Acceleration of Genetic Algorithms for Sudoku Solution on Many-core Processors

Yuji Sato Hosei University 3-7-2 Kajino-cho Koganei-shi Tokyo 184-8584 JAPAN +81-42-387-4533

yuji@k.hosei.ac.jp

Naohiro Hasegawa Hosei University 3-7-2 Kajino-cho Koganei-shi Tokyo 184-8584 JAPAN +81-42-387-4533

naohiro.hasegawa.af@stu.hosei.ac.jp

Mikiko Sato TUAT 2-24-16 Naka-cho Koganei-shi Tokyo 184-8588 JAPAN +81-42-388-7139

mikiko@namikilab.tuat.ac.jp

ABSTRACT

In this paper, we use the problem of solving Sudoku puzzles to demonstrate the possibility of achieving practical processing time through the use of many-core processors for parallel processing in the application of genetic computation. To increase accuracy, we propose a genetic operation that takes building-block linkage into account. As a parallel processing model for higher performance, we use a multiple-population coarse-grained GA model to counter initial value dependence under the condition of a limited number of individuals. The genetic manipulation is also accelerated by the parallel processing of threads. In an evaluation using even very difficult problems, we show that execution times of several tens of seconds and several seconds can be obtained by parallel processing with the Intel Corei7 and NVIDIA GTX460, respectively, and that a correct solution rate of 100% can be achieved in either case.

Categories and Subject Descriptors

D.0 [Computer Applications]: General

General Terms

Algorithms, Performance, Design, Experimentation, Verification.

Keywords

Genetic Algorithms, Parallel Processing, Sudoku Puzzles, Graphics Processing Unit.

1. INTRODUCTION

Research on the implementation of genetic computing methods on massively parallel computing systems to attain faster processing [1-5] has been going on since about 1990, but it has not come into widespread use. On the other hand, multi-core processors, graphics processing units (GPU) and other such many-core processors have been coming into wide use in ordinary personal computers in recent years. The features of many-core processors include suitability for small and medium scale parallelization of from several to several hundreds of nodes, and low-cost compared to massively parallel computing systems. This environment has

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'11, July 12-16, 2011, Dublin, Ireland.

Copyright 2011 ACM 978-1-4503-0690-4/11/07...\$10.00.

stimulated research on parallelization of genetic computing on many-core processors [6-10]. Current reports, however, focus on benchmark tests for genetic computing using typical GPUs.

The objective of our research is to use an actual and practical problem to demonstrate that practical processing time is possible through the use of a GPU for parallelization of genetic computation, even for problems for which the use of genetic computing has not been investigated previously because of the processing time problem.

As the first step towards that objective, we take the problem solving Sudoku puzzles [11] and investigate acceleration of the processing with a GPU. The reasons for this approach are listed below.

(1) Sudoku puzzles are popular throughout the world.

(2) Assuming a single core processor, the processing time for genetic algorithms is much higher than for backtracking algorithms [12]. On the other hand, backtracking algorithms pose problems for parallelization, whereas genetic computation is suitable for parallelization. Therefore, increasing the number of GPU cores may make the processing time for genetic algorithms equal to or less than that for backtracking algorithms.

(3) The use of multi-core processors has recently expanded to familiar environments like desktop PCs and laptop computers, and it has become easy to experiment with parallelization programs on multi-core processors through thread programming. GPUs are designed for the processing of computer graphics in games. But research on General-Purpose computation on Graphics Processing Units (GPGPU) has begun, and GPUs can be used to support solving a logical game.

(4) Although the processing time for the backtracking algorithm increases exponentially as the puzzle size increases from 9×9 to 16×16 , the fact that a genetic algorithm is a stochastic search algorithm opens up the possibility of reversing the processing time ratio.

High-speed genetic computing by a GPU requires the design of a parallel-processing program that is dependent on the number of cores and memory capacity in the GPU. This means adjusting the degree of parallelization and the amount of processing allocated to each task according to GPU specifications. In the case of Corei7, however, we focus only on the degree of parallelization afforded by the number of cores in a multi-core processor and propose a method for speeding up Sudoku problem solving by generalpurpose thread programming conforming to POSIX specifications [13].

In this paper, we show the possibility of a large reduction in processing time for genetic computation by parallelization using a many-core processor. In the following section 2, we show an improvement in the accuracy of Sudoku puzzle solution for by using a genetic operation that takes building-block linkage into account [14]. In section 3, we propose an implementation of a parallel genetic algorithm on a GPU. Section 4 describes a comparative evaluation of the solution of a difficult Sudoku puzzle executed on a CPU and on a many-core processor. Section 5 presents a discussion and section 6 concludes this paper.

2. IMPROVED ACCURACY IN SUDOKU SOLUTION USING GENETIC OPERATION THAT TAKES LINKAGE INTO ACCOUNT 2.1 Genetic Operations that Takes Linkage Into Account

A number of studies on application of GA to solving Sudoku have already been made. On the other hand, there seems to be relatively few scientific papers. References [15, 16], for example, defines a one-dimensional chromosome that has a total length of 81 integer numbers and consists of linked sub-chromosomes for each 3x3 sub-block of the puzzle, and applies uniform crossover in which the crossover positions are limited to the links between sub-blocks. References [17, 18] compares the effectiveness for different crossover methods, including one-point crossover that limit crossover points to links between sub-blocks, two-point crossover, crossover in units of row or column, and permutation-oriented crossover operation. In these examples, optimum solutions to simple puzzles are easily found, but the optimum solutions for difficult puzzles in which the starting point has few givens are often not obtainable in realistic time. We believe the reason for the failure of this design is that the main GA operation, crossover, tends to destroy highly fits, schemata (BB) [19]. To avoid that problem, we defined 9×9 two-dimensional arrays as the GA chromosome and proposed a crossover operation [14] that takes building-block linkage into account. An example of this crossover is shown in Fig. 1.



Figure 1. An example of the crossover considering the rows or the columns that constitute the sub-blocks.

In this figure, we assumed that the highest score of each row or column is 9. Therefore, the highest score of each row or column in sub-blocks becomes 27. Child 1 inherits the row information from parent 1, parent 2, and parent 1 in order from top to bottom. Child 2 inherits the column information from parent 1, parent 2, and parent 2 in order from left to right. Mutations are performed for each sub-block. Two numerals within a sub-block that are not given in the starting point are selected randomly and their positions are swapped. We added a simple local search function in which multiple child candidates are generated when mutation occurs, and the candidate that has the highest score is selected as the child. These experiments use tournament selection. The fitness function, Eq. (1), is based on the rule that there can be no more than one of any numeral in a row or column. The score is the number of different elements in a row (g_i) or column (h_i) , and the sum of the row and column scores is the score for the individual.

$$f(x) = \sum_{i=1}^{9} g_i(x) + \sum_{j=1}^{9} h_j(x)$$

$$g_i(x) = |x_i|, h_j(x) = |x_j|$$
(1)

Here, |.| indicates the number of different numerals in a particular row or column. Therefore, maximum score of the fitness function f(x) becomes 162.

2.2 Sudoku Solution Accuracy by GA

For the puzzles used to investigate the effectiveness of the genetic operations proposed in [14], we selected two puzzles from each level of difficulty in the puzzle set from a book [20]: puzzles 1 and 11 from the easy level, 29 and 27 from the intermediate level, and 77 and 106 from the difficult level, for a total of six puzzles. We also used the particularly super difficult Sudoku puzzles introduced in reference [21]. The experimental parameters are population size: 150, number of child candidates/parents: 2, crossover rate: 0. 3, mutation rate: 0. 3, and tournament size: 3.

The relation between the number of givens in the starting point and the number of generations required to reach the optimum solution is shown in Table 1 and Fig. 2. For the three cases in which only mutation is applied (a kind of random search), when mutation and the proposed crossover method are applied (mut+cross), and when the local search improvement measure is applied in addition to mutation and crossover (mut+cross+LS), the tests were run 100 times and the averages of the results were compared. The termination point for the search was 100,000 generations. If a solution was not obtained before 100,000 generations, the result was displayed as 100,000 generations. When the search is terminated at 100,000 generations, the proportion of obtaining an optimum solution for a difficult puzzle was clearly improved by adding the proposed crossover technique to the mutation, and improved even further by adding the local search function. The mean number of generations until a solution is obtained is also reduced.

In Table 2, the number of times the optimum solution was obtained in 100 test runs using the super difficult Sudoku puzzles is shown [21]. Without any trial limitation the method was solved every time. On the other hand, when using a limit of 100,000 trials, our method was solved 99 times, 83 times, and 74 times out of 100 test runs for three super difficult problems, respectively.

On the other hand, Figure 3 shows the relation of the average number of generations and dispersion. For puzzles that have the same number of initial givens, there is a dependence on the locations of the givens, and a large variance is seen in the mean number of generations needed to obtain the optimum solution. Furthermore, for difficult puzzles that provide few initial givens, there were cases in which a solution was not obtained even when the search termination point was set to 100,000 generations. The reason for that result is considered to be that the search scope for the solution to a difficult puzzle is large and there exist many high-scored local solutions that are far from the optimum solution. Another possibility is that there are puzzles for which the search scope is too broad and there is a dependence on the initial values. The processing time was still very poor compared to the backtracking algorithm.

Table 1. The comparison of how effectively GA finds solutions for the Sudoku puzzles with different difficulty ratings.

Difficulty rating	Givene	mut+cross+LS		mut+cross		Swap mutation	
Difficulty rating	divens	Count	Average	Count	Average	Count	Average
Easy (No. 1)	38	100	62	100	105	100	223
Easy (No. 11)	34	100	137	100	247	96	6627
Medium (No. 27)	30	100	910	100	2274	86	26961
Medium (No. 29)	29	100	3193	100	6609	66	42141
Difficult (No. 77)	28	100	9482	100	20658	35	77573
Difficult (No. 106)	24	96	26825	74	56428	9	94314

Table 2. The number of times the optimum solution was obtained in 100 test runs using super difficult problems.

Sudoku puzzle	SD1	SD2	SD3
Count	99	83	74

The numbers represents how many times out of 100 test runs each method reached the optimum.



Figure 2. Relationship between givens and the average of how many GA generations were needed to find the solution.



Figure 3. The difficulty order of tested Sudoku. The minimum and maximum generations needed to solve each Sudoku from 100 test runs as a function of generations needed.

3. ACCELETING GENETIC COMPUTING WITH MANY-CORE ARCHITECTURE 3.1 System Architecture for Sudoku Solution

3.1.1 GTX460 and CUDA Programming

The NVIDIA GeForce GTX460 GPU that we use in this study has 1 GB global memory and seven streaming multi-processors (SMs), and each SM has 48 processors that share a 48-KB high-speed memory. The basic CUDA operations are broadly grouped into the four classes: (1) Reserving GPU memory, (2) Data transfer from CPU to GPU, (3) Parallel execution on the cores of the GPU, and (4) Data transfer from the GPU to the CPU.

CUDA has three units of processing: a thread corresponds to a single process, a block is a number of threads, and a number of blocks of the same size constitute a grid. In CUDA, a thread array of up to three dimensions can be made into a block and a grid can include an array of blocks of up to two dimensions. The unit for the execute instruction from the host is the grid. All of the threads in a grid are executed by the same program, which is called the kernel. The CUDA programming model is a kind of multi-thread model. Each thread is allocated an element in a data array, and the data array serves in the management of the parallel execution of those threads. Threads within the same block share the shared memory inside the SM, so the number of threads within a block is limited to 1024.

3.1.2 System Architecture for Corei7

A conceptual diagram of the homogeneous multi-core processor architecture and system software targeted by this research is shown in Figure 4. Intel, AMD, and other semiconductor companies have recently been marketing quad-core products for a wide range of computers from PCs to servers. In this research, we target a homogeneous multi-core processor that has recently come to be used in PCs and attempt to speed-up Sudoku puzzle solving on a commonly available system. The OS (Ubuntu 10.04) used in our research provides a POSIX thread interface as an application programming interface (API) and can execute a multi-thread program on a multi-core processor.



Figure 4. The system architecture for multi-core processors

3.2 Parallel GA Model and Implementation for Many-core Architecture

3.2.1 Parallel GA model and Implementation for GPU Computation

Figure 5 shows the parallel GA model for GPU computation. Because the grid is the unit of execution for instructions from the host, we conducted experiments with seven blocks in a grid to match the number of SM and with the number of threads in a block equal to three times the number of individuals ($3 \times N$) for parallel processing in units of rows or columns that consist of Sudoku region blocks. We allocate the population pools *PP* and *WP*, and some working pools to the shared memory of each SM. Here, *PP* is the population pool to keep individuals of the current population, and *WP* is a working pool to keep newly generated offspring individuals until they are selected to update *PP* for the next generation.

The procedure of the parallel GA model for GPU computation is as follows.

(1) In the host machine, all individuals are randomly generated and then sent to the global memory of the GPU.

(2) Each SM copies the corresponding individuals from global memory to its shared memory, and the generational process is repeated until the termination criteria are satisfied.

(3) Finally, each SM copies the evolved individuals from its shared memory to global memory.

When applying the genetic algorithm to the solution of Sudoku, the general procedure is to define an 81-bit one-dimensional chromosome that consists of the joined sub-chromosomes of the various puzzle regions and then perform crossover with the crossover points limited to only the joints between regions. Crossover of this type, however, is believed to easily destroy building blocks. As one way to solve that problem, we define a two-dimensional chromosome, taking building-block linkage into account. Crossover is then performed by assigning a score to each row or column, each of which consists of region blocks, comparing the scores for the two parent individuals, and then passing the row or columns that have the highest scores on to the child. Generally, local search functions are effective for constraint satisfaction problems such as Sudoku. Our objective, on the other hand, remains as the solution of Sudoku puzzles in a practical time within the framework of genetic computing. Accordingly, we added only a simple local search function in which multiple child candidates are generated when mutation occurs, and the candidate that has the highest score is selected as the child. This is equivalent to the selection function in (μ , λ)-ES and is an operation in the genetic computing framework.



Figure 5. Parallel GA model for GPU computation

For efficient execution of CUDA programs, we take the following measures.

3.2.1.1 Implementation that takes measures against the initial value dependency problem into account

From Fig. 3, we can see that when the number of individuals is 150, the number of generations until the optimum solution is found depends on the initial values. If we do not consider parallelization, the processing time is considered to be determined by the product of the number of individuals and the number of generations. Accordingly, from the relation of the proportion of correct solutions obtained to the processing time, we set the number of individuals to 150 on the basis of preliminary experiments. We could conjecture that the effect of parallelization using the GPU would be that increasing the number of individuals would not greatly affect the processing time. On the other hand, increasing the parallelism requires that the data on individuals and other data required for genetic computing be stored in the shared memory of the SM rather than in the global memory, but the shared memory capacity is low and may not hold the data for a sufficient number of individuals. Furthermore, the data transfer speed between SM in the GPU is more than 100 times slower than the communication within SM, so an implementation that requires frequent communication between SM is not suitable. Accordingly, we adopt an implementation method in which each SM runs the

same genetic computation program, changing only the initial values of the individual data, etc., and whichever SM finds a solution terminates. In other words, the genetic computing programs running in the SMs using threads are executed in parallel, and the execution of the same program in each SM with different initial values is considered to serve as a measure against initial value dependency. This can be considered to be SIMD-type parallel processing.

3.2.1.2 Sparing use of high-speed shared memory

In previous experiments that involved the evaluation of benchmark tests implemented on multi-core processors [10], tables for random number generation were placed on each core to reduce the time required for random number generation. On the other hand, the shared memory of the GTX 460 is small, with a maximum of 48 KB, so if random number tables are placed in each block, the required number of individuals cannot be defined. Therefore, the CURAND library function is used instead of the random number tables. Random number generation with CURAND is slower than using random number tables on each core, but it is much faster than generating the random numbers on the host and transferring the values to each SM. Furthermore, because 8 KB of the 64 KB read-only constant memory can be accessed at high-speed as a cache, the initial arrangement of the Sudoku puzzle (four bytes (int) \times 81 = 324 bytes) is stored in that cache memory. The data allocated to the shared memory is as follows.

- Area for storing individual data: 1 byte (char) \times 81 \times N \times 2
- Work area for tournament selection: 4 bytes (int) $\times N$
- Work area for crossover: 4 bytes (float) $\times N/2$
- Work area for mutation: 1 byte (char) \times 81 \times N

3.2.1.3 Parallel processing in units of subchromosome

Figure 6 shows an example of the swap mutation within a subblock and the thread assignment. For a Sudoku puzzle that comprises 3×3 region blocks, the genetic manipulation for each region block is performed in parallel, so nine threads are allocated to the processing for one individual. Because of the limited shared memory capacity, however, we assign here three threads to the processing for one individual. The processing for crossover, mutation or other such purpose for each line or column that consists of three region blocks is accelerated by the parallel processing of three threads.



Figure 6. An example of the swap mutation within a sub-block and the thread assignment.

3.2.2 Parallel GA Model and Implementation for Multi-core Processors

It can be seen from Figure 3 that the number of generations needed to find an optimal solution for the same Sudoku problem is highly dispersed. This indicates that a Sudoku solution using genetic computations is dependent on the initial value when a sufficient number of individuals cannot be set due to insufficient memory capacity or other constraints. With this in mind, we generate the same number of threads as cores in the target processor and propose a method that executes genetic operations with each core having a different initial value. We also adopt the value of the core that finds a Sudoku solution first.

The procedure of the parallel GA model at each core processor is as follows.

(1) All individuals are randomly generated.

(2) The generational process is repeated until the termination criteria are satisfied.

(3) The core that finds a Sudoku solution first cancels the operations in the other cores.

4. EVALUATION EXPERIMENTS

4.1 Execution platform

The specifications of the execution platform used in experiments for Intel Corei7 and GTX460 are listed in Table 3 and 4 respectively. The specifications of the GTX 460 GPU used in these experiments are listed in Table 5.

Table 3. Multi-core processor execution environment

CPU	Intel Corei7 920 (2.67GHz, 4 cores)
OS	Ubuntu 10.04
C compiler	gcc 4.4.3 (optimization "-O3")

Table 4. GPU execution environment

CPU	Phenom II X4 945 (3 GHz, 4 cores)
OS	Ubuntu 10.04
C compiler	gcc 4.4.3 (optimization "-O3")
CUDA Toolkit	3.2 RC

Table 5. GTX460 specifications

Board	ELSA GLADIAC GTX 460
#Core	336 (7 SM x 48 Core/SM)
Clock	675MHz
Memory	1 GB (GDDR5 256 bits)
Shared memory/SM	48 KB
#Register/SM	32768
#Thread/SM	1024

4.2 Scalability

The system described here has scalability with respect to the number of cores. Increasing the number of cores in an SM is also considered to improve robustness against the dependence on initial values. For this reason, in the case of Corei7, we varied the number of threads to be executed in parallel from 1 to 8 and surveyed (1) solution rate, (2) average number of generations until the correct solution was obtained, and (3) average execution time. The number of cores in the multi-core processor is 4, but since 2 threads can be executed in parallel in one core by hyper-threading technology, we performed the experiment by executing a maximum of 8 threads in parallel. On the other hand, in the case of GPU, we varied the number of SM ranging from one to seven.

The results are presented in Tables 6, 7, and 8 for Corei7, and in tables 9, 10, and 11 for GTX460. The values shown in the results are the averages for 100 experiments that were conducted with 150 individuals and a cut-off of 100,000 generations.

Table 6. The rate of correct answers, the number of average generations, and the average execution time (Intel Corei7: SD1)

#Threads	Count [%]	Average Gen.	Exec. time
1	94	32,858	22s 19
2	100	15,268	13s 87
4	100	7694	13s 11
8	100	3527	7s 39

Table 7. The rate of correct answers, the number of average generations, and the average execution time (Intel Corei7: SD2)

#Threads	Count [%]	Average Gen.	Exec. time
1	82	42,276	28s 41
2	98	25,580	22s 48
4	100	13,261	21s 47
8	100	5,992	12s 12

Table 8. The rate of correct answers, the number of average generations, and the average execution time (Intel Corei7: SD3)

#Threads	Count [%]	Average Gen.	Exec. time
1	69	60,157	39s 88
2	93	46,999	40s 43
4	100	19,982	30s 79
8	100	8,795	17s 13

Table 9. The rate of correct answers, the number of average generations, and the average execution time (GTX460: SD1)

#SM	Count [%]	Average Gen.	CPU time
1	62	57,687	16s 728
2	80	40,820	11s 845
4	98	19,020	5s 527
7	100	10,014	2s 906

Table 10. The rate of correct answers, the number of average generations, and the average execution time (GTX460: SD2)

#SM	Count [%]	Average Gen.	CPU time
1	50	70,067	20s 199
2	69	58,786	16s 958
4	93	31,254	9s 260
7	97	22,142	6s 391

Table 11. The rate of correct answers, the number of average generations, and the average execution time (GTX460: SD3)

#SM	Count [%]	Average Gen.	CPU time
1	32	82,742	23s 958
2	59	68,050	19s 722
4	77	47,811	13s 879
7	95	30,107	8s 727

4.3 Experiments on Increasing the Number of Individuals

For a large number of individuals, initial values in a Sudoku solution are highly diverse. To investigate the relationship between diversity in initial values and the results of a Sudoku solution, we varied the number of individuals on Intel Corei7 from 150 to 400 and surveyed (1) solution rate, (2) average number of generations until the correct solution was obtained, (3) average execution time, and 4) minimum number of generations.

The result for SD2 is presented in Table 12. The values shown in the results are the averages for 100 experiments. In these experiments, we limited the number of threads to 4 and 8 and set the search termination point to 100,000 generations.

Table 12. The result on increasing the number of individuals (SD2)

#Individuals	Count [%]	Ave. Gen.	Exec. Time	Best Gen.
100	100	8,641	11s 63	644
150	100	5,992	12s 12	243
200	100	7,115	19s 20	229
300	100	9,441	38s 29	123
400	98	15,441	84s 76	86

4.4 Minimum Number of Generations

To estimate the performance in the case that the initial value dependence problem has been solved, we determined the minimum numbers of generations and the execution times required to solve SD1 through SD3 (Table 13).

Table 13. The minimum numbers of generations and the execution times required to solving SD1 through SD3.

Sudoku	Minimum Gen.	Exec. time
SD1	83	25 ms
SD2	158	47 ms
SD3	198	76 ms

5. DISCUSSION

5.1 Scalability and the Dependence on Initial Values

From Table 6 through Table 8, we can see that increasing the number of threads reduces the execution time and increases the correct solution rate. Furthermore, we can see that the reduction rate of the average number of generations and average execution time with respect to execution by one thread decreases as the number of threads increases. In other words, the problem of initial value dependence tends to be eliminated as the number of threads is increased for both the processing time and the correct solution rate.

On the other hand, in the case of GPU, from Table 9 through Table 11, we can see that increasing the number of SM reduces the execution time and increases the correct solution rate. In other words, the problem of initial value dependence tends to be eliminated as the number of SM is increased for both the processing time and the correct solution rate. From Table 13, increasing the number of SM or any other means of solving the initial value dependence problem makes it fully possible to solve super-difficult Sudoku puzzles within one second in a stable manner by parallelization of genetic computation using a GPU.

In other words, we consider that the performance of a multi-core processor is scalable in relation to number of threads and that the performance of a GPU is scalable in relation to the number of SMs. Parallelization using the GTX460 GPU finds solutions faster than that using Corei7 multi-core processor, but we consider this to be due simply to a difference in number of cores. On the other hand, Corei7 exhibits higher solution rates, which we think are due to the fact that random numbers in GTX460 are generated using the CURAND library function.

5.2 Setting the Number of Individuals

Generally, we can consider that the effect of parallelization will become large as the number of individuals increases. From Table 12, in case of the Corei7, increasing the total number of individuals increased the number of individuals that came closer to the correct Sudoku solution but also increased the number that deviated from the correct solution. This is considered to be the reason why the average number of generations until the correct solution was obtained also increased. Increasing the number of individuals also increased the processing time for one generation thereby increasing the average execution time until the correct solution was obtained. The value of best generations also decreased. Accordingly, if the number of cores in the processor can be increased and processing performance by parallelization increased in future multi-core processors, we can expect processing to accelerate to the point where it will be possible to derive correct solutions for even super-difficult Sudoku problems in less than a few seconds.

In the processing-acceleration technique by GPU that we have been developing in parallel with the above technique, programming must take into account upper limits such as task parallelization and memory capacity based on the hardware specifications of the GPU to be used. There is therefore a limit as to how far the number of individuals can be increased to solve the problem of initial value dependence. On the other hand, the technique introduced in this paper, while inferior to the GPU technique in terms of parallelization, enables a parallel program to be executed without limitations in number of threads or memory capacity by virtue of using general-purpose multi-core processors. Looking forward, we believe that accelerating genetic computations for solving Sudoku puzzles by a high-parallelization system should be effective for either GPUs or multi-core processors while solving the problem of initial value dependence by increasing the number of individuals.

On the other hand, in the case of GTX460, the amount of data allocated to the shared memory as described in section 3.2.1.2 is equal to 249*N*, so the maximum number of individuals for which data can be stored in the 48 KB shared memory is 192. Furthermore, considering that 192 is a multiple of the number processors within the SM and also a multiple of the CUDA thread processing unit, the execution time and the correct solution rates for when the number of individuals is set to 192 are presented in Table 14.

Table 14. The execution time and the correct solution rates for when the number of individuals is set to 192.

Sudoku	Count [%]	Average Gen.	CPU time
SD1	100	9072	2s 751
SD2	100	13,481	4s 530
SD3	100	22,799	6s 862

Compared with the case in which the number of individuals is set to 150, the processing time was reduced by approximately 5% while the correct solution rate remained at 100%. For SD2, the correct solution rate increased from 97% to 100% and the processing time decreased by approximately 29%. For SD3, the correct solution rate increased from 90% to 100% and the processing time decreased by approximately 21%. From this data, we can see that setting the number of individuals to an appropriate value for parallel execution of a genetic computing program written in C on a GPU accelerates the processing relative to processing on a CPU by a factor of 25.5 for the SD1 problem, by a factor of 19.1 for SD2, and by a factor of 16.2 for SD3. We can also see that a correct solution rate of 100% can be attained, even for problems in all three of the super-hard categories.

It can therefore be seen that super-difficult Sudoku problems can be solved in realistic times by the parallelization of genetic computing using the Corei7 multi-core processor commonly used in desktop PCs or the inexpensive, commercially available GTX460 GPU.

These experiments also show that the GPU can find solutions faster than the multi-core processor by making use of a higher degree of parallelization. As new GPUs with a higher number of SMs and a higher degree of integration come to be developed, we can expect even faster execution times. At the same time, the GPU suffers from limitations such as the need for programming that must consider upper limits in task parallelization and in the memory allocated to each task due to hardware constraints. In other words, it is more difficult to use a GPU than a multi-core processor which can execute programs in parallel without having to worry about limitations in number of threads or memory capacity. Furthermore, when using a library function such as CURAND for random-number generation due to limitations in shared-memory capacity within a SM, problems with the cycle length of random numbers generated in this way must be taken into account.

Thus, when trying to decide whether to use a GPU or multi-core processor when attempting to accelerate genetic computations by parallelization, due consideration must be given to the target application problem.

6. CONCLUSIONS

We have used the problem of solving Sudoku puzzles to show that parallel processing of genetic algorithms in a many-core environment can solve difficult problems in practical time. Specifically, we implemented parallel genetic computing on the NVIDIA GTX 460, a commercially-available GPU, or the commercially available Corei7 multi-core processor from Intel. Evaluation results showed that execution acceleration factors of from 10 to 25 relative to execution of a C program on a CPU are attained and a correct solution rate of 100% can be achieved, even for super-difficult problems. In short, we showed that the parallelization of genetic computing using a multi-core processor commonly used in desktop PCs or a GPU that can be purchased at low cost can be used to solve problems in realistic times, even in the case of problems for which the application of genetic algorithms has not been studied in the past because of excessive processing times. Future work will use many-core architecture of a higher integration scale to challenge problems of a larger scale and compare the results with the backtracking algorithm.

7. ACKNOWLEDGMENTS

This research is partly supported by the collaborative research program 2010, Information Initiative Center, Hokkaido University.

8. REFERENCES

- V. S. Gordon, V. S. Gordon, D. Whitley, and D. Whitley, "Serial and parallel genetic algorithms as function optimizers," in Proc. of the 5th International Conference on Genetic Algorithms. Morgan Kaufmann, 1993, pp. 177–183.
- [2] H. Mühlenbein, "Parallel genetic algorithms, population genetics and combinatorial optimization," in Proc. of the 3rd International Conference on Genetic Algorithms, 1989, pp. 416–421.
- [3] H. Mühlenbein, "Evolution in time and space the parallel genetic algorithm," in Foundations of Genetic Algorithms. Morgan Kaufmann, 1991, pp. 316–337.
- [4] R. Shonkwiler, "Parallel genetic algorithm," in Proc. of the 5th International Conference on Genetic Algorithms, 1993, pp. 199–205.
- [5] E. Cantu-Paz, Efficient and Accurate Parallel Genetic Algorithms. Kluwer Academic Publishers, 2000.
- [6] J.-H. Byun, K. Datta, A. Ravindran, A. Mukherjee, and B. Joshi, "Performance analysis of coarse-grained parallel genetic algorithms on the multi-core sun UltraSPARC T1," in Southeastcon, 2009. SOUTHEASTCON'09. IEEE, March 2009, pp. 301–306.
- [7] R. Serrano, J. Tapia, O. Montiel, R. Sep'ulveda, and P. Melin, "High performance parallel programming of a GA using

multi-core technology," in Soft Computing for Hybrid Intelligent Systems, 2008, pp. 307–314.

- [8] S. Tsutsui and N. Fujimoto, "Solving quadratic assignment problems by genetic algorithms with GPU computation: a case study," in GECCO '09: Proc. 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference, 2009, pp. 2523–2530.
- [9] Asim Munawar, Mohamed Wahib, Masaharu Munetomo, Kiyoshi Akama. Theoretical and Empirical Analysis of a GPU Based Parallel Bayesian Optimization Algorithm. In Proceedings of the International Conference on Parallel and Distributed Computing, Applications and Technologies. IEEE, December 2009, pp. 457-462.
- [10] M. Sato, Y. Sato, and M. Namiki. Proposal of a Multi-core Processor from the Viewpoint of Evolutionary Computation. In Proceedings of the IEEE Congress on Evolutionary Computation. IEEE, July 2010. pp. 3868–3875.
- [11] Wikipedia. Sudoku. Available via WWW: http://en.wikipedia.org/wiki/Sudoku (cited 8.3.2010).
- [12] Wikipedia. Backtracking. Available via WWW: http://en.wikipedia.org/wiki/Backtracking (cited 1.11.2011)
- [13] IEEE, ISO/IEC 9945-1 ANSI/IEEE Std 1003.1, 1996.
- [14] Y. Sato and H. Inoue. Solving Sudoku with Genetic Operations that Preserve Building Blocks. In Proceedings of the IEEE Conference on Computational Intelligence in Game. IEEE, August 2010, pp. 23–29.
- [15] T. Mantere and J. Koljonen, "Solving and Ranking Sudoku Puzzles with Genetic Algorithms," in Proceedings of the 12th Finnish Artificial Conference STeP 2006, October, 2006, pp. 86-92.
- [16] T. Mantere and J. Koljonen, "Solving, Rating and Generating Sudoku Puzzles with GA," in Proceedings of the IEEE Congress on Evolutionary Computation, September, 2007, pp. 1382-1389.
- [17] A. Moraglio, J. Togelius, and S. Lucas, "Product Geometric Crossover for the Sudoku Puzzle," in Proceedings of the IEEE Congress on Evolutionary Computation, July, 2006, pp. 470-476.
- [18] Edgar Galv'an-L'opez and Michael O'Neill, "On the Effects of Locality in a Permutation Problem: The Sudoku Puzzle," In Proceedings of the IEEE Symposium on Computational Intelligence in Game. IEEE, September 2009, pp. 80–87.
- [19] D.E. Goldberg, and K. Sastry, "A practical schema theorem for genetic algorithm design and tuning," In Proceedings of the 2001 Genetic and Evolutionary Computation Conference, Morgan Kaufmann Publishers, pp. 328-335, 2001.
- [20] Number Place Plaza (eds.), "Number Place Best Selection 110," vol. 15, ISBN-13: 978-4774752112, Cosmic mook, December, 2008.
- [21] Super difficult Sudoku's. Available via WWW: http://lipas.uwasa.fi/~timan/sudoku/EA_ht_2008.pdf#search ='CT20A6300%20Alternative%20Project%20work%202008' (cited 8.3.2010).