PARALLEL EVOLUTIONARY ALGORITHMS FOR MULTIOBJECTIVE PLACEMENT PROBLEM

ABSTRACT

Non-deterministic iterative heuristics such as Tabu Search (TS), Simulated Evolution (SimE), Simulated Annealing (SA), and Genetic Algorithms (GA) are being widely adopted to solve a range of hard optimization problems [1]. This interest is attributed to their generality, ease of implementation, and their ability to deliver high quality results. However, depending on the size of the problem, such heuristics may have very large runtime requirements. One practical approach to speeding up their execution is parallelization. This is all the more true for multi-objective cell placement, where the need to optimize conflicting objectives (interconnect wire-length, power dissipation, and timing performance) adds another level of difficulty [2]. In this paper, we present parallelization of TS and SimE, for multiobjective VLSI standard cell placement problem. Profile analysis of sequential code is conducted to assist in selecting and engineering earlier proposed strategies. Fuzzy logic is used to integrate the costs of multi-objectives. The TS implementation is a based on a synchronous candidate list partitioning model. The implementation of parallel SimE is based on random distribution of rows to processors [3, 4]. For comparison purposes, a parallel genetic algorithm (GA) based on the island model [5, 6], and a parallel SA based on the Asynchronous Multiple-Markov Chain Model [7] were implemented. Results of experiments on ISCAS-85/89 benchmark circuits are presented, with solution quality and speedup used as metrics for the comparative/relative evaluation of the heuristics.

1. INTRODUCTION

General Non-deterministic Iterative Heuristics such as Tabu Search, Simulated Evolution, Simulated Annealing, and Genetic Algorithms are getting more widely adopted to obtain near optimal solutions to numerous hard problems [1]. For small problems, stochastic heuristics have reasonable runtime requirements. For instance, a VLSI cell placement problem with few hundred modules, it is possible to find very good solutions in reasonable time. However, most practical circuits are very large and require several hours of computer time to solve [8]. One way to adapt iterative techniques to solve large problems and traverse larger search spaces in reasonable time is to resort to parallelization [4, 9].

In this paper we are addressing the problem of parallelizing non-deterministic iterative heuristics to solve the multiobjective VLSI standard cell placement problem by using a cluster of low cost PCs. The eventual goal being to achieve either lower run-times for same quality solutions, or higher quality solutions in the same amount of time as the serial approach.

This however is easier said than done: an effective parallelization strategy must consider issues such as proper partitioning of the problem to facilitate uniform distribution of computationally intensive tasks,

while also enabling a more thorough traversal of the complex search space. It is worth mentioning, that in certain design problem, sometimes even the smallest savings in time are worth several times the cost, or conversely, the importance of achieving a specific higher quality becomes essential, irrespective of time taken or resources used. In the subsequent paragraphs we present a brief review of earlier efforts towards the parallelization of TS and SimE.

Parallelization of TS: A number of parallelization techniques for TS have been reported in literature [10]. In the most straightforward and widely adopted approach k tabu search processes are spawned and run concurrently on k processors where each processor carries out independent search [10, 11]. Malek suggested linking independent searches where each slave runs a copy of a serial TS but with different parameter settings [12]. After a specified time the slave processes are halted and the main process selects the best solution found and broadcasts to all slave processes to start the entire search again from this new solution.

Another approach to parallelize search within an iteration is when each process is given a task of exploring a subset of the neighborhood of current solution. Two approaches are followed: *synchronous* and *asynchronous*. In the synchronous approach the various processes are always working with the same solution, but exploring different partitions of the current local neighborhood. A *master* process orchestrates the activities of the *slave* processes [11]. In the asynchronous approach, all processes are peer and usually are not all working with the same current solution [10]. Both approaches require that the set of possible moves be partitioned among the available processors so that each processor will be exploring a distinct sub-region of the current solutions neighborhood.

Suggestions to increase efficiency of TS by parallelizing also include partitioning the search space, which is difficult, or partitioning the problem into smaller sub-problems, determining the best moves for each sub-problem, and then performing a compound move [10].

Attempts to solve several practical NP-hard problems have been reported in literature. For example, parallel implementations include those for the vehicle routing problem [11], and quadratic assignment problem (QAP) [13]. For the QAP, the parallelization strategy has been implemented on the Connection Machine CM-2, a massively parallel SIMD machine. A reduction in the runtime per iteration was achieved when compared to other sequential and parallel implementations [14, 13].

Parallelization of SimE: Unlike other iterative heuristics the parallelization of SimE has not been the subject of much research. Three ways of speeding up the SimE algorithm have been suggested in literature [3, 4]: (1) a distributed memory MIMD parallel algorithm, (2) a shared memory MIMD parallel algorithm that is based on an extension of basic SimE using the concepts of windowing and hierarchy, and (3) hardware acceleration which consists of implementing time consuming parts in hardware (namely, goodness computation).

A parallelization strategy for VLSI cell placement for a single objective (wirelength) was attempted on a network of workstations [3], where each station is assigned a number of rows of the placement problem, in a pre-determined order. The stations executes one iteration of the SimE algorithm on the cells of the rows assigned to it. In each iteration, the rows are redistributed among the processors in a predetermined order [3].

In the following section, we briefly mention our NP-hard, multiobjective, VLSI cell placement problem. Problem formulation and models for estimating the costs for the various objectives to be optimized are referred. In Section 3, details of the parallelized algorithms is presented. Experimental setup, re-

sults obtained on ISCAS benchmark Circuits and other observations are given in Section 4, followed by Conclusion in Section 5.

2. PLACEMENT OPTIMIZATION PROBLEM AND COST FUNCTIONS

Due to its complexity, VLSI design process is divided into several intermediate levels of abstraction. Placement is a phase in physical level, and is a process of arranging circuit cells (components) on a layout surface with a goal of optimizing certain design objectives while meeting the given constraints.

In this work, we deal with standard cell placement, where all the cells are constrained to have the same height, while the width of the cell is variable and depends upon its complexity [2]. Cells are arranged in rows with routing channels between the rows. Due to varying width of cells, row widths may be unequal depending on the type and number of cells placed in a row. An approximation would be to treat cells as points, but in order to have a more accurate estimate of wirelength, widths of cells are taken into account. Heights of routing channels are estimated using the vertical constraint graphs constructed during the channel routing phase. With this information, a fairly accurate estimate of power dissipation, delay and total wirelength can be obtained [2]. The aggregate cost of the solution is computed using fuzzy rules used in [8].

The multiobjective placement problem under consideration is NP-hard, and conventional constructive techniques have often proved inadequate. Even the wire-length only optimization problem is known to be NP-Complete [1].

3. PARALLELIZATION & IMPLEMENTATION OF ITERATIVE HEURISTICS

In this section we present the parallelization strategies and the implementation details of the heuristics. We discuss the how the computationally intensive tasks are handled in each heuristic to distribute the workload amongst the available processors.

3.1. Parallel Tabu Search

A generic intuitive strategy for parallelization is to partition the data into small subsets that are distributed among the processors. Each processor is responsible for a data subset and implements a sequential version of the concerned function (or the heuristic) over this data subset.

The sequential implementation of TS was analysed using profiling tools (GNU profiler) to obtain insight into determining the time consuming operations of the code and the usage of resources. For the circuits experimented on, between 60-80% of time was spent on computation of cost of the objectives and their fuzzification. Furthermore, experiments with parameters revealed that for our hard optimization problem with conflicting multiobjectives, large sizes of candidate list (upto 120) were required to obtain high quality solutions. Since the computation of cost for all moves in the candidate list was the most time consuming operation, (in each iteration) the algorithm was designed to partition this workload.

Therefore, the parallel Tabu Search strategy adopted in this work employs dividing the operations within a TS iteration. According to taxonomy given by Crainic et. al [15], our approach can be classified as a synchronous master-slave (one master and remaining slaves), 1-control (each process is responsible for its search), Rigid Synchronous (RS) (all processes are forced to establish communication and exchange information at specific points) and Single Point Single Strategy (SPSS) (all the processes start with the same initial solution and follow the same strategy).

```
Algorithm  MasterProcess;
Begin
    (* S_0 is the initial solution. *)
    (* BestS is the best solution. *)
    (* PCL is the Partial Candidate List. *)
    (* p is the number of slave processors. *)
    (* OBM is the Overall Best Move. *)
    Generate S_0 and p number of PCLs;
        Send S_0 and a PCL to each slave process;
        While   iteration-count < max-iterations
            Receive best move and cost from each slave;
                Find OBM subject to tabu restrictions;
                Generate P number of PCLs;
            Send OBM and a PCL to each slave process;
            Update BestS; /*by applying OBM on BestS*/;
        EndWhile
    Return (BestS)
End. /*MasterProcess*/
```

Fig. 1. The master process in parallel TS.

```
Algorithm  SlaveProcess;
Begin
    Receive S_0 and a PCL from the master process;
    CurS = S_0; (* Current Solution *)
    While   iteration-count < max-iterations
        Try each move in PCL and compute cost;
        Send the best move and its cost to the master process;
        Receive OBM and a PCL from the master process;
        Update CurS /* by applying OBM on CurS */;
    EndWhile
End. /*SlaveProcess*/
```

Fig. 2. The slave process in parallel TS.

In this implementation, there is an initialization step during which, the master process (Figure 1 generates and sends an initial solution and a disjoint (non-overlapping) partial candidate list (PCL) to each slave process. A move in a PCL assigned to a slave in a particular iteration does not appear in PCLs assigned to other slaves. Each slave process searches its local neighborhood by trying each move in the partial candidate list on the initial solution and computes gains due to them. Then it sends the best move and its corresponding cost (or gain) to the master process. The master process selects the overall best move (OBM) among the moves it received from slave processes subject to tabu restrictions.

Then in each subsequent iteration, the master process sends the overall best move and a new partial candidate list to each slave process. Each slave process now starts by performing the received overall best move so that all the slave processes start their iteration from the same solution. Each slave process searches its local neighborhood and sends the best move and its cost to the master process. The pseudo code of the slave process is given in Figure 2.

3.2. Parallel Simulated Evolution

The parallelization of the SimE algorithm is carried out by partitioning the workload among available processors. The partitioning is done according to rows. The workload for each slave in the cell placement problem is the computation of SimE operations of Evaluation, Selection, and Allocation on it's assigned rows [3].

Algorithm Slave_Process($CurS, \Phi_s$)
Notation
 (* B is the bias value. *)
 (* $CurS$ is the current solution. *)
 (* Φ_s are the rows assigned to slave s. *)
 (* m_i is module i in Φ_s. *)
 (* g_i is the goodness of m_i. *)
Begin
    Receive Placement_ And_ Indices
 Evaluation:
    ForEach $m_i \in \Phi_s$ evaluate g_i;
Selection:
    ForEach $m_i \in \Phi_s$ **DO**
      Begin
       If $Random > Min(g_i + B, 1)$
       Then
        Begin
         $S = S \cup m_i$; Remove m_i from Φ_s
        End
      End
Sort the elements of S
Allocation:
    ForEach $m_i \in S$ **Do**
      Begin
       Allocate(m_i, Φ_s)
 (* Allocate m_i in local partial solution rows Φ_s. *)
      End
Send_Partial_Placement_Rows
End. (*Slave_Process*)

Fig. 3. Structure of the Distributed Simulated Evolution Algorithm.

The row allocation pattern that was proposed in [3] is made up of two alternating sets of rows. In the even iterations, each slave gets a slice of $\lceil \frac{K}{m} \rceil$ rows, (where m is the number of slaves, and K is the total number of rows in the placement) while in the odd iterations the j^{th} slave gets the set of rows j, $j + m$, $j + 2m$, and so on. It has been shown that with the above fixed pattern of assigning rows to slaves in alternate steps, each cell can move to any position on the grid in at most two steps [3]. The consequence of row partitioning however is that the each processors has only a partial view of the placement. This hinders free cell movement, making it more difficult for cells to reach their optimal locations. Results from implementing this strategy on our multiobjective optimization problem revealed that even when given a large amount of time, the best solution obtained was poorer than one achieved by the serial implementation.

Though the lack of a global placement view will always exist in case of a parallel algorithm, the effects of restrictive cell movement can be alleviated by using a better row allocation pattern. The use of a pattern that facilitates a variety of combination among the rows sounds intuitively better. In this work we propose an enhanced random row allocation scheme. The pseudo code of the parallel simulated evolution is illustrated in Figures 3 and 4. As can be seen, one of the processors (the master) is in-charge of running SimE on a particular partition as well as performing the following tasks periodically at the end of each iteration: (1) receive the partial placements from all other processors and combine them into a new solution and evaluate its fitness, (2) partition the new solution to obtain a new row allocation, and finally, (3) distribute the resulting sub-populations among the processors. The number of rows randomly

```
  Algorithm Parallel_Simulated_Evolution
     Read_User_Input_Parameters
     Read_Input_Files
Begin
     Construct_ Initial_ Placement
     Repeat
     Generate Random_ Row-Indices
         ParFor
            Slave_ Process(CurS, Φ_s)
     (* Broadcast Cur Placement And Row-Indices. *)
         EndParFor
         ParFor
            Receive_Partial_Placement_Rows
         EndParFor
     Construct_Complete_Solution
     Calculate_Cost
     Until (Stopping Criteria is Satisfied)
  Return Best_Solution.
End. (*Parallel_Simulated_Evolution*)
```

Fig. 4. Outline of Overall Parallel Algorithm.

assigned depends on the size of the placement and the number of processors. This is repeated for all iterations until the termination condition is met.

4. EXPERIMENTS & RESULTS

4.1. Experimental Setup

The experimental setup consists of the a homogeneous cluster of 8 machines, x86 architecture, Pentium-4 of 2 GHz clock speed, and 256 MB of memory. These machines are connected by 100Mbit/s Ethernet switch. The operating system used in RedHat Linux 7.3 (kernel 2.4.7-10). The paradigm used for parallelization is MPI (Message Passing Interface). Specifically, MPICH (a portable implementation of MPI standard 1.1) is used in our implementation. In terms of GFlops measure, the maximum performance of the cluster, with NAS Parallel Benchmarks was found to be 1.6 GFlops, (using NAS's LU, Class A, for 8 processors). Using this same benchmark for a single processor, the individual performance of one machine was found out to be 0.3 GFlops. The maximum bandwidth that was achieved using PMB was 91.12 Mbits/sec, with an average latency of 68.69 μsec per message.

4.2. Results & Discussion

In this section we present the performance of the implemented heuristics. ISCAS-85/89 circuits are used as performance benchmarks for evaluating the proposed parallel TS placement technique. These circuits are of various sizes in terms of number of cells and paths, and thus offer a variety of test cases.

For comparison purposes, a parallel genetic algorithm (GA) which is a derivative of a standard distributed GA and follows the island model, with independently evolving sub-populations and periodic exchanges of solutions through migration [5, 6], was implemented. A pseudo-diversity approach is taken, wherein similar solutions are not permitted in the population at any time. This diversity serves to widen the search, while limiting the possibility of premature convergence in local minima solution space. The

initial population is constructed at the master process and distributed among N slave processes which start running serial GA on their allocated population for a predefined number of iterations called the Migration Frequency (MF). Then each slave process sends MR (Migration Rate) number of its best solutions to the master process, which selects MR overall best solutions and broadcasts them to all slave processes. Each slave process absorbs the incoming best solutions into its population (if they are not already found) by replacing the weakest solutions. Each slave process then continues with the serial GA for another MF iterations. Standard PMX crossover is used to generate offsprings [1].

We begin with a comparison of TS and GA. The quality of solution obtained and runtime required using different number of processors for both TS and GA are tabulated in Table 1. For each circuit, the number of cells are given in the table. The '$\mu(s)$ TS' and '$\mu(s)$ GA' columns show the aggregate fuzzy membership of solution obtained by TS and by GA respectively, whereas 'p' denotes the number of processors used. It should be noted that runtimes shown are for achieving a certain fixed quality.

Parallel TS clearly outperformed Parallel GA, both in terms of quality and runtimes. As can be observed, in case of large circuits, parallel GA was unable to find a reasonable quality solution even after running for a large amount of time. Even for smaller circuits, the solution quality obtained using TS is significantly superior to that obtained using GA. The proposed parallel TS has shown a consistent trend in terms of speedup with increasing number of processors. On the other hand, parallel GA did not show such performance or trend.

Table 1. Run times and solution quality $\mu(s)$ for achieving a target membership for serial and parallel TS/GA approaches. UH indicates unreasonably high runtime requirement.

Circuit Name	# of Cells	$\mu(s)$ TS	Time for Serial TS	Time for Parallel TS						$\mu(s)$ GA	Time for Serial GA	Time for Parallel GA		
				p=2	p=3	p=4	p=5	p=6	p=7			p=3	p=5	p=7
s386	172	0.688	52	28	20	17	16	15	14	0.504	15	9.9	5.7	6.7
s641	433	0.785	934	472	332	239	205	171	151	0.616	793	307	390	289
s832	310	0.644	74	40	33	23	22	20	19	0.479	128	43	37	39
s953	440	0.661	195	98	71	53	46	41	36	0.511	309	136	91	108
s1196	561	0.653	374	187	132	97	88	78	67	0.484	988	327	262	205
s1488	667	0.603	259	131	93	69	63	55	49	0.482	1883	677	435	418
s1494	661	0.601	268	137	96	72	65	57	51	0.496	1405	847	638	479
c3540	1753	0.665	2142	1146	703	547	440	370	344	-	UH	UH	UH	UH
s3330	1961	0.699	1186	590	451	313	245	210	184	-	UH	UH	UH	UH
s5378	2993	0.669	1850	914	601	467	371	312	264	-	UH	UH	UH	UH
s9234	5844	0.631	5571	2855	2006	1525	1272	1062	849	-	UH	UH	UH	UH

Table 2 illustrates the quality of solution obtained for the two SimE parallelization schemes, namely random row distribution strategy and fixed row distribution strategy. The table illustrates the amount of time for the taken to reach a predefined fuzzy membership with increasing number of processors. For the strategy proposed in this paper, as can be seen, there is a constant decrease in runtime, for all circuits, with increase in number of processors. Better trends are observed for medium to large circuits, than for smaller ones, as can be seen in Figure 5(a). Speedup is also illustrated in the bar-chart given in Figure 5(b). Due to space restrictions, and scaling factor limitations, not all results have been included in the same figure for sake of clarity.

The fitness values achieved with the proposed random row allocation are consistently higher in all test cases when compared to the fixed row allocation scheme, as shown by the *Qual Fixed* column in Table 2, the fixed row allocation never equals 100% of the solution quality obtained by the proposed
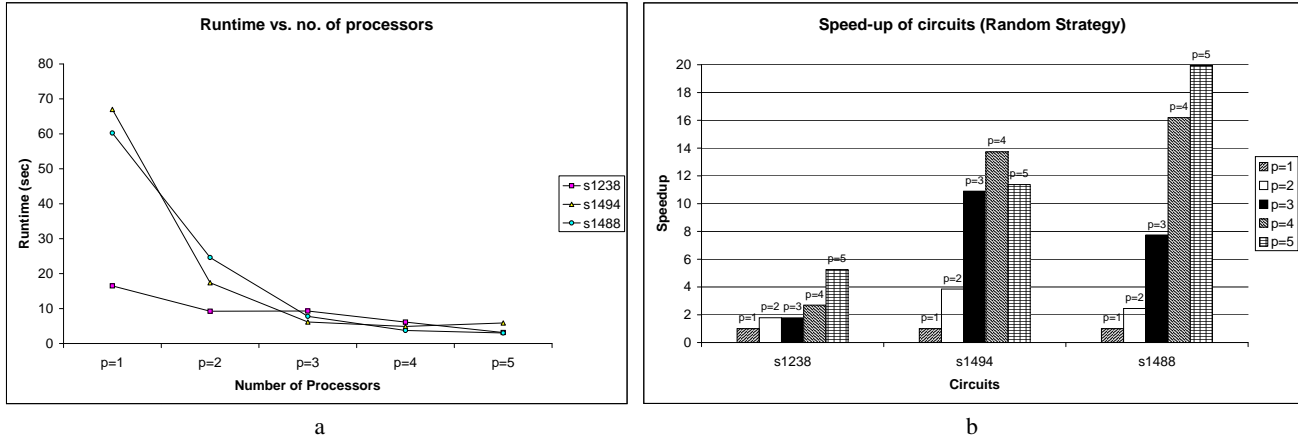
Fig. 5. (a) The decrease in runtime to reach a pre-defined fitness objective with increasing number of processors; (b) Speedup versus number of machines.

scheme. Further, the run times are far better, and the speedup is super linear in most cases for the random row distribution strategy. This can be attributed to modified working space of the selection and allocation operators on each slave, as in each iteration different sets and combination of rows are addressed. This has resulted in even more reduced times to obtain desired solution quality than with workload partitioning alone.

Table 2. Run times of Parallel SimE for a achieving a target fitness, for serial and parallel implementations, for both random and fixed row allocation strategies. UH indicates unreasonably high times.

| Circuit | # of | Time for Random Row Distribution | | | | | Qual | Time for Fixed Row Distribution | | | |
Name	Cells	$p=1$	$p=2$	$p=3$	$p=4$	$p=5$	Fixed	$p=2$	$p=3$	$p=4$	$p=5$
s641	433	UH	4.99	4.97	3.99	3.87	79.7%	9.14	1.08	0.76	0.55
s1238	540	16.5	9.24	9.29	6.12	3.14	95.8%	17.83	8.47	11.30	5.71
s1494	661	67	17.4	6.15	4.88	5.89	82.3%	2.77	1.85	1.76	4.34
s1488	667	60.23	24.6	7.78	3.72	3.02	96.6%	22.0	4.89	5.1	16
s3330	1961	UH	678.02	115	108.5	49.14	33.8%	316	215	4.6	3.4
s5378	2993	UH	1620	338.2	286.6	178.6	46.8%	UH	UH	124.3	95.0

When SimE was compared to other heuristics, the following was observed. For GAs, the time for completion to obtain solutions of a certain pre-specified quality were exorbitantly high. And in some cases, for the given run-time, acceptable solutions could not be obtained. For example, for the S1494, the serial GA implementation took 1883 Seconds, and when the parallel version was executed on 7 processors the best time was 418 Seconds (with 8% inferior quality than that obatined by SimE).

When comparing Parallel SimE with parallel TS, better quality was obtained in some cases at the cost of high computation time using TS, for the same quality the run-time requirements for TS were over three times more than that required by parallel SimE. For example, for s1494, the time taken by serial TS was 268 Seconds, and when parallel TS was run on 6 processors, the runtime was 57 Seconds, (compared to 5 Seconds by SimE) with slightly better quality, and TS took over 15 Seconds to obtain solutions of same quality as SimE. A similar trend was seen for all circuits.

Table 3. Run times and solution quality $\mu(s)$ for achieving a target membership for serial and parallel SA approaches.

Circuit Name	# of Cells	$\mu(s)$ SA	Time for Serial SA	Time for Parallel SA					
				p=3	p=4	p=5	p=6	p=7	p=8
s386	172	0.659	18	9.5	7.2	6.7	5.5	5.0	4.7
s641	433	0.755	221	189.3	169.7	129.6	125.2	120.5	117.7
s832	310	0.638	45	35.0	25.3	19.6	15.6	14.9	14.2
s953	440	0.704	122	96.6	83.4	60.1	56.5	54.5	53.4
s1196	561	0.675	190	145.9	130.9	110.3	96.9	98.2	94.8
s1488	667	0.650	275	151.4	118.4	112.6	98.8	94.0	92.6
s1494	661	0.647	214	131.4	116.2	101.9	98.1	92.3	89.1
c3540	1753	0.734	2445	2153.0	1950.5	1510.2	1306.2	1288.2	1124.4
s3330	1961	0.793	2137	1875.5	1658.6	1572.9	1419.5	1254.5	1081.2

For simulated annealing, the asynchronous multiple-Markov chain parallelization strategy was chosen [7]. In this strategy, each slave processor runs the Metropolis loop for a fixed number of iterations. Upon completion of the loop, the slave replaces the best solution in the master only if it has found a better one, else, it receives the solution at the master to be used as an initial solution for subsequent iterations. Since all slaves are running the same Metropolis loop for the same number of iterations, the decrease in the runtime to find a solution of a specified quality is not significant, however a decreasing trend in required runtime is observed. As can be see from Table 3, given enough time, parallel SA was also able to achieve slightly better quality solutions than SimE. However, for a fixed slightly lower quality, SimE was seen to be increasingly faster than SA as processors were increased. For instance, for s1494, with 3 processors SA took 131 Seconds to achieve the desired quality, while SimE took only 6.1 Seconds. With 5 processors, SA required 101 Seconds on average, while SimE needed only 5.9 Seconds. Similar trends are seen for most circuits.

5. CONCLUSIONS

In this work, we presented the parallelization of non-deterministic heuristics, namely, Tabu Search and Simulated Evolution, with the objective of accelerating the solution to a constrained multiobjective VLSI cell placement problem. Fuzzy logic was resorted to assist in integration of the costs of multi-objectives. Profile analysis of sequential code was conducted to assist in selecting and engineering earlier proposed parallelization strategies. The TS implementation was based on a synchronous candidate list partitioning model whereas the implementation of parallel SimE was an enhancement of the approach proposed in [3, 4].

For comparison purposes, a parallel genetic algorithm (GA) based on the island model [5, 6], and a parallel SA based on the Asynchronous Multiple-Markov Chain Model [7] were implemented. Results of experiments on ISCAS-85/89 benchmark circuits revealed that for the proposed and implemented strategies, while TS exhibited linear speedups, it was possible to obtain super-linear speedups with SimE. However the quality of TS and SA was slightly superior to other techniques. While SA required the largest run times, it provided the best quality solutions. However, the efficiency (defined as Speedup/p, where p is the number of processors) was far below 100% for SA.

Finally, it can be concluded, that for appreciable quality solutions, simulated evolution with random row distribution SimE has exhibited dramatic speedups with increase in number of processors, even when compared to other, more established heuristics. The results obtained suggest that in scenarios

where placement quality considerations are overridden by design time constraints, the proposed parallel SimE algorithm with Random Row Distribution should be favored.

6. REFERENCES

[1] Sadiq M. Sait and Habib Youssef. *Iterative Computer Algorithms and their Application to Engineering*. IEEE Computer Society Press, December 1999.

[2] Sadiq M. Sait and Habib Youssef. *VLSI Physical Design Automation: Theory and Practice*. World Scientific, Singapore, 2001.

[3] Ralph M. Kling and Prithviraj Banerjee. ESP: A new standard cell placement package using simulated evolution. *Proceedings of 24th Design Automation Conference*, pages 60–66, 1987.

[4] Prithviraj Banerjee. *Parallel Algorithms for VLSI Computer-Aided Design*. Prentice Hall International, 1994.

[5] M. Toulouse, T. G. Crainic, and M. Gendreau. Issues in Designing Parallel and Distributed search Algorithms for Discrete Optimization Problems. *Publication CRT-96-36, Centre de recherche sur les transports, Université de Montréal, Montréal, Canada*, 1996.

[6] Erick Cant-Paz. A survey of parallel genetic algorithms. *Calculateurs Parallles, Reseaux et Systems Repartis*, 1998.

[7] John A. Chandy, Sungho Kim, Balkrishna Ramkumar, Steven Parkes, and Prithviraj Banerjee. An evaluation of parallel simulated annealing strategies with application to standard cell placement. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 16:398–410, April 1997.

[8] Sadiq M. Sait, Mahmood R. Minhas, and Junaid A. Khan. Performance and low-power driven VLSI standard cell placement using tabu search. *Proceedings of the 2002 Congress on Evolutionary Computation*, 1:372–377, May 2002.

[9] Van-Dat Cung, Simone L. Martins, Celso C. Riberio, and Catherine Roucairol. Strategies for the Parallel Implementation of metaheuristics. *Essays and Surveys in Metaheuristics*, pages 263–308, Kluwer 2001.

[10] I. De Falco, R. Del Balio, E. Tarantino, and R. Vaccaro. Improving search by incorporating evolution principles in parallel tabu search. In *Proc. of the first IEEE Conference on Evolutionary Computation- CEC'94*, pages 823–828, June 1994.

[11] Bruno-Laurent Garica, Jean-Yves Potvin, and Jean-Marc Rousseau. A parallel implementation of the tabu search heuristic for vehicle routing problems with time window constraints. *Computers & Operations Research*, 21(9):1025–1033, November 1994.

[12] M. Malek, M. Guruswamy, M. Pandya, and H. Owens. Serial and parallel simulated annealing and tabu search algorithms for the traveling salesman problem. *Annals of Ops. Res.*, 21:59–84, 1989.

[13] J. Chakrapani and J. Skorin-Kapov. Massively parallel tabu search for the quadratic assignment problem. *Annals of Operations Research*, 41:327–341, 1993.

[14] E. Taillard. Robust tabu search for the quadratic assignment problem. *Parallel Computing*, 17:443–455, 1991.

[15] T. G. Crainic, M. Toulouse, and M. Gendreau. Towards a taxonomy of parallel tabu search heuristics. *INFORMS Journal of Computing*, 9(1):61–72, 1997.