Project Deep Blitz: AMD Opteron PCs Take on IBM's Deep Blue

by Ron Herardian, Global System Services Corporation

In the 1990's world chess champion Gary Kasparov played two historic chess matches against IBM's Deep Blue chess supercomputer. Kasparov won the first match but lost the second by a single point. As a graduate student at Stanford writing a thesis on artificial intelligence at that time I was fascinated by the match. I'd been a computer hobbyist since the 1980s as well as a chess buff. Over the years I've tried practically every commercially available chess program on every platform, including Sargon, Socrates, Chessmaster, and others. I used TRS 80 and Apple II computers, and then IBM PCs running DOS, followed by Macintosh and Windows systems. Since that time I've wanted a chess computer as powerful as Deep Blue: my own world champion level sparring partner. This year my company, Global System Services Corporation (GSS), which is an Advanced Micro Devices Solution Provider based in Mountain View, California, sponsored a research project (Project Deep Blitz) to determine if current 64-bit x86 compatible processors could be used to build a chess computer equivalent to IBM's Deep Blue supercomputer using off-the-shelf parts from AMD and affiliated vendors.

Sizing up Deep Blue

Deep Blue was the brainchild of Feng-Hsiung Hsu who began researching computer chess at Carnegie Mellon University where he received his Ph.D. in Computer Science in 1989. Chess had been viewed as a fundamental challenge in the field of Artificial Intelligence. Hsu's idea was to obtain orders of magnitude improvements in performance by parallelizing the processing of chess positions. In addition to parallel algorithms, implementing Hsu's idea required close integration of software algorithms and hardware circuits. By beating the World Champion in a six game match in 1997 it was finally proved that a brute force search of chess moves is superior to the most sophisticated human conceptual understanding of the game and superior to the ability of the most skilled humans to calculate chess moves. The meaning of these results continues to be debated today because searching through all possible moves up to 8 moves ahead is definitely not how humans play the game. In his book Deep Blue: Building the Computer that Defeated the World Chess Champion, Hsu equates Deep Blue to any other tool devised by humans that can perform a specific task better than a human. Deep Blue did not possess an intellect or consciousness and was literally just a machine. What is more important, according to Hsu, is that the creation of Deep Blue is a human accomplishment.

Figure 1. Kasparov vs. Deep Blue 1997 Match

In the past, building a personal computer equivalent to Deep Blue was not a realistic goal. IBM had spent millions on Deep Blue (the cost of the Deep Blue project from 1985 to 1997 is estimated to have been over $100 million), which was a massively-parallel RS/6000 SP based computer with 32 processors that could evaluate 200 million chess positions per second. Setting aside the multi-million dollar price tag, Deep Blue consisted of a pair of 6 foot 5 inch black towers weighing 1.4 tons. Deep Blue's processors, designated "P2SC", integrated eight older Power2 chips on a single die with a total of 15 million transistors. Thus in terms of processor chips alone Deep Blue contained 480 million transistors; but the Deep Blue team did not stop there. In 1997 Deep Blue also contained 512 Application Specific Integrated Circuits (ASICs), each with 1.3 million transistors for an additional 666 million transistors resulting in a grand total of 1.15 billion transistors.

Since there are an estimated 10^{120} possible positions in a chess game, playing chess well, from a computing standpoint, depends on how many positions can be compared within a time limit, such as a chess tournament time control of 40 moves in 2 hours. Computing speed and brute force calculations are the only way computers can challenge human understanding of the game: machines don't understand the game in terms of human ideas; but they can calculate good moves. Table 1 shows how Deep Blue's calculations evolved between 1985 and 1997.

Table 1. Deep Blue Positions per Second

Year	Positions per Second
1985	50,000
1987	500,000
1988	720,000
1989	2,000,000
1991	7,000,000
1996	100,000,000
1997	200,000,000

If Deep Blue's computing power is compared with its chess rating, assuming Deep Blue was a chess master (rating 2200) in 1985 and equal to the World Champion in 1997 (chess rating 2800+), we can see a dramatic acceleration of the number of chess positions per second necessary to achieve a significant gain in chess rating. On average, the number of positions per second evaluated by Deep Blue increased by a factor of 2.2 times the number of positions per second every year, and an average of 2 years passed for every 100 point chess gain in Deep Blue's rating. As a result, if Gary Kasparov's chess rating had been 2900, rather than 2820, it would have taken IBM at least another 2 years to develop a computer that could beat him. What is interesting, however, is that it would have required calculating nearly 1 billion positions per second (969,289,665) to reach the chess rating of 2900 (see Table 2). This is not surprising given that the number of possible chess positions expands exponentially in a tree of possible positions; and considering that Deep Blue ultimately represents a brute force approach to the problem of making computers play chess. Calculating all possible moves for 10 moves, for example, involves roughly 10 trillion possible positions. It's no wonder that IBM denied Kasparov a rematch because if Deep Blue had lost IBM would have to have invested substantially more money to win the next match.

Table 2. Chess Ratings by Positions per Second

AMD Takes on Deep Blue

With the emergence of 64-bit x86 compatible processors from AMD and multi-core technologies, as well as improvements in commercially available chess software, it might be possible to build a microcomputer—a PC—that could rival Deep Blue. Multiprocessor solutions using AMD x64 and dual core technologies represent an enormous advance over the single-processor, single-core, 32-bit systems available from Intel. Project Deep Blitz involved creating a specialized chess similar to a high-end game PC but using multiple dual-core AMD Opteron processors. Four 64-bit AMD Opteron cores running at 2.2GHz yield an aggregate 8.8GHz of 64-bit x86-compatible microprocessor power. A dual-core Opteron chip contains approximately 233 million transistors. To rival Deep Blue’s 1.15 billion transistors, theoretically, only 5 such chips would be needed (which would equal roughly 1.17 billion transistors); but it’s not that simple. Off-the-shelf PCs, even the most powerful ones in the world, do not contain the custom chess processing ASICs that Deep Blue utilized—Hsu “chess chips”—and Deep Blue used no less than 512 specialized chips. There would be no chess chip or chess coprocessor card in the Deep Blitz machine.

The processor architecture and clock speeds of IBM’s P2SC chips (135 MHz for the two processors allocated to I/O and 120MHz for the 30 processors used for computation) are very different from the AMD Opteron. In terms of industry standard processor benchmarks, the Standard Performance Evaluation Corporation (SPEC) found that Deep Blue’s IBM P2SC chips scored 6.5 on the SPECint95 benchmark and 17.3 on the SPECfp95 benchmark. The modern successor of these benchmarks, SPEC CPU2000, shows that the Opteron 252 (a single core chip) running on a Socket 940 Tyan Thunder K8SE Pro (S2882) motherboard (which is similar to the board in the Deep Blitz machine) scored 1796 (peak) on the SPECint2000 benchmark and 1742 (peak) on the SPECfp2000 benchmark. Obviously, the numbers are not exactly comparable (The SPEC CPU95 benchmark was retired in July 2000), but they do reflect an order of magnitude increase (268 times for integer operations, and 104 times for floating point operations), in step with the 9 years of advances in computing technology since the Deep Blue versus Kasparov chess match in 1997. At the same time, processor clock frequencies, measured in megahertz, have increased an astounding 20 million times. Of course clock frequencies cannot be used to compare different types of processors. Interestingly, the SPEC CPU2000 benchmark includes a chess engine, Crafty. Internally, Crafty is a 64-bit chess engine but it is designed to run on 32-bit machines. Chess “engines” generate and evaluate chess positions using integer, rather than floating point calculations and logical operations such as AND, OR, and XOR (exclusive OR). Crafty runs a reproducible set of chess move searches and therefore can be used to compare the integer, branch prediction and pipe-lining facilities of microprocessors. In any case, how well the Deep Blitz machine would compare to Deep Blue's 32 RISC processors and 512 ASICs would depend primarily on the number of chess positions per second that the Deep Blitz machine could generate and evaluate.

The motherboard used in the Deep Blitz machine was a Tyan Thunder K8SE (S2892) with two 2.2GHz 64-bit, dual-core AMD Opteron (model 275) microprocessors and 4GB of low latency PC 3200 Kingston HyperX DDR 400 ECC SDRAM (KRX3200AK2/2GB). The motherboard supports up to 16GB of memory but not 16GB of high speed DDR 400 memory. 4GB was the maximum amount of low latency DDR 400 memory supported. For the intended application, 4GB of memory (based on 1GB per CPU core) was sufficient since there are limitations to the hash table size supported by the chess software selected (2GB). The board was mounted in a SilverStone Temjin III Nimitz server case (SST-TJ03) with a 520W OCZ PowerStream power supply (OCZ-520ADJ). The Deep Blitz machine, which is bare aluminum on the outside, emerged from the workbench as an imposing silver monolith (see Figure 2).

Figure 2. GSS' Deep Blitz Machine (the Silver Monolith)

The Deep Blitz Machine Versus High-end "Game PCs"

The basic limitation of a dual Opteron system from a game perspective is that it is not the fastest possible machine for single-threaded games that execute on a single processor core. To get the benefits of a four-core system it was necessary to accept that, by comparison, a top-end AMD Athlon 64 FX system (widely preferred by computer game players) would be superior for single-threaded games because Athlon 64 chips are available at higher clock speeds than the dual-core Opteron chips but are otherwise very similar. Dual processor boards and dual-core CPUs do not currently offer any benefit for computer games, but the reverse is true for multi-threaded applications where more processor cores translate to more threads running in parallel thus more total processing.

Compared with the AMD Athlon 64 FX chips, AMD Opteron chips are designed with three "HyperTransport" links for input/output and communication between processors, rather than a single HyperTransport link. Each HyperTransport interface provides up to 24GB per second peak bandwidth per processor to reduce I/O bottlenecks. Intel's alternative "front side bus" architecture provides only 14GB per second peak bandwidth per processor in the best case. In other words, AMD-based systems have far more bandwidth (theoretically, up to 72GB per second per chip) for memory access, multiprocessor configurations, and input/output compared to Intel chips. AMD Opteron chips are designed to produce stable performance at high levels of load involving concurrent memory access and I/O functions. In other words, it's a server chip, not a desktop PC chip. But the Opteron's multiple HyperTransport links boost overall performance for any application.

While the amount of memory supported per CPU by way of a 128-bit processor-to-memory interface is the same on the Athlon 64 FX and Opteron, the Opteron requires registered ECC memory while the Athlon does not. The only performance loss for a single-threaded application running a dual-core Opteron system instead of a top-end Athlon 64 FX is clock speed (about 600MHz differential at the time of this writing). This limitation is partially offset by the additional HyperTransport links of the Opteron chip. For multithreaded applications the Deep Blitz machine would far exceed the performance of any single processor AMD Athlon or comparable Intel based system.

Video, Audio, and Disk Subsystems

Audio and video performance is not significant for chess compared with computer games in general but it's possible to outfit a server motherboard with high-end (PC game quality) video and audio cards. The Deep Blitz machine was given a Sapphire ATI Radeon X1800 Crossfire PCI-E x16 with 512MB GDDR3 memory, a 625MHz Core, 16 pixel pipelines and two 400MHz RAMDACs. Although the X1800 is not as fast as an nVidia SLI solution with two top-end cards its performance is comparable to a single PCI-E nVidia 7800GT and it doesn't require an SLI motherboard, which is not an option for a dual-processor server motherboard. With the ATI X1800 card, the Deep Blitz machine can handle any single threaded computer game with respectable frame rates (roughly 60 fps with the most demanding applications). In terms of moving chess pieces on a screen, the Deep Blitz machine would not suffer from a video performance bottleneck.

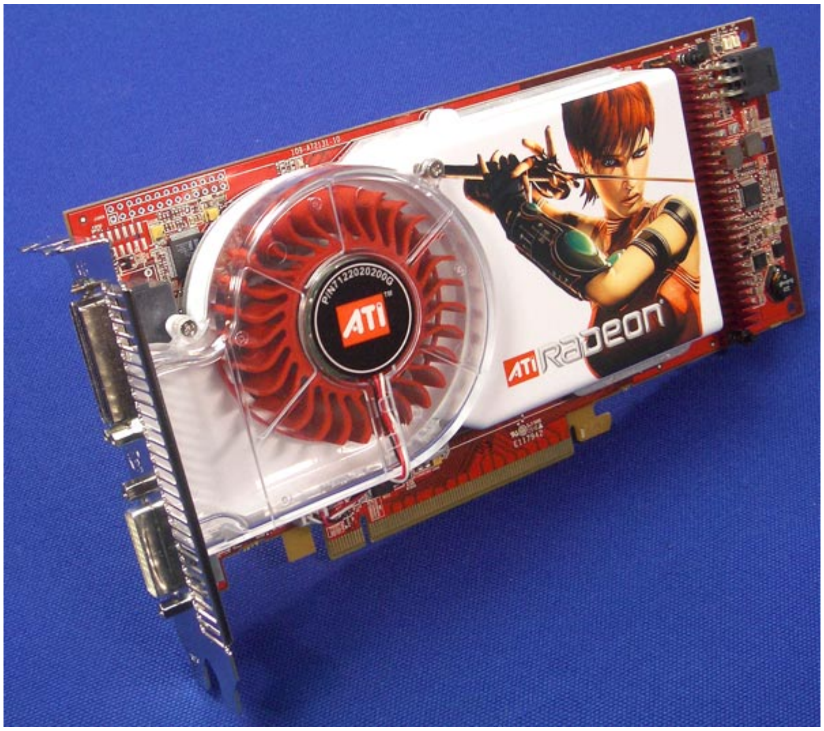

Figure 3. ATI X1800 Video Card

A PCI sound card was used to offload audio processing from the Deep Blitz machine's CPUs. On a dual processor board, or with multiple CPU cores, there may not be a clear benefit to using a separate sound card as would be in a single processor/single core system. In the latter, an on-board audio chipset could sap CPU cycles. Dual-processor server motherboards don't have on-board audio chipsets as most desktop PCs today do. Although the best technical option was probably no audio, the Deep Blitz machine was given a standard off-the-shelf 8-channel Creative Labs Audigy 2 ZS card supporting 24-bit THX and DVD audio.

CD-ROM speed is not an issue for modern games including the hottest chess engines on the market but a CD-ROM is a required component to install software, so the Deep Blitz machine was given a Sony DRU-810A dual-layer 8.5GB DVD drive (DVD+R 16x, DVD+RW 8x, DVD+R DL 8x), which is also a 48x CD-ROM reader. The Sony DRU-810A is not the fastest CD-ROM drive on the market but it does it all. To make up for the read speed deficiency the Deep Blitz machine was also given a Sony CRX320AE CD-RW/DVD combo drive which provides 52x read and 32x (CD-RW) write speed along with a 16x DVD-ROM read speed. If needed, the machine would be highly effective in making disk-to-disk CD or DVD copies using the faster drive to write each media type (using the DRU-810A to write DVDs at 16x speed and the CRX320AE to write CDs at 32x speed).

Disk performance is very important for computer games, which can have multiple gigabytes of data on disk and where a half second lag could mean “game over.” For chess playing programs, disk speed is less of an issue but for a large database of chess games, such as the ChessBase version 9 “Big ’05 database” containing roughly 2.9 million historical games, disk performance is vital for fast indexing and searching. The Deep Blitz machine would also serve as a massive chess database.

The Tyan K8SE (S2892) comes with on-board SATA RAID thanks to the nVidia nForce4 chipset supporting up to 4 SATA disk drives in RAID 0, RAID 1 and RAID 0+1 configurations. Western Digital's 10,000 RPM Raptors are the fastest SATA drives at the time of this writing but they generate more heat and noise than other drives and they tend to be smaller in capacity. A two-drive RAID 0 array of Raptors is practically the performance standard for high-end game PCs but RAID 0 is less reliable. The Deep Blitz machine employs a RAID 0+1 array (striping across drives with mirroring for reliability) using four 7200.7 RPM Seagate Barracuda 160GB drives (ST3160827AS) because, like the Raptors, they also support SATA/150. A four-drive array of Barracudas provides acceptable performance with greater capacity (298GB of disk space using the NTFS 5 file system with 4K clusters) and superior reliability compared with RAID 0, with far less heat and noise. Since the SilverStone Temjin III Nimitz server case has plenty of room (6 internal drive bays with a dedicated oversized fan), a cold spare 160GB Barracuda was racked up for easy recovery of the array in case of an individual drive failure. At the same time, striping over 4 disks partially compensates for the slower rotational speed of the Barracudas as compared with a pair of Raptors. Running, ChessBase version 9, the Deep Blitz machine can search 2.9 million chess games by opening in roughly 2 seconds.

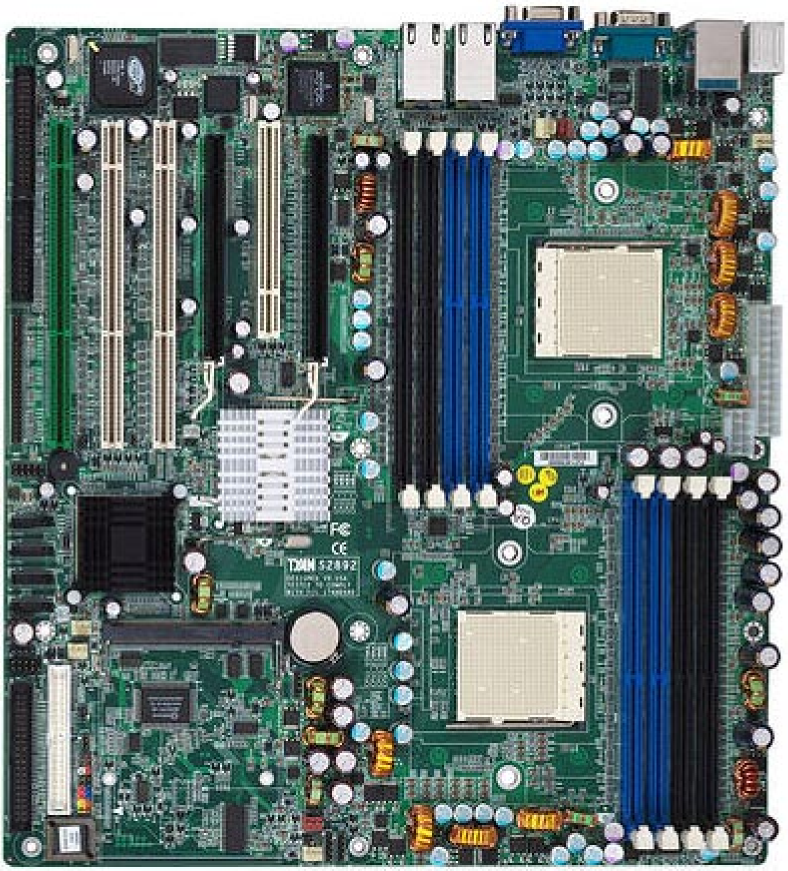

Figure 4. TYAN K8SE Dual-processor, Dual-core AMD 64 Motherboard

The Strongest Microcomputer Chess Software in the World

Project Deep Blitz was limited to commercial chess software thus the chess software of choice, was Deep Fritz Version 8 Multiprocessor Version, available in the US from ChessBase USA (http://www.ChessBaseUSA.com). Deep Fritz is widely regarded as the strongest commercially available microcomputer chess program in the world in human versus computer competition and it is one of only two multithreaded commercial chess packages that can take advantage of multiple CPUs. Shredder, the other multithreaded commercial chess program, written by Stefan Meyer-Kahlen, is the strongest program in computer versus computer competition and it has repeatedly won the title of World Microcomputer Chess Champion. Both programs were installed on the Deep Blitz machine. Deep Shredder 9 Multiprocessor Version was donated to Project Deep Blitz by ChessBase USA. Fritz plays chess, especially “blitz chess” (games that take place in 15 minutes or less) above the master level on a high-end Intel based PCs. The Deep Fritz program, running on a 933MHz dual processor Pentium III computer, played a match against World Chess Champion Vladimir Kramnik in 2002, which ended in a draw. Kramnik defeated Kasparov in a 16 game match in 2000.

Unfortunately, Deep Fritz doesn't run on 64-bit Linux or on Sun’s Solaris x86, both of which can be used on AMD Opteron PCs. Linux, or Solaris x86, because of their more efficient kernels, handling of multiple processors, multithreading capabilities and generally better performance compared with Windows, would have been ideal for Project Deep Blitz. Deep Blue’s code, for example, was written in the C programming language and ran under the AIX, IBM's version of the UNIX operating system. Using a UNIX or Linux based application would be an important step towards running a chess program on a cluster or grid, for example, Sun Microsystems' N1 Grid technology. As a result, Microsoft Windows Server 2003 Enterprise x64 Edition was installed on the Deep Blitz machine. Since Microsoft worked closely with AMD on its x64 technology, it is reasonable to expect significant performance gains despite the fact that both Deep Fritz and Deep Shredder are 32-bit applications.

Fritz Versus the Deep Blue Program

Since there are too many possible positions for computers to evaluate in a reasonable amount of time, one challenge for chess programs is to "intelligently" limit the number of moves evaluated so that only good moves are evaluated in depth. Chess playing programs are based on the idea of a game tree that consists of all possible chess positions starting with the current position. In Figure 4, each player takes 1 turn having 3 possible moves resulting in 12 positions or "nodes" in the game tree. A 'move' in chess consists of 2 plies where each player moves one piece and each ply represents one position. Each position is evaluated by an algorithm that takes different factors into account, such as material and control of space on the chessboard. In a real chess game, each player typically has roughly 30 possible moves and to play at the master level a depth of 6 to 8 moves (12 to 16 plies) must be considered (30^8) resulting in roughly 656 billion possible positions. In a chess tournament, players typically have 3 minutes per move, thus a pure brute force search would require processing 3.6 billion positions per second. Compared to the magnitude of the challenge, Deep Blue's 200 million positions per second seem inadequate.

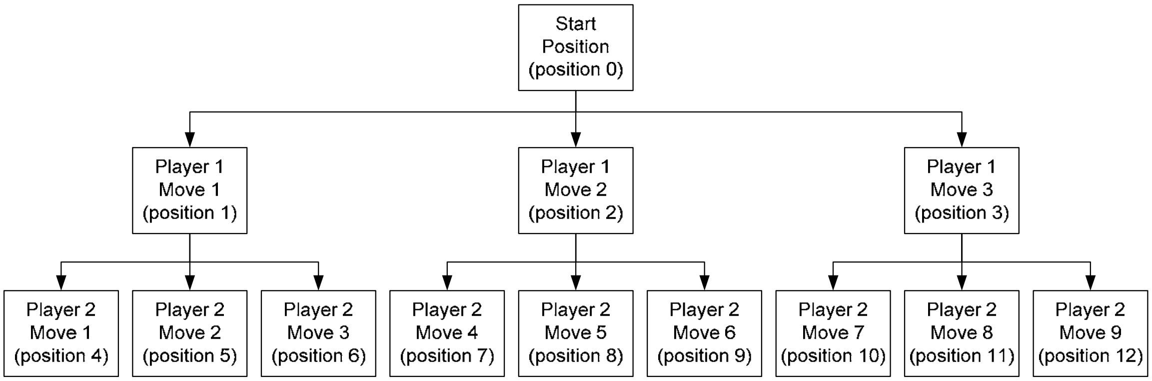

Figure 5. Game Tree

As each position is evaluated moving forward through the game tree, chess programs choose for its color the position that maximizes the value of its position and for the opposing color the position that minimizes the value of its position. This approach is the minimax method. Minimax algorithms make computer programs very good at avoiding tactical mistakes and at refuting unsound attacks but it does not help them with strategy. To cut down the tree of possible positions to a number that could be computed in a reasonable amount of time Deep Blue used an "alpha-beta pruning" algorithm to narrow its search. Alpha-beta pruning improves on minimax by ignoring branches of the game tree that do not contribute further to the outcome, thus reducing the number of positions evaluated while still using a fundamentally brute force approach. According to Frtiz' creators, Frans Morsch and Mathias Feist, the program is built around another pruning technique known as the null-move search where one side moves twice (the other side makes a null-move). This allows Fritz to detect weak moves and to avoid searching them in depth, i.e., a null move pruning algorithm. Move generators, evaluation functions, and data structures in Fritz are designed to maximize the effectiveness of null-move search. Of course the idea underlying Deep Blue was to show that performance could be improved by orders of magnitude using specialized hardware designed for its alpha-beta pruning algorithm, while Deep Fritz and Deep Shredder are designed to run on general purpose x86 compatible computers.

Tuning and Heuristics

Since there are time limits in chess, making the right search is just as important as being able to search deeply enough. This is essentially what humans do when they dismiss “obviously” weak moves and focus on potentially good ones. Historically, chess programs used heuristics to narrow their searches but this approach leads to a collection of highly specific rules that may be incorrect under certain circumstances. Although humans refer to principles of chess there are many exceptions that humans have no difficulty identifying. For example, a knight is generally worth more than a bishop in a closed position, so it would seem to follow that the exchange of bishop for knight would be beneficial in closed positions; but there are many situations where an exchange of pieces would be a mistake. If an exchange of pieces allowed a pawn to begin advancing down the board and eventually become a queen or if the position could later become open, the exchange would be a mistake; but the consequences could be 10 or more moves ahead. Many of the most memorable games in the history of chess involve so-called ‘brilliancies’ that typically begin with utterly counterintuitive moves, thus chess evaluation algorithms cannot be derived directly from the principles of chess in the same way that propositions in geometry can be proved based on axioms.

Today, chess programs use a combination of brute force searches and non-heuristic algorithms that narrow the search space, such as null move or alpha-beta pruning. At the same time, chess programs have evolved into sophisticated expert systems that contain extensive representations of human knowledge, such as a database of chess “openings”, opening traps, and specialized rule sets for chess “endgames.” Kasparov’s comments during the 1997 match to the effect that Deep Blue was making “human” moves, and his suspicion that humans had intervened during the games suggests that human expertise played a central role in Deep Blue’s programming. It seems likely that Kasparov was mistaken in his suspicions but Deep Blue had been “tuned” by a team of human chess experts specifically to play against Kasparov. For example, the Deep Blue team made sure to program specific chess openings and variations based on Kasparov’s record. In fact, Deep Blue won the final game of the 1997 match (and the match) because of an obscure opening trap.

Deep Fritz and Deep Shredder have undergone tuning as a part of their development cycles but their tuning did not focus on a specific opponent. Regardless of the algorithms used, the standard of comparison for the Deep Blitz machine would be the number of chess positions per second evaluated compared with Deep Blue’s 200 million positions per second. There are obviously differences in software but at the end of the day chess programs have to evaluate positions and compare them in order to select a move. For this reason, “chess positions per second” is an excellent cross-platform CPU benchmark. The results are totally dependent on the processor and memory (main memory and cache) performance.

The Shoot Out

So how did the Deep Blitz machine measure up? The Fritz chess software product uses a built-in benchmark, the “FritzMark”, which includes the number of positions per second. While the Deep Blitz machine hits approximately 2000 FritzMarks using a 2GB hash (the maximum supported size)—about 4 times the typical score of a high-end 32-bit Intel Pentium 4 PC—and an impressive 4.55 million chess positions per second, it doesn’t come close to Deep Blue’s 200 million chess positions per second. Additional FritzMark tests were run using the Deep Shredder 9, Crafty 19.17 and Fritz 5.32 chess engines with similar results. The Fritz 5.32 engine hit 5 million positions per second on the FritzMark. The ChessBase company has stated that high end Intel based PCs can process as many as 2 million positions per second so the Deep Blitz machine’s 4-5 million is impressive. However, no comparative benchmarks were run with a dual processor, dual core Intel Xeon based machine. Based on multiple chess engines, the Deep Blitz machine represents approximately 2.275% of the Deep Blue supercomputer’s chess processing power. Of course the Deep Blitz machine is also roughly 2.275% the height and weight of Deep Blue so, pound-for-pound, it’s on equal footing. It’s interesting to consider that 32 IBM RS/6000 SP chips plus 512 ASICs could grind out 200 million positions per second in 1995 while 2 dual-core AMD Opteron chips can handle about 4.5 million positions per second 9 years later. The data suggest that a computer using only 44 dual-core AMD Opteron 64 chips would equal Deep Blue in chess performance.

What is most impressive about the Deep Blitz machine’s results is not how quickly x86 compatible microprocessors are catching up to the big iron, but rather how powerful Deep Blue was 9 years ago. 44 dual core Opteron 275 chips add up to a staggering 10.25 billion transistors. The fact that roughly 10 times the number of transistors is needed to equal Deep Blue’s performance is a testament to the genius of Hsu and the Deep Blue team. What is often overlooked in discussions about Deep Blue is that Hsu’s basic theory, that order of magnitude improvements could be achieved through software/hardware synergy, was clearly proven. In a larger context, however, solving the problem of developing a program to play chess through a brute force approach shows that computers do not understand chess in the way humans do: they cannot apply general principles taking situational variables into account. Although Deep Blue was a triumph for artificial intelligence in one sense, there’s nothing intelligent about brute force calculation.

Checkmating Deep Blue

A computing cluster using only 24 machines identical to the Deep Blitz machine (dual processor/dual core Opteron x64 processors running at 2.2GHz) would surpass Deep Blue’s performance but currently available microcomputer chess software does not run on computing clusters or on a grid. At the moment, the only multi-threaded commercial chess programs are Deep Fritz and Deep Shredder. Deep Fritz 9 Multiprocessor Edition will ship from ChessBase in the second quarter of 2006 but it’s unlikely that chess engine optimizations or incremental performance improvements will make a significant difference in terms of chess positions per second. Either additional processors or specialized chess chips (or both) are necessary to catch up to Deep Blue.

Setting aside hardware issues and brute force benchmarks, chess software has improved in the past few years. Evaluating the games from the 1997 Kasparov versus Deep Blue match using the Deep Blitz machine shows that Deep Fritz and Deep

Shredder consistently find very strong moves for either player within chess tournament time controls and generally do not make the moves considered to have been mistakes in post game analysis of the historic match. In positions such as the 21st move from game 1 of the match, Deep Fritz (with the black pieces) plays the same move as Deep Blue and anticipates the same continuation (Figure 6). The evaluation is to a depth of 16 plies (8 moves) topping out at 4.737 million positions per second with all 4 CPU cores at 100% utilization after roughly 3 minutes of analysis.

Kasparov,Garry 2785 - **Comp Deep Blue 1-0**
A07 New York man vs machine (1) 03.05.1997

1.Nf3 d5 2.g3 Bg4 3.b3 Nd7 4.Bb2 e6 5.Bg2 Ngf6 6.0-0 c6 7.d3 Bd6 8.Nbd2 0-0 9.h3 Bh5 10.e3 h6 11.Qe1 Qa5 12.a3 Bc7 13.Nh4 g5 14.Nhf3 e5 15.e4 Rfe8 16.Nh2 Qb6 17.Qc1 a5 18.Re1 Bd6 19.Ndf1 dxe4 20.dxe4 Bc5 21.Ne3 Rad8 22.Nhf1 g4 23.hxg4 Nxg4 24.f3 Nxe3 25.Nxe3 Be7 26.Kh1 Bg5 27.Re2 a4 28.b4 f5 29.exf5 e4 30.f4 Bxe2 31.fxg5 Ne5 32.g6 Bf3 33.Bc3 Qb5 34.Qf1 Qxf1+ 35.Rxf1 h5 36.Kg1 Kf8 37.Bh3 b5 38.Kf2 Kg7 39.g4 Kh6 40.Rg1 hxg4 41.Bxg4 Bxg4 42.Nxg4+ Nxg4+ 43.Rxg4 Rd5 44.f6 Rd1 45.g7
A well-played game by Kasparov, from both the chess and the psychological point of view.
1-0

```
1. = (0.18): 21...Rad8 22.Nhf1 g4 23.Nd2 gxh3 24.Bxh3 Bg6 25.Ndc4 Qc7 26.Nf5 b5 27.Nd2 Bxf5
2. = (0.21): 21...Red8 22.Kh1 Kh7 23.f4 gxf4 24.gxf4 exf4 25.Nf5 Qc7 26.e5 Nd5 27.Bxd5 cxd5
3. = (0.24): 21...Kh7 22.Qd2 Rad8 23.Nhg4 Bxe3 24.Nxf6+ Nxf6 25.Qxe3 Qxe3 26.Rxe3 a4 27.c4 b6 28.b4 Rd2
4. ± (0.27): 21...Rab8
```

A07: Réti Opening: New York and Capablanca Systems

Figure 6. Deep Fritz 8 Makes the Same Move as Deep Blue

Despite its vastly inferior brute force, the Deep Blitz machine could already be a match for Deep Blue because of improvements in chess software. Deep Fritz is able to evaluate lines of play to a similar depth because it successfully narrows its search only to the strongest lines of play. The data suggest that Deep Blue spent a lot of time evaluating bad moves but overcame this weakness through brute force. In a match between Deep Blue and the Deep Blitz machine running Deep Fritz or Deep Shredder, it seems unclear which machine would win. Obviously, Kasparov did not evaluate 200 million chess positions per second when he defeated Deep Blue in game 1 of the 1997 match, thus the 200 million positions per second number is not a requirement to play chess at the word championship level. It seems likely that Deep Fritz, which is more efficient at filtering out weak moves, is a far more 'intelligent' chess program compared to Deep Blue's software.

Current developers of top chess programs like Deep Fritz and Deep Shredder believe that modern programs, running on PC hardware, are definitely stronger than Deep Blue.

The use of hash tables and other software techniques made possible by the large amounts of memory available today, along with improvements in chess engines mean that far fewer positions per second need to be evaluated for a program to play chess at the grandmaster level. This means that the number of chess positions per second evaluated by the fixed algorithms embedded in Deep Blue's hardware is not directly comparable to the number of chess positions per second evaluated by the best chess programs today. Deep Blue consisted of multiple nodes, each with its own microprocessors, RAM memory and disk storage. As a result, Deep Blue could evaluate very large numbers of chess positions in parallel, but was less able to compare large numbers of positions with each other compared to computers using large hash tables that require several gigabytes of RAM. Playing chess well is a question of evaluating the right positions, not simply more positions (if they are the wrong ones).

Setting aside, the improvements in chess software, the goal of challenging Deep Blue using a PC or a small cluster of PCs, if it has not already been achieved, is definitely within reach. A cluster of machines or a future system similar to the Deep Blitz machine (for example, with quad 64-bit multi-core processors on the motherboard and 8 or more cores per socket running at 4.0GHz or faster) would obviously surpass Deep Blue in terms of brute force. It's only a matter of time; perhaps another 5 years from a hardware perspective. In terms of software, there is a Linux-based Open Source project, ChessBrain, that aims to develop a massively scalable chess program that runs on a cluster of computers (http://www.chessbrain.net). A program like ChessBrain running on a cluster of machines like the Deep Blitz machine would be a logical step toward surpassing Deep Blue. A project of this kind using ChessBrain would cost roughly $500 thousand dollars today–less than 1% of the cost of IBM's Deep Blue project. In another 5 years, the cost will be a fraction of that number. In the mean time, Deep Blue's baby brother, Deep Blitz, plays a "pretty good" game of chess.

Table 3. Deep Blitz Hardware Components

Component	Description	Manufacturer
Motherboard	Tyan Thunder K8SE (S2892)	Tyan
CPU	2 ea. 2.2GHz 64-bit, dual-core AMD Opteron (model 275) microprocessors	AMD
Memory	4GB Kingston HyperX DDR 400 ECC SDRAM (2 ea. KRX3200AK2/2GB)	Kingston
Case	SilverStone Temjin III Nimitz server case (SST-TJ03)	SilverStone
Power Supply	520W OCZ PowerStream power supply (OCZ-520ADJ)	OCZ
Video Card	Sapphire ATI Radeon X1800 Crossfire PCI-E x16 with 512MB GDDR3 memory	Sapphire (ATI)
Audio Card	Creative Labs Audigy 2 ZS	Creative Labs
CD/DVD ROM	Sony DRU-810A dual-layer 8.5GB DVD drive	Sony
RAID Controller	On-board nVidia nForce4 chipset SATA RAID controller	nVidia
HARD Drives	4 ea. Seagate Barracuda 160GB 7200.7 RPM SATA/150 hard drives (ST3160827AS)	Seagate

Table 4. Deep Blitz Software Components

Component	Description	Manufacturer
Chess Engine	Deep Fritz Version 8 Multiprocessor Version	ChessBase
Chess Engine	Deep Shredder 9 Multiprocessor Version	ChessBase
Chess Database	ChessBase Version 9	ChessBase
Operating System	Microsoft Windows Server 2003 Enterprise x64 Edition	Microsoft

Table 5. Chess Titles and Ratings

Titles (FIDE and US)	Rating
World Champion	2800+ FIDE rating (estimated)
Grandmaster (GM)	2500+ FIDE rating
International Master (IM)	2400-2499 FIDE rating
Fédération Internationale des Échecs (FIDE) Master (FM)	2300-2399 FIDE rating
National Senior Master (SM) / US “Senior Master”	2400+ USCF rating
National Master (Master or NM) / US “Master”	2200-2399 USCF rating
National Expert or Candidate Master (E or CM) / US ”Expert”	2000-2199 USCF rating
U.S. Chess federation (USCF) Amateur Classes	1999 and below USCF ratings

About Project Deep Blitz

The Deep Blitz Project was sponsored by Global System Services Corporation (GSS), a Sun Premier Partner and AMD Solution Provider based in Silicon Valley. GSS is a leading provider of software and solutions for e-mail and messaging, directory services, groupware, and wireless. With customers worldwide, GSS provides a full spectrum of services ranging from technology strategy consultation to system architecture and design as well as project management and implementation. For more information visit GSS on the Web at http://www.gssnet.com, send email to info@gssnet.com or call +1 (650) 965-8669.