. Knuth has a claim to some portion of it, and writing about it, but not in a real program. And there was another chess program that said they invented it. So you know it's sort of lost in history who actually invented or maybe it was invented more than once. But there's two forms of it. There's a so called deep cut-off version and a shallow cut-off version and I think that the two versions came about separately.

Q: Okay. But it was widespread and used then in the 70's, was that ?

Ken Thompson: Yes it was.

Q: Yeah, okay. So that had gotten in – I think that's certainly one of the things we've been interested in for the chess exhibit was to show the power of such algorithms and why they do make a difference. So that's good to hear that extra commentary on it. Okay, so then after doing mostly software with a little bit of hardware, somewhere in there you really got going on doing Belle with some help from Joe Condon then and I guess some other folks?

Ken Thompson: Right. Joe Condon and I teamed up and we built a small chess machine, which is in a display case in the lobby of Bell Labs now, that is maybe a cubic foot. It had a LSI-11 in it and three cards, which were about a square foot each. One card was a move generator. It would generate the moves and sort them, and play them. It would actually – it had a micro-machine that would actually play. It had another card that would evaluate the board and then it had a third one which was made out of memory chips -- 64K bit memory chips -- that was the transposition table to… like a big cache so that when you found the position you've seen before you can just cough up the answer instead of re-evaluating it, and cut off the search. That machine made its debut in the U.S. Championships in Washington, D.C. in '78 and it won them essentially for the first time, cutting off CHESS X.Y's dynasty.

Q: Okay, and what kind of depth of search?

Ken Thompson: It was searching about eight ply.

Q: Eight ply at that point?

Ken Thompson: Yeah.

Q: Okay.

Ken Thompson: Six to eight.

Q: Okay. All right, so then where did you go from there?

Ken Thompson: We played it again the next year in – I can't remember where -- Minneapolis and we tied for first, because of course they were on monster general purpose computers that were getting faster all the time. And Joe and I then went off and designed another machine after that. So it played those two tournaments. The next one was about the size of a small refrigerator, you know, a hip-high refrigerator.

It was built on the exact same principles but was much more parallel. It had three sections; one section was four boards, which was move generation. One section was evaluation, which was four more boards. There was one more board that was microcode that ran the whole thing. It was conducted by a PDP-11 – an LSI-11 -- and it had commercial memory at this point, 1 megabyte -- big, monster -- of transposition memory that was run by the microcontroller. It was approximately – it was probably -- 100 times faster than the other one. It ran about a 160,000 positions per second. Typical software ran about 6,000 positions per second; that's on a fast machine.

[Start of tape 2]

Ken Thompson: …anyway I took this off to its very first… It went to the chess club a couple of times and I actually became the Westfield Chess Club champ <laugh>. It was clearly playing in the master level area, it-- we had probably four or five masters and one or two senior masters and then a couple of visiting IMs (International Masters) at this chess club. It was a very strong local chess club. Since the early 70s Bill was a member and a frequent sponsor of things. Let me digress for second. There was a hostility about computers and chess, and humans and chess in the same tournament and whether they cheated or not, and it was all big philosophical question, some people were just adamant. And mostly they were just apprehensive and afraid. And so I made very, very slow inroads into the chess club and built on relationships and things like that, sponsored tournaments and always refused prizes, always moved the prizes down if I ever won prizes and then played friendly games and played simultaneous exhibitions with the computer, and then pass out analysis that the computer would print during its game with the people., It built it in so it was a part of the structure of the club and then when new people it was just there and they always played it and everything, everybody played it, while other people would go into a club and demand that they had the right and generate lots of bad feelings. If you have a chess machine like this or a program and you're trying to find out how well you're doing the best way to do that is to take it and play it and get it rated, and the rating is a nice number that says how well you're doing. And if you haven't got a nice friendly place to do that, first off if you're getting rating points from fear factor, which I don't think is fair, and second you're generating more heat than necessary in doing it. So anyway I had this local chess club, Westfield Chess Club, that I cultivated as a resource to play and I went nearly every week for ten years.

Q: And how did its ratings improve in that period?

Ken Thompson: From 1200 to 22-something in a period of about six years. So anyway this monster machine went to the world championship in Linz Austria that was…

Q: Which year was that?

Ken Thompson: That was '80…

Q: 1980?

Ken Thompson: Yeah. And that was -- I don't know -- that was a lot of fun. It was too important to actually take the machine although it was portable enough to do that. So we left it at home and called it up by phone. Most of the computers which were big main frames were also called up by phone. And after four rounds there was a tie – five rounds? -- and so they ran out of money for phones after that <laugh>. And so what they did is they had both the winners -- co-winners were in the United States, one was Belle and one was Chaos was their names -- and had them call each other up in the United States and then once every five or six moves they'd call us and read the moves off to us and we'd pretend like it was real-time <laugh>. So anyway it… And Belle won that, won the play-off. These are Swiss tournaments, very short Swiss tournaments, so actually who wins is remarkably close to the notion of who I thought was the strongest, even though they're pretty random events because of the very short tournaments and "one loss and you're out" kind of thing.

Q: Okay so then where did things go from there, because Belle was certainly a force in chess for a number of years?

Ken Thompson: Yeah, I played it in a bunch of tournaments, a bunch of exhibitions; took it to Moscow. That was that famous…

Q: The famous -- infamous -- trip, yes since you mentioned it, why don't you describe this?

Ken Thompson: This is a hard thing to describe. I got invited by Botvinnik. In case you don't know who Botvinnik was, he was a world champion, the grandfather of chess. He was the… He was actually the span across the war, he was a strong player before the war and a strong player after the war. Well, all the other world champions who had been real strong players were either pre-war or post-war. He spanned it. He was kind of the inventor of the Soviet school of chess and a Soviet dominance of chess. His instituted a set of camps for finding chess talent and building it, and everything. He's just the Soviet chess player of the old time. He invited-- mostly invited my computer because he had a chess program and he was having trouble funding it and it was largely regarded as a-- computer chess was largely regarded as a fraud and he wanted to generate some publicity that this was really was working and that there really was -- for his own personal benefit you know for his own fall out -- and I was allowed to go along with my computer. <laugh> So I show up at Moscow, and I had arranged-- I'd brought the computer out all packaged to be in the hold of this plane that I was going in, and made all the arrangements and everything so that when I arrived it would be there and I could take it out and have -- it wasn't separately shipped -- so I had to go out to Kennedy twice; once to package it the day before, and

then once to go myself. So I landed and I go to the … while they're undoing this plane, and it's not there. I don't now why, it's just not there. Disappeared. And so we make a trace on it and the trace comes back in a few hours and I get it back in the hotel and they say 'it missed the plane, it'll be on the next plane'. So I tell my sponsors, and we go out to the airport and we meet the next plane and wait and it doesn't come off. Don't know what happened and we call and they say 'we don't understand, it should've been on that-- it must be on the next one'.. So we meet the next one, it wasn't on the next one - it's like two days now and there's a bunch of-- there is a massive schedule, I'm there for a week and there's a massive schedule of where I'm supposed to take this machine every day and then there's like little pieces of touristy stuff in between these chess schedules and they keep pushing the schedule back and making it all tourist but you know so I'm touristing. And at some point I get a call from Joe Condon from Bell Labs and he says "I had a—" I hand carried a bunch of spare parts for this machine -- I had a disc and a CPU and you know a whole bunch of spare parts -- and Joe called me up and said 'Don't bring the spare parts back, throw them away'. And I said 'why?' And he says 'because you're probably going to be arrested when you get back in the United States' <laugh> Oh, good. And I said 'what for?' 'For smuggling computers into Russia' and I said 'oh God' <laugh>. So what happened is that an ex-Bell Labs guard was working as a night guard at Kennedy, and happened to notice that a Bell Labs package that said 'computer' on it was cordoned off by the security police -- the Feds from Washington -- and that nobody at Kennedy, none of the local customs people, were allowed to touch it, it was just roped off, it was like another world. And he called up his friends in the guard department at Bell Labs and said 'you know maybe you'll want to know about this' and that rippled around and finally got to Joe Condon who found out that this thing was confiscated out of the hold. What had happened is the day before I left, Ronald Reagan gave his famous "hemorrhage of technology" speech to congress saying 'we're just bleeding our lifeblood out to these commie pinko bastards' and so the customs department they just jump up and they run out to all the airports and what do they see? They see computer, destination Moscow, they just grab it <laugh> and they lock it in a corner and don't tell anybody about it, there's nothing… So I… Next day I call a meeting of all my sponsors and it was the most uncomfortable position I've ever been in my life. It was in the chess club and it was a very long dark drapery ladened deep purple red under lit room, with a table, an oak table probably this thick that was maybe ten yards long. And the head of the Soviet, guy his name is Krogius, he was a grand master, was at the far end, and I was at the far end, I'm sitting all alone and he's got like four minions, clustered him and then there's about 20 spare seats and I'm sitting at the other end at the head of the table. And he says-- I explained to him that it's not going to come, it's confiscated. And then he said something very, very strange. He said 'you know Ayatollah Khomeini?' And I says 'yeah I've heard of him, I'm not a good friend.' <laugh> He says 'well Khomeini has outlawed chess in Iran because it it's against God and do you suppose Reagan did this to outlaw chess in the United States?' And I-- <laugh> I can just see the headlines now. "No I don't think so, I think this is some low level custom guy who probably didn't even know how to play chess." <laugh> And he says 'oh okay" That was the end of this meeting. And then I was a tourist and they wouldn't talk to me, they-- I mean I was pariah, you know but I-- you can't leave except on the… the reservations are like forever, you have to have them months in advance to go in and out of there and you can't-- don't change them so I was there for another half a week as a complete pariah. These guys just didn't want to talk to me. I finally went home and I stopped off in Germany and dumped all my spare parts <laugh> and then flew back to the United States.

After that I tried to get the computer back and you could go out to Kennedy and it was there, but nobody knew who had it, what they wanted, what they were hoping for, what it was all about. And so we had this import/export company that kind of worked with Bell Labs for going to shows and things like that, and they worked full time on it, couldn't find anything about it. At one of the chess club meetings, this Westfield Chess Club, I was telling the story about how it's taken and you can't get it back, you don't know who wants it, you call and there's this telephone number on it, it says 'for information call' -- right on the computer -- and you call this number and they take your name and address and say 'they'll call you back', and they don't and that's it, that's all you can get out of these people, this telephone number. So I relayed this story in the frustrations of trying to get this back, this was about six weeks after it was taken, still sitting at Kennedy. And one of the guys there worked part time in one of these weekly throw-away magazines on Long Island, the ones that you get, you know, a little local community newspaper, he says 'you mind if write this?' And I said 'no, go ahead write it, it's true'. So he wrote it and it came out in one of these magazines and then a clipping service got it and some guy from the Washington Post called me by way of this magazine. And I told him the whole story and he says 'well let me look into it' and I said 'well good luck you know, I've been looking for six weeks, I've been trying very hard for six weeks'. <laugh> So he hung up and 15 minutes later -- 15 minutes later! -- he called me back and he knew who had it, what they wanted, everything about it. He had contacts. Somehow he figured it out from what I told him, and he says that they claim that if customs has it and they claim that it can be used for weapons research. I said, 'you know there's no way it can be used. Maybe if you threw it out of a plane it could kill somebody but you…'. <laugh> He says 'can I quote you?' <laugh> and I said "sure, I said it". So the next day there was a Washington Post article, fairly large, one about this, with this quote about killing somebody by throwing it out of a plane. And the next day it was picked up by nearly every newspaper source. The day after that the Customs called me, they finally called me, -- I'm sure it had nothing to do with these newspaper articles that made them look foolish -- and they said 'what is it worth?' I said 'technically'… They said to me 'technically you're in violation of the Export Act.' The HP terminal that I was taking with this to run the computer was on their list of bad things that you can't take overseas right, and their list is so old -- I mean this was an obsolete terminal you know -- the computers were not on the list, <laugh> all the chips weren't on the list but this lousy HP terminal was on the list. And they said 'so technically you're in violation, and typically what has to be done is you have to pay a fine that's a percentage of the worth of the equipment, so how much is the equipment worth?' And I didn't understand what he was saying but I said 'oh well we-- the budget -- we kept a budget for it -- and it was seventeen thousand dollars to build this thing'. He says 'well you don't seem to understand it was probably used parts now, right? And they probably weren't worth their full weight, and you must realize that we're going to have to fine you a percentage <laugh> of your _____' and I said 'oh well yeah they were used. <laugh>, It's probably only worth a thousand dollars' he says 'okay, well we have to charge you 50 percent, so a five hundred fine'. So Bell Labs sent them five hundred dollars and they sent the computer back <laugh>. That happened in the next day so suddenly it just all broke loose. So anyway that's the Moscow story.

Q: So that's where the origin of the quote is, I had heard that quote of course but I had thought that it was earlier in the sequence, not at the end.

Ken Thompson: No, no it's actually what broke it loose, otherwise it would probably still be at Kennedy <laugh>.

Q: Which year was that?

Ken Thompson: Probably '82 or '83 - '83 - it was June.-While I was there it was VE day -- which is a big, big holiday in Moscow because of the war and it was later than our VE day by one day because good news didn't travel, it was a day late getting there

Q: But you seemed to have lived through that experience with a good story of it anyway. So then after that, let's see, I think as I recall Belle won a number of additional times in those years…

Ken Thompson: Yeah, yeah it sort of dominated; it had a hardware problem that I could never fix, the-- it was wire-wrapped and a batch of wires that we got were nicked What they do is they pre-cut the-- they're in lengths with pre-cut insulation on the ends and their machine was ill-calibrated and it actually nicked the wire when it stripped the wires. So when you put it in a gun and wrap it on the post it would stretch it -- that's how wire wrap works -- and it would break inside the insulation. If you didn't notice it -- if it didn't actually separate the wire so it would cause an error -- you'd have these nicked wires under stress on these posts and they… If they failed you could find them but if they were intermittent which one of them turned out to be, and I never found it, then it would sometimes run and sometimes not run and when you chase it down it would start to work and it would be catastrophic during a game because it would fail in the middle of a game. It stopped working around '85 or 6; and actually the last time it won was at the '86 US championship; it won '80, '81, '82 -- didn't win '83 -- and '86.

Q: So at that stage how many plies was it going and what kind of rating did it have?

Ken Thompson: It was going eight full plies easily, the other one was probably going six and was easily master, probably 2300. I never worked for rating. You can squeeze rating points out by studying your opponents and booking up and taking away-- put a punch of trap kind of stuff in. But I wanted to just play nice solid chess, so it was probably 2300.

Q: And I think at some point I recall you talking about the ratings compared to compute power and you had I think extrapolated for and about what kinds of compute power it would take to get certain improvements and ratings

Ken Thompson: Yeah…

Q: I can't remember when that was, it must have been early '80s or something…

Ken Thompson: Yeah it was early '80s. —It was asymptotic but it was in the high rise part of the asymptote so you can't really tell whether-- where it rolled over. But it would, certainly in that era you were gaining -- for every doubling of horsepower you'd gain a hundred points. I came with some really fakey back of the envelope calculations that said at around 12 or 13 ply you'd be world class. Maybe not-- you'd compete for the world title, you may not win it but you'd be comparable.

Q: Presumably the ratings in some sense get a little less accurate at the very top I would think…

Ken Thompson: Yes, yes.

Q: So it last won in '86 and then at that point did you retire it or were you using it internally for things, what were you doing with it then? There was a period as I recall when you were doing a lot of work on endgames, maybe it's worth talking about the endgames.

Ken Thompson: Okay the endgames. I did it under the name 'Belle' but I didn't really use any hardware, that was all software…

Q: So why don't you talk about that some, because I recall that's been an area where computers have certainly augmented our understanding of chess.

Ken Thompson: It started-- it's a very long story, it goes back-- the first tournament I was in was in… No, the second tournament I was in was in San Diego in about '75 -- '74 or '75. In that tournament David Levy, who's a famous chess personality, was the tournament director. After the games we were in the bar talking and he was saying that computers can't play endgames, even simple endgames and they never will. And he says 'I'm an expert in the rook and pawn against rook endgame, and a computer will never play a rook and pawn against a rook endgame'. So I went to my room that evening and I was calculating the numbers and came to the conclusion that this was doable. That you could solve that game, absolutely solve it by a different mechanism; not by normal computer chess but by a different mechanism. You could just have the answer and look it up, that you could make a table for everything that you're supposed to do. I came back the next day and told him about it and he says 'no, takes too many plies". And I said 'it's ply independent; this is a different method'. He says 'ach, no'; so he just poo- pooed me and I got sort of -- angry is not the right word but I got you know… So I went home and I worked probably for after that ten years on endgames <laugh> from that point encounter one evening after the bar.

So the method is to lay out a chess board that has two kings on it and then look for all the positions that two kings can -- all combinations of positions of the two kings can be there -- and find which ones are mates. And surprise!, you'll find none of them are mates, right. So then you put a king and a queen against the king and you find all the mates in there, so then you have a list of all the mates in one. And then you build -- instead of a move generator you build an unmove generator where you find all the positions that can arise from that. A move is like an unmove except check is a little different, uncapture is a very strange move, sometimes uncastling and stuff like that. But basically it's just a move generator with different rules. So you find all positions that you can force the mates that you found in the last pass in two ply and then you do that again and again and again and again, and when you find no more, everything you didn't find is a draw. So you do it by exhaustive search, one ply at a time back from the terminal positions, from the end positions. It's called retrograde analysis. There's a lot of chess problems like this. So I worked on it and… The computers weren't just big enough to really-- this was a data manipulation problem, it wasn't a computation problem, you just couldn't hold the answers. The biggest disk was like 5 megabytes at the time and this was… Each of these positions of three pieces -- a king, a king and a thing -- were megabytes and then you had another piece on it, it becomes.. it gets bigger and bigger and then when you add a pawn, which implies that it can promote which implies that it can be any number. So anyway I saw a bunch of these three and four piece endgames and some of them turned out to be very, very tricky and actually overturned conventional wisdom.

Q: Say some more about that?

Ken Thompson: Well the hardest one of the four piece endgames is king and queen against king and rook. It's ,the conventional wisdom that the rook cowers like a chicken <laugh> and the queen comes over eventually separates it and takes it-- checks and takes the rook. But the defense it comes up with is not cowering -- it's very active and very pro-active and it kind of keeps the rook a… It keeps the opponent king away from your king by an empty square -- an empty row. No one's every seen this defense before. So I played this at the local chess club against the masters and then I played it at the '77 world championship against some very strong Canadian masters. And then there were a US championship going on in Chicago at the Park Hotel and a friend of mine was there playing. Called him and he arranged and had a bunch of the real strong masters from the US championship-- US Open -- play the computer simultaneously queen against rook … and it held them all to a draw. They couldn't win with a queen. And then he said 'well let's find a real strong grand master'. The US champion at the time was Walter Brown and a friend knew Walter and says "you can't get him to play this, he'll think it's stupid, but if you bet him money, -- he's a bettor -- <laugh> then you've got him"…

Q: …human ego, yes, comes through for the machine.

Ken Thompson: So we bet him fifty bucks he couldn't win queen against rook and so we set up a date. He had friends over and he was going to … He allocated 30 minutes. He was just going to just blow it

away in 30 minutes and you could… He had people over to watch and… About half way through he kicked the people out . <laugh> He was silent, he started going way over on time but we didn't call him on time. Then at some point -- we knew exactly how far we were to mate at every point and we knew he wasn't going to win because it was going to take him more time than we had. We said that the rules were 50 moves, he has to win to the 50th move or less and you know computer chess-- this is human rules, 50 moves.

Q: Sure, that's the 50 moves that required a piece to be taken right?

Ken Thompson: A piece to be taken or a pawn to be moved but there are no pawns here. He lost, and so he wanted a rematch - 50 more bucks … And I said "I'll give you your 50 bucks back I've made my point'. <laugh> He says, 'no, no 50 bucks'. <laugh> We gave him from the current… There are two best positions, or worst positions, whatever you call these things -- maximal positions -- and we gave him one and kept the other one a secret, and from the one we gave him we gave him perfect play to capture. He went home and studied for a week on perfect play and he's a studier, and he studied I think the whole week. The next week -- because this was a Saturday -- the next Saturday we set up again and it was much more serious. We played the second maximal position. And he attacked like… And the computer when it had equal moves it would play them at random. Anyway he got… he said he wanted.. he made some more conditions to: 50 moves he wins. 50 to 55 it's a draw and more than 55-- some number like that I can't remember the exact conditions -- lose. So anyway he actually he did it -- first guy who did it over the board, did it in exactly 50 moves right on the nose, one more move and it would've been a draw by the rules. There's one point right around 18 moves to the win where he… it's like a wall where people just hit it and bounce, they just don't know how to cope with this defense. The computer makes like two or three random moves —selections -- and one of them turned into a repeating move where he could move the whole thing down a notch by doing a set of checks -- it wasn't optimal but it was human and that's how he got through that wall. Anyway it was fun, it was written up in the Minnesota Chess Magazine in an article called Beer in Your Ear. <laugh> -- very, very well written.

Q: So interesting, so it sounds like… What's the current state of endgames and how-- it sounds like we've gotten there with a lot of software.

Ken Thompson: Well they grow exponentially with the number of pieces on a board. So this was four pieces -- five pieces were done after that. Right now the state of the art is six, and six generates gigabytes of data and every more piece after that goes up by a factor of 64. So if you think that something… You've got to wait for the computer industry -- Moore's Law -- to give you a 64 in order to try the next step, and we're between steps from between six and seven now.

Q: As I recall, unless I misheard this, I thought at some stage did you not discover that there were cases I guess that it was impossible to do it in the 50 moves, it was still a win, were there cases like that?

Ken Thompson: Well there are tons and tons of them. The computer… Tthe human rules were 50 moves to-- but they didn't want to outlaw wins that were real wins and there were like two known positions: one was two knights against a pawn -- I think that might have been the only exception -- two knights against the pawn which is the old what's his name, gee I can't think of his name, the guy worked this out, human, where two knights alone can't win because they can't mate a king without stalemating him first, and so that's why the king has the prawn and the pawn is there just so that when you stalemate the king the pawn moves and then you go through.

Q: Yeah, yeah okay.

Ken Thompson: Trotsky, this is the Trotsky position. That was an exception to the rule, where for that game, if you reached that game you had a 100 moves instead. And the idea was that you got to demonstrate that you're going towards a solution but they don't want to take the win away from you if you are in fact wining in the Trotsky position and you just don't have enough time they don't want to take that away from you. This was their theory… is that anyone that was proved to be more than 50 moves they'd add to the list of exceptions and so I started publishing all these lists of exceptions and the list got longer and longer and longer and they were all inhuman I mean they were just amazing solutions…

Q: But you said they were inhuman…

Ken Thompson: <laugh> these were just-- they're just crazy. The hardest one is a queen and pawn against queen. People avoid this like the plague because the queens are just wild on the board and they're interpolated with a million checks before every pawn move. In the '60s endgames there's one that-- this runs about-- depending on where the pawn is, this can run up to about a hundred moves, so they put in a hundred and fifty move exception for a couple of these very long positions of Queen and pawn against queen. They started putting in more and more exceptions and finally they just canned it -- and so they changed these rules about four times adding exceptions -- and then they said 'that's it, 50 moves and if you can't win, tough'. <laugh> So they're back now to: no exceptions 50 moves. The current record is king and two rooks against a king, a rook and a knight and it's like two hundred and fifty moves. I'll give you exact numbers sometimes

Q: I just find that's an interesting interaction between the software finding these things and human-- extra human play

Ken Thompson: Enough of these now are in the literature, are on the web, or there are databases around where you can go find the answers that they become celebrities. When they show up in a famous human game and a human doesn't do them right, suddenly everybody on the web's an expert and they

..."Fisher did this". .. In one of the Fisher-Spassky games that ended up in a draw it was in fact a win based on…

Q: …on these cases, yeah interesting. So let's see… I guess an interesting question looking back on this is to what extent did the structure of your software change from when you had a program that was playing something that's 12 or 1300 to Belle at the end. There's clearly the parts where you replaced software by the sort of equivalent hardware. Were there other particular structural changes that went on?

Ken Thompson: Oh the first ones I didn't really understand alpha beta and it was kind of a mess, and you'd compare things and quit and it was just really not under control. But then as soon as I really boiled down the essence of alpha beta and got that, then they were all the same after that. They had little different extension rules where you search deeper and keep track of how deep you've gone and stuff.

Q: How about evaluation functions, has that changed over time?

Ken Thompson: They just got more and more… When you have a 1200 program there's certain… material dominates and that's it -- hardly anything else. Anything else is: if it plays roughly 1200 points as far as tactically goes, it just doesn't matter whether you understand weakest square complexes and <laugh> -- it doesn't matter. As the computer plays deeper, things that we just learn, like forks and pens and stuff like that, it'll learn, but things that are just outside of its range -- you want to teach it and you put those in the evaluation function. Search does a lot -- it amplifies the evaluation function and so subtle things mean more at bigger depths. There's also another thing -- that there's a randomness involved-- there's a famous experiment I'll-- it's not well-- well-- not many people know it but… Suppose you have a chess program that plays everything except the evaluation function, it's a function "f", okay and you put in some terms like "material is good" you know, whatever you want, all right. Suppose you replace that with a random number -- just rand() -- will that outplay a program that has some weak evaluation real function-- or zero -- will it outplay a program that has zero? If it has a big random number generator as opposed to little random number generator will it outplay it and the answer is yes. Random number generators are good evaluation functions, and the reason is, it amplifies…If there's only one way to play, right, you're going to get a random number evaluation back, right? You're going to play-- this guy must play this, this guy must play this, this guy must play this -- random numbers, so you've got like zero to one that's your random number and you bring it back, that's going to be your function at the end right? But if you had like two or three ways to play what you're going…

End of Tape 102630661_Tape 2128K001.mp3

Ken Thompson: So anyway if you have this… Random number generators in fact are good evaluation of the shape of the tree. And that if you have – there's various technical reasons -- but if you have some real evaluation function, if you throw some little random number generator on, it'll make it better. It will take the shape of the tree into the account of the position and it will give you more flexibility in the tree. It's the way of breaking ties in certain cases. And that shows up in fact that it – a four ply search can have like four coefficients; four patterns, whatever you call them. Four isn't good for deeper things because it'll find different ways to find those four. What you want to do is the deeper you go you want more and more subtleties in your random number generator – in your evaluation function. And so yeah, I guess the question was as you go deeper and deeper and deeper, you must keep up with a valuation function.

Q: So now as I recall I think at some point didn't you have instances of Belle playing itself or something like that? Or instances of the earlier software?

Ken Thompson: That was to get the curves that we talked about earlier about what a ply means. Why you'd have Belle playing at four-ply against Belle at five-ply, and four against six, and six against four and five. And then you pretend like these different Belles, at different plies, are in a tournament and you'd rate the tournament, and you'd get a rating per ply for each of the – and then you'd get this typical asymptotic thing. And the range that we were doing was 100 points per ply.

Q: Okay. That's where that came from, okay. So actually, some of that brings me to one of the more meta-questions in here which is the brute force versus artificial intelligence thing. So, you were clearly on one side of this, perhaps it's worth talking about how you came to believe that was the right thing to do.

Ken Thompson: Because you can't afford to make a mistake. And the deeper you go, the more you can't. Plus, alpha beta is so good that it takes you one unit of time to evaluate one move if you think it's the best move, right? It takes you one more unit of time to – with that as a basis value in the alpha beta, the alpha value -- in the alpha beta search it takes you one more unit of time to prove that all the rest are less at the same depth. So I mean it's so – this algorithm is so dominating that it just doesn't make sense to try to do better than that because once you have a good move and on evaluation it's so much – so easy in computation to use that value to prove a bounds proof on all the other moves that you don't have. And you never have to "you're sorry" when you do that. So it's… the truth is that there is a melding of these two ways now. The current method is to do massive extensions down the main lines and then search the – with some different kind of proof. There's these known-move generators is the big algorithm now. That, if in this exchange, this …, inside the computer you're playing yourself, … you make these moves and the idea is that if at some point you think you're doing okay and so you just pass and let the opponent have a move without you doing anything. It's like testing for the position for a threat. And if he can't do anything, if he can't exceed this alpha value with two moves in a row, then he's got a lousy position.

Q: He's got a lousy –

Ken Thompson: And so you don't have to search that one very deep because with two moves in a row the idea is that with… if you really took your move you can't do better than that. If you can't do it with two moves in a row, he can't do with it –after you move, which is not true in these endzone positions. But you try to avoid the pitfalls. And so this gives you a place where you can severely cut down – what you do is: if you can do a null move and get away with it then you don't search that to a big depth. And if you do a null move and you don't get away with it, either he has a threat then you search back to the full depth.

Q: You have to work for it.

Ken Thompson: And so you can take and throw away most of these positions and search them at, like two ply less. And so computers now, if they applied the horsepower to the – today's horsepower to the algorithms of the 80's they'd be doing eight or nine ply. That now these PC's are doing ten and twelve plies with this null move threat analysis.

Q: So that's interesting. That's another, yet another algorithmic efficiency approach to this.

Ken Thompson: Yes.

Q: Yeah. So, as I recall, once upon a time, chess was viewed as sort of the great AI thing, right? But that seemed to have died off fairly early in the chess playing game, is that fair or is that –?

Ken Thompson: No, it was always touted as the drosophila I guess. That was the word they always used, it's the drosophila of AI. And the – you know _____ and all of those still – they talk about the computer chess being waged over there in brute force. But from afar, they just look at the results, so -- they're getting, they're getting better, they're getting better -- as proof that AI is working without looking at what they're doing. And the… As soon as it became very clear that they're just brute force, plus they're getting very good, and that no AI program ever competed, they slowly dropped the drosophila. Now they're talking about Go.

Q: Go, okay.

Ken Thompson: Where brute force doesn't work on Go.

Q: Have you ever tried the Go turf at all?

Ken Thompson: No, I never got into the game. I guess you have to wire your brain at an early age to – you know I certainly understand the – I mean it's –

Q: The rules are simple.

Ken Thompson: You know it's a trivial game.

Q: Yeah.

Ken Thompson: In some _____ yes. But no I never got into the game. And I tried to read and understand the game but this was in my 20's and I think it's too late then.

Q: So, sort of finishing up here. Any particular moments in these matches where there were surprises, either positive or negative, where Belle surprised you, or the opponent surprised you? Any particular things that stand out?

Ken Thompson: Oh there's some – oh yeah. Almost every match, nearly every game had some montrous surprise. When it was flaky and it would just throw away a piece in the middle of – you know, it had – it was an attacker. I mean I tuned the evaluation function to weight king safety really high, and it would just go for the throat. I loved the way it played. And I did it mostly for entertainment not for results. And so it would be just tearing this guy's king apart just going in there and laying waste and then suddenly drop a piece and announce that it had a mate. And he's looking, where's the mate, you know. And.. No, those are horrible. And then the – all the missed opportunities where you upgrade a program or hardware that's like maybe ten or 100 times faster. And then you go look at the tournament that the previous program played and you say, oh no how did I miss that? And then there's a game – I'll show it to you if there's a chess board around, but it's a 20 move game that is just spectacular that it played. This was in Washington D.C.; it beat – because I was paired – I was a dark horse there and I was paired in the center. And so the first round I was paired in the -- you know, in the Swiss tournaments -- I was paired against the top player. I played with Chess X.Y and I beat him, first round. And the second round I was paired against the second player and beat him. And the third round I beat somebody in the middle and had won the tournament, because they'd all lost and couldn't catch up. And so I'd won the tournament after the fourth round. After the fourth round I'd already won the tournament, I had another game to go. I was more than a point ahead. And this last game I played and it was just – and it was one of the most beautiful games I've ever seen. And you know to watch that on top of coasting to a win and watch it play this marvelous rook sack right out of the opening, where nobody saw it on the board, it was just beautiful. It was a – there was a move in it that was – his move was pawn takes queen, my queen, which was a sack discovering a check.

Q: Okay.

Ken Thompson: And the reply was a double checkmate. You know, block the check and deliver double checkmate.

Q: Deliver check –

Ken Thompson: Double checkmate. It was just – I've never seen anything like that on the board.

Q: Interesting. Well it is – I guess when I was watching all this the thing that struck me was that… the funny things that happened where this was supposed to be artificial intelligence but it was more partnership between human designers of chess programs and the programs. Again, sometimes generating surprises or spending all this analysis of endgames to find the cases.

Ken Thompson: There's very – the endgame stuff -- almost every program now that competes has those endgames just buried in their tables, so that if they run into them they get the answer. Have – they're actually rarely useful, you never run into the endgames in a real game.

Q: Now how is that? Usually you have an endgame sooner or later, correct?

Ken Thompson: No. Not – computer chess games are won in the middle game. They're – rarely – I'm not –

Q: Okay.

Ken Thompson: Not never, but rarely. They're – everybody has them because they think they have to have them. But they're not that useful. Opening book is useful but unfortunately it can be used a weapon against you. If somebody knows your opening book they can book against you and essentially win the game before you –

Q: Get out of book, yeah.

Ken Thompson: - before you're computer starts to play.

Q: Yeah, okay.

Ken Thompson: And the – these commercial programs, the ones that are – sell at Christmas for – it's a big feather to be the best of the commercial programs.

Q: Sure, of course, yes.

Ken Thompson: And so the – and of course the opening books are commercially available. So somebody who comes out a week later has an anti-book to the everything on the market. So anybody who goes and plays these two programs, A beats B 100% of the time because B has A's book. And so I mean there's, you know, there's all of this stuff … Well in real play, except to not fall into opening book traps, opening books are not that useful. If you want to play real high level chess you need to have them. You've got to save the time. They save a lot of time and they avoid a lot of –

Q: Traps.

Ken Thompson: - traps and stuff. And they can direct you into the kind of positions that computers are good at. Computers are still good at wide open tactical positions and they're pretty lousy at block positions. So the combination of saving time, avoiding traps, and steering the game into the types of positions that computers are good at; make them really good for, you know, like Kasparov against Deep Blue, those kinds of positions.

Q: So let's see. Okay, so sort of the current state of the world then, I don't know, what sort of size of opening book do the typical programs have?

Ken Thompson: I don't know. I honestly don't know.

Q: Well what size did you have when you ...

Ken Thompson: Mine was monstrous. I had essentially all of ECL [???].

Q: Okay.

Ken Thompson: And the reason is, is because I would play everybody's favorite opening. Again, if you open E4, because that's your book, your book says E4, you've wasted maybe 80% of your book by not

having, by also having D4 and you know and C4 might have three in there right that don't transpose. And that's exactly the same proportion that alpha beta cuts off. Do you understand? If you're going to make this move you don't have to analyze the other ones.

Q: Sure, sure.

Ken Thompson: And so I had an opening book during that era. It was an order of magnitude larger than the next – but mainly because it was – it had all the openings and I really loved to mix it up and especially in “simuls” I'd play random openings.

Q: Oh okay. And how many different ones would be in the mix of the random openings if you had white? Let's see, there's not that many.

Ken Thompson: I'd probably – I had where I'd tell it to make a random choice, maybe 20 or 30 places.

Q: Okay.

Ken Thompson: Real early in the book.

Q: Okay.

Ken Thompson: So it really – it rarely played the same game twice or near it.

Q: Oh interesting, okay.

Ken Thompson: And then I put the weird stuff in. You know I put, you know, countering gambits and king gambits. And, you know, just weird, weird stuff. Loved it.

Q: Yeah. Now how much storage did that take at that point?

Ken Thompson: Not much at all.

Q: Not too much, yeah.

Ken Thompson: If you store it by tree it takes a byte per move.

Q: Okay.

Ken Thompson: And – but I stored it by position and hashed into the positions and it took maybe 50 bytes per position.

Q: Okay. All right.

Ken Thompson: It would also – on one game it would always convert every position it stored into white, you know. It would rotate – if it was black to move – every position in the book was white to move.

Q: Okay.

Ken Thompson: And so if I wanted – if I had a position that was black to move I'd rotate the board and flip it and flip the colors and then look it up white to move. And there's many instances in the opening books where the same position occurs.

Q: Okay.

Ken Thompson: Black or white to move.

Q: Okay.

Ken Thompson: Where essentially somebody's – white's lost half a move.

Q: Okay, yeah.

Ken Thompson: And if you follow the analysis it's completely different. And the reason is, is because there's a psychological point of view that you're white, you attack.

Q: Okay.

Ken Thompson: If you're black you defend.

Q: Oh I –

Ken Thompson: And so you'd have the exact same position, where the moves… where one side would be attacking and the other side would be defending in the same, you know… the same side would attack and the same side would defend. And once in a game I actually transposed into a black side of a Sicilian that turned into a white side of a English and started attacking. It was beautiful.

Q: Okay. So I guess the question is, are there any other particular things that come to mind about this whole experience that are worth talking about? Or any other particular high points that we haven't caught already sort of as closing words?

Ken Thompson: Oh gosh I don't know. I don't know it was a – I miss those days and I miss the tournaments and I miss the people. I don't know if you know it, [International Master Michael J.] Valvo passed away –

Q: No I haven't heard that.

Ken Thompson: - last year, the end of last year, not long ago at all. They – it was just fun. I don't know, you know, serious chess is pretty serious.

Q: Yes.

Ken Thompson: And you know you can't talk and, you know, –nothing. But these chess tournaments is nothing but talking, you know. And the chess is almost an aside, it's almost like a symposium with some little chess on the side. And that has actually changed; it's much more cut throat now. And they hide their algorithms and the programmers don't talk. They typically have heavy commercial concerns and you saw that changing in the 80's. And now it's really changed. I miss the way it used to be. It was almost like a fraternity or something. It was, you know, people who knew – I don't know how to describe it.

Q: Well I think that'll make an interesting center point then for actually the – some of this for the chess exhibit. Because there's certainly a phase in there when a lot of knowledge comes out and the exact sort of thing you're talking about was happening. So…

Ken Thompson: Well there was – I have a – maybe I shouldn't say this, but there –in a Washington tournament, this guy, a chess programmer, Schwartz, Fred Schwartz, got up from the table to go pee, all right. And he doesn't come back and he doesn't come back, and he doesn't come back, and he doesn't come back, and everyone's worried about it. So we go look in the bathroom and he's in this bathroom in his underwear. And what happened is he got robbed at gunpoint in downtown Washington. And the guy took all of his clothes and his money and then ran out.

Q: With that – so there are probably more surprises than one can possibly describe. But I think that's a – that's an interesting one. Okay, any other last words on this?

Ken Thompson: No.

Q: You're done, okay.

Ken Thompson: Thank you sir.

Q: Thank you sir.

END OF INTERVIEW

