Comparative evaluation of platforms for parallel Ant Colony Optimization

Ginés D. Guerrero · José M. Cecilia · Antonio Llanes · José M. García · Martyn Amos · Manuel Ujaldón

Published online: 18 March 2014 © Springer Science+Business Media New York 2014

Abstract The rapidly growing field of *nature-inspired computing* concerns the development and application of algorithms and methods based on biological or physical principles. This approach is particularly compelling for practitioners in high-performance computing, as natural algorithms are often *inherently parallel* in nature (for example, they may be based on a "swarm"-like model that uses a population of agents to optimize a function). Coupled with rising interest in nature-based algorithms is the growth in *heterogenous computing*; systems that use more than one kind of processor. We are therefore interested in the performance characteristics of nature-

G. D. Guerrero

J. M. Cecilia (⊠) · A. Llanes Computer Science Department, Universidad Católica San Antonio de Murcia, Murcia, Spain e-mail: jmcecilia@ucam.edu

A. Llanes e-mail: allanes@ucam.edu

J. M. García Computer Engineering Department, University of Murcia, Murcia, Spain e-mail: jmgarcia@ditec.um.es

M. Amos School of Computing, Mathematics and Digital Technology, Manchester Metropolitan University, Manchester, UK e-mail: m.amos@mmu.ac.uk

M. Ujaldón Computer Architecture Department, University of Malaga, Malaga, Spain e-mail: ujaldon@uma.es

National Laboratory for High Performance Computing, University of Chile, Santiago, Chile e-mail: gguerrero@nlhpc.cl

inspired algorithms on a *number* of different platforms. To this end, we present a new OpenCL-based implementation of the Ant Colony Optimization algorithm, and use it as the basis of extensive experimental tests. We benchmark the algorithm against existing implementations, on a wide variety of hardware platforms, and offer extensive analysis. This work provides rigorous foundations for future investigations of Ant Colony Optimization on high-performance platforms.

Keywords Heterogeneous computing \cdot Ant Colony Optimization \cdot CUDA \cdot OpenCL \cdot APU \cdot GPU

1 Introduction

Algorithms inspired by *natural* processes are gaining increasing acceptance, and are now used in a wide variety of application domains [28]. Many nature-inspired methods (such as the genetic algorithm [16], or particle swarm optimization [20]) are *population-based*, meaning that they maintain a *collection* of individual solutions which evolves or is modified as the computation proceeds. This structure naturally lends itself to *parallelisation*, and many parallel versions of such algorithms now exist [1].

One nature-based method that is proving to be increasingly popular is *Ant Colony Optimization* (ACO) [8,10,13]. This algorithm is based on foraging behaviour observed in colonies of real ants and has been applied to a wide variety of problems, including vehicle routing [32], feature selection [6] and autonomous robot navigation [15]. The method generally uses simulated "ants" (i.e., mobile agents), which first construct tours or paths on a network structure (corresponding to solutions to a problem), and then deposit "pheromone" (i.e., signalling chemicals) according to the quality of the solution generated. The algorithm takes advantage of emergent properties of the multi-agent system, in that positive feedback (facilitated by pheromone deposition) quickly drives the population to high-quality solutions.

The original ACO method (called the *Ant System* [11]) was developed by Dorigo in the 1990s, and this version (or slight variants thereof, such as the MAX-MIN Ant System (MMAS) [31]) is still in regular use [5,19,21]. Parallel versions of the Ant System have been developed [7,23,30,33] (see also [26] for a survey), and, in recent work, we present a graphics processor unit (GPU)-based version of ACO that, for the first time, parallelizes *both* main phases of the algorithm (that is, tour construction *and* pheromone deposition) [3,4].

The original version of our algorithm was developed for the CUDA (Compute Unified Device Architecture) platform,¹ which offers easy access to the parallel processing capabilities of GPUs (thus facilitating so-called "GPGPU" or "general purpose GPU" computation). Although it laid the foundations for general GPU-based computing, CUDA is proprietary to Nvidia, one of the dominant manufacturers in the GPU market. With that in mind, an alternative open-standard was developed, which became known as *OpenCL* (Open Computing Language) [29]. This standard provides a com-

¹ Full technical details at http://docs.nvidia.com/cuda/index.html.

mon language, programming interfaces and hardware abstractions over a *wide range* of devices (CPUs, GPUs and other accelerators), and has contributed significantly to the growth of *heterogeneous computing* [2]. Importantly, OpenCL offers *portability* across combinations of operating system, GPU and other processors, which, in turn, have their own hardware costs and performance characteristics. It is therefore possible to write a portable, parallel algorithm for a specific problem, which may run on a hardware/software combination that meets multiple constraints (cost, performance, and so on).

With that in mind, we present a new OpenCL-based version of our ACO algorithm, which may run on a variety of platforms (from laptops to high-end servers). Our aim is to demonstrate how such an implementation may be used as the foundation for *high-performance, portable* ACO-based solutions. We benchmark our algorithm on a range of platforms and give an analysis about its scalability on high-end platforms.

The paper is organized as follows: in Sect. 2 we briefly describe our ACO-based algorithm and the process of migrating it to OpenCL. We then present the results of experimental investigations in Sect. 3, offer some analysis in Sect. 4, and then conclude in Sect. 5 with a brief discussion of our findings.

2 ACO algorithm

Our ACO-based solution to the Travelling Salesman Problem (TSP) is described in detail in [3,4], so here we simply give a brief overview to highlight specific issues arising from the migration to OpenCL.

The TSP is a well-known *NP*-hard optimization problem, and is often used as a standard benchmark for heuristic algorithms [18]. Indeed, it was the first problem to be solved using ACO [11], and our own work is a natural development of this. Briefly, the TSP involves finding the shortest ("cheapest") round-trip route that visits each of a number of "cities" exactly once. In what follows, we address the symmetric TSP on *n* cities, which may be represented as a complete weighted graph, *G*, of *n* nodes, with each weighted edge, $e_{i,j}$, representing the inter-city distance $d_{i,j} = d_{j,i}$ between cities *i* and *j*. The ACO algorithm for TSP uses a number of simulated "ants" (or *agents*), which perform distributed search on a graph. Each ant moves on the graph until it completes a tour, and then offers this tour as its suggested solution. To achieve this latter step, each ant deposits "pheromone" on the edges that it visits during its tour. The quantity of pheromone deposited, if any, is determined by the *quality* of the solution relative to those obtained by the other ants. Pheromone levels on each edge "evaporate" over time (i.e., they are gradually reduced), to prevent the algorithm from being locked into sub-optimal solutions.

While building a tour, each ant probabilistically chooses the next city to visit based on two different sources of information: (1) heuristic information, obtained from intercity distances, and (2) the pheromone trail, which facilitates indirect communication between ants via their *environment* (a process known as *stigmergy* [9]). The combination of local search and global signalling enables a process of directed positive feedback, by which the population quickly converges to a high-quality solution to the problem. The main body of the algorithm, therefore, has two main phases: (1) *tour construction*, and (2) *pheromone deposition*.

During tour construction, a number of ants build tours in parallel. Ants are initially placed at random, and they then repeatedly apply a probabilistic action choice rule to decide which city to visit next. Pheromone deposition occurs once all ants have constructed their tours; first, the pheromone levels on all edges are reduced by a constant factor (to simulate evaporation), and then pheromone is deposited on edges that ants have included in their tours (the precise amount for each edge in a particular tour being inversely proportional to the tour's length). In this way, edges that are used by many ants (and which are part of short tours) receive more pheromone, and are therefore more likely to be selected by ants in subsequent rounds (thus implementing the positive feedback process that we have already described).

2.1 Original CUDA implementation

We first briefly review the main characteristics of CUDA [24], for the benefit of readers who are unfamiliar with the programming model. CUDA is based on a hierarchy of abstraction layers; the *thread* is the basic execution unit; threads are grouped into *blocks*, each of which run on a single multiprocessor, where they can share data on a small but extremely fast memory. A *grid* is composed of blocks, which are equally distributed and scheduled among all multiprocessors. The parallel sections of an application are executed as *kernels* in a SIMD (Single Instruction Multiple Data) fashion, that is, with all threads running the same code. A kernel is therefore executed by a grid of thread blocks, where threads run simultaneously grouped in batches called *warps*, which are the scheduling units.

We now consider the implementation of each phase of the algorithm. The "traditional" task-based parallelism approach is based on the observation that ants run in parallel while searching for the best tour [4] (that is, parallelism is expressed at the level of individual ants). Within the basic model, each ant is associated with an individual thread, but this approach has three main drawbacks:

- 1. **Low degree of parallelism**. Because the number of ants used is generally a (linear) function of the problem size, the number of threads required is generally too low to fully exploit the resources of the GPU.
- 2. **Control dependencies**. *Warp divergences* (a situation where threads take different control-flow paths) can often arise when ants check the so-called *tabu list*—the record of cities already visited. Put simply, different threads in a warp may need to do different things, depending on which cities the different ants have visited, and this is expensive.
- 3. **Irregular memory access**. Because the ACO algorithm is inherently stochastic, this can produce an unpredictable *memory access pattern*. This prevents the GPU from taking advantage of *caching* schemes and other techniques for reducing memory access latency.

In previous work, we developed an alternative approach that places more emphasis on *data parallelism* [3]. We now briefly describe this algorithm, to establish the differences between the CUDA and OpenCL implementations. When an ant makes a decision on which city to visit next, it must calculate heuristic information, as previously described. The heuristic information available to any one ant at a given time is *the same*, regardless of which ant is making the query, so it makes sense to separate out the computation of heuristic values into a separate *heuristic info kernel*, which is then executed prior to tour construction. Transition probabilities are stored in a two-dimensional *choice matrix*, which is used to inform "roulette wheel" (Monte Carlo) selection by each ant.

In the *tour construction* kernel, each ant is associated with a *thread block*, such that each thread represents a city (or cities) that the ant may visit. This avoids the problem of warp divergences, and enhances data parallelism, as all threads within a block may *co-operate*. The degree of parallelism improves by a factor of 1 : w, where w is the number of CUDA threads per block.

Finally, the *pheromone kernel* performs evaporation and deposition, as described earlier. Evaporation is straightforward, as a single thread can independently lower each entry in the pheromone matrix by a constant factor. Deposition is more problematic, as each ant generates its own private tour in parallel, and will eventually visit the same edge as another ant. Therefore, to prevent race conditions, we require the use of CUDA atomic operations when accessing the pheromone matrix.

The efficiency of a parallel implementation is also affected by the *types* of operation on which it relies; in our code, scatter/gather operations [17] predominate (i.e., those which either write or read a large number of data items). As Table 1 reflects, the vast majority of operations are of the "gather" type; algorithms of this type are memorybounded and amenable to optimization via methods such as *coalescing* (Nvidia GPUs) and the use of SSE *vector instructions* (Intel CPUs). A comparative study [22] of these optimisations reveals similar impact on performance across platforms, which suggests that the experimental sections of the current paper will not suffer too much from platform-specific biases.

2.2 OpenCL migration

In this section, we briefly describe various issues that arose during the migration from CUDA to OpenCL. The foundations of OpenCL are based on the CUDA threading model, but with differences in terms of naming schemes and identifiers. We therefore used source-to-source translation to migrate our CUDA-based kernels to OpenCL. This mapping requires in-depth knowledge of both application programming interface (API) models, as it is considerably more complex than simple instruction conversion. Also, OpenCL is still relatively young compared to CUDA, and does not provide the same functionality offered by its more mature partner.

The process of setting up a device for kernel execution differs *substantially* between CUDA and OpenCL. The APIs for context creation and data copying use different conventions for mapping the kernel onto the device processing elements, which may substantially affect the programming effort required to code and debug a parallel application. CUDA provides several libraries to enhance the functionality of its API. For example, our ACO algorithm uses the CURAND library [25] to generate pseudo-random numbers. This library is not directly implemented in OpenCL, where the main

Algorithm stage	Operator	Key features	CUDA kernel
Generation of choice_info array	Gather	Data parallelism fully exploited	choice_info
Tour construction	Gather	Optimized via choice_info array	Next_tour
Tabu list update	Scatter	Optimized via an array in register file	Next_tour
Pheromone evaporation	Scatter	Concurrent updates, no queries	Pheromone
Pheromone deposit	Gather	Single update using atomic operations	Pheromone

Table 1 Characterization of the stages involved in our ACO implementation on GPUs

alternative is an implementation of the RANLUX pseudo-random number generator, called RANLUXCL.² Unfortunately, we found this library to be fairly wasteful in terms of memory, so we decided to implement our own, taking a C counterpart as a departure point [14].

3 Experimental results

In this section we give the results of extensive comparative evaluations of ACO-based solutions to the TSP on different CPU, APU and GPU platforms. The underlying hardware platforms we tested are specified in Table 2.

For validation purposes, we use a baseline comparison with the sequential ANSI C code provided in [12]. The experimental setup (in terms of hardware/software) is listed in Table 3. We run our three ACO implementations (ANSI C, CUDA and OpenCL) on selected benchmark TSP instances from the well-known TSPLIB library [27]. All instances are defined on a complete graph, and distances are given as integers. Table 4 specifies the instances used; they were selected to ensure a representative sample, from "small" to "medium" and "large" (for reasons of practicality, we test only the high-end platforms on pr2392; these results are used for the later scalability analysis). Importantly, we note that our methods solve all instances *to optimality*; for the purposes of this paper, we are less interested in the *quality* of solutions produced, so to ensure a fair comparison we use instances that are solvable to optimality by our implementations as described in [3].

For all runs, we set the ACO parameters according to the values recommended in [12]; $\alpha = 1$, $\beta = 2$, $\rho = 0.5$, and m = n, meaning that the number of ants, m, is equal to the number of cities, n. We run each algorithm for 1,000 iterations, and average timings over 1,000 runs. CUDA times are obtained with a block size of 128 threads, and OpenCL local size is also set to 128.

Before discussing the results of our experiments, we consider several issues with respect to performance. First, APUs are much more limited in terms of thermal design power, as they must also include the CPU. This means that execution units will need to be removed to keep power consumption down. Second, because the APU is a costeffective solution, it does not have its own dedicated global memory, but instead it

² See https://bitbucket.org/ivarun/ranluxcl/.

	CPU	GPU	GPU
(a) Processors found in high-	end servers		
Release date	Q4 2009	Q4 2009	Q1 2010
Codename	Intel Westmere	Nvidia Fermi	ATI Cypress
Commercial model	Xeon E5620	Tesla C2050	FirePro V8800
No. cores @ speed	4 @ 2.4 GHz	_	-
No. stream processors	_	448 @ 1.15 GHz	1,600 @ 925 MHz
L2 cache size	12 MB	768 KB	512 KB
DRAM memory size	16 GB	3 GB	2 GB
DRAM type	DDR3	GDDR5	GDDR5
Memory bus width	128 bits	384 bits	256 bits
Memory clock	1,066 MHz	$2 \times 1.5 \text{ GHz}$	$4 \times 1.15 \text{ GHz}$
Memory bandwidth	17 GB/s	144 GB/s	147.2 GB/s
	CPU on APU	GPU on APU	
(b) Processors found in deskt	top PCs		
Release date	Q1 2010	Q1 2010	
Codename	AMD Llano	ATI Redwood	
Commercial model	E-350	ATI HD 6310	
No. cores @ speed	2 @ 1.6 GHz	-	
No. stream processors	-	80 @ 492 MHz	
L2 cache size	$2 \times 512 \text{ KB}$	-	
DRAM memory size	4 GB (shared)	4 GB (shared)	
DRAM type	DDR3	DDR3	
Memory bus width	64 bits	64 bits	
Memory clock	1,066 MHz	1,066 MHz	
Memory bandwidth	8.5 GB/s	8.5 GB/s	
	CPU on APU	GPU on APU	GPU
(c) Processors found in lapto	ps		
Release date	Q2 2011	Q2 2011	Q1 2011
Codename	AMD Llano	ATI Redwood	ATI Redwood
Commercial model	A6-3420	Radeon HD 6520	Radeon HD 6650M
No. cores @ speed	4 @ 1.4 GHz	-	-
No. stream processors	-	320 @ 400 MHz	480 @ 600 MHz
L2 cache size	4 MB	-	-
DRAM memory size	4 GB (shared)	4 GB (shared)	1 GB (exclusive)
DRAM type	DDR3	DDR3	DDR3
Memory bus width	64 bits	64 bits	128 bits
Memory clock	1,333 MHz	1,333 MHz	900 MHz
Memory bandwidth	10.6 GB/s	10.6 GB/s	14.4 GB/s

 Table 2
 Summary of hardware features for the CPUs, APUs and GPUs used during our experimental survey

Target hardware	Software tools
Intel Xeon CPU	gcc compiler, 4.3.4 version with the -O3 flag set
Nvidia Tesla GPU	CUDA compilation tools, release 4.0
ATI FirePro GPU	Software Suite 8.85.7.2 and OpenCL runtime v831.4
AMD APUs and dedicated GPUs	AMD's APP SDK 2.6, Catalyst driver 11.12, OpenCL runtime version 793.1

Table 3 Software resources used for each hardware platform in our experimental study

Table 4 TSP instances used in our s	tudy
-------------------------------------	------

Small dataset				Medium/large dataset					
Graph name	d198	a280	lin318	pcb442	rat783	pr1002	pcb1173	d1291	pr2392
Number of cities	198	280	318	442	783	1,002	1,173	1,291	2,392
Best tour length	15,780	2,579	42,029	50,778	8,806	259,045	56,892	50,801	378,032

relies on an emulated global memory located in system memory. While this is good for performance when transferring data directly between the CPU and GPU, it means that it will also suffer in terms of overall bandwidth, as even low-end GPUs have more memory bandwidth.

We present a summary of our results in Fig. 1. For each row (i.e., each platform, or hardware/software combination), we show execution times averaged over the small (top bar) and medium/large (bottom bar) instances. Note that times are measured in milliseconds (ms), and represent the elapsed time for a *single iteration* of the platform-specific algorithm, averaged over 100 runs of 1,000 iterations each (as opposed to the average run time for the whole algorithm). We focus on the average time for a single iteration precisely because we are interested in the *overall kernel performance* on each platform, so this fine-grained approach gives us the insights that we require.

4 Analysis

We now give an analysis of the performance of each category of hardware platform.

4.1 Desktop PCs

Beginning with the E-350 APU, we see that the CPU does not perform particularly well. This is expected, based on the architecture's emphasis on power consumption over performance for this consumer market. However, when moving to the GPU we see that, for small problem instances, it actually scales better in terms of overall computational time than the FirePro V8800 for the same base architecture.

Looking closely at the numbers, the E-350 APU, which is outclassed by factors of 37 and 17 for computational power (*execution resources* \times *clock speed*) and memory



Fig. 1 Summary of experimental results. X axis shows each platform, Y axis (logarithmic plot) shows execution time (ms) for one iteration. Bars are ordered from the smallest (*left*) to the largest (*right*) instances

bandwidth, respectively, manages to only perform at roughly 1/10th the speed. We attribute this to the APU's ability to quickly transfer data to and from the CPU to the GPU. However, as the input size increases this advantage disappears, as raw computational throughput and bandwidth become more important than latency. Comparing these results to the Tesla C2050 GPU, the APU is at an even greater disadvantage, due to its VLIW architecture (compared to the scalar and compute-oriented architecture of the C2050). This should change, however, with AMD future generations of APUs, which consider a GPU based on their newly released Graphics Core Next (GCN) architecture. GCN greatly improves computational throughput, by moving scheduling from the compiler to the hardware.

4.2 Laptop computers

Moving to the A6-3420M APU, we see very similar results as with the E-350 APU. Here, our integrated GPU (iGPU) has roughly 3 times the amount of computational resources, but only 1/4 more bandwidth. This is evident in the scaling of the algorithm, as we go from 200 ms. with the E-350 APU to 148 ms. with the iGPU, a near exact scaling of the bandwidth advantage that the iGPU possesses.

As we increase the size of the problem instance, we then see that performance becomes constrained more by computational resources than by bandwidth. Our simplest comparison is to the A6-3420M's dedicated GPU (which has 2.25 times the computational power), where, as the input size increases, the difference between the two solutions approaches this limitation. This shows that, while the memory system influences ACO performance, computational resources become the dominating factor in overall performance. For PCI-express 2.0, the maximum bandwidth (unidirectional) is 8 GB/s, while using zero copy the APU is able to reach nearly 16 GB/s. If this had been taken into account, the results for the APU and dedicated GPUs would in fact be much closer, as this type of workload/data transfer is playing to the APU's strength.

Scalability → Language/API HW platform		Short range	Mid range	Long range
		Time(pr2392)/ Time(rat783)	Time(rat783)/ Time(d198)	Time(pr2392)/ Time(d198)
С	CPU Xeon	31.24 <i>x</i>	82.30 <i>x</i>	2,571.46 <i>x</i>
CUDA	GPU Tesla	24.43 <i>x</i>	45.74 <i>x</i>	1,117.88 <i>x</i>
OpenCL	GPU Tesla	24.16 <i>x</i>	44.83 <i>x</i>	1,083.52 <i>x</i>
OpenCL	GPU FirePro	22.70 <i>x</i>	17.09 <i>x</i>	388.12 <i>x</i>

 Table 5
 Scalability on high-end-platforms depending on hardware and programming methods

FirePro behaves better on larger problem instances, followed by Tesla using OpenCL (with CUDA very close), and finally Xeon using C

Ending with the dedicated GPU (dGPU), we see a similar speedup increase, just as we did for the E-350 and A6-3420M APUs. Again, for small input sizes, latency and bandwidth are much more important than the computational abilities of the device, as there are fewer threads to interleave to hide memory accesses. This is visible in the dGPU, which has 7 and 2 times the amount of computational power and memory bandwidth as the E-350 APU, while performing just over twice as quickly for the d198 dataset.

As we increase the complexity of the workload, we again see that memory bandwidth becomes a less important issue, and computational power becomes the main contributing factor for overall performance. Comparing once again to the Tesla C2050, the APU solution does not perform as well as we had hoped.

4.3 High-end platforms

High-end processors usually cover large-scale applications, and our performance analysis emphasises *scalability*. Table 5 shows the behaviour of the execution time when the problem size increases. We compare execution times on small, medium and large instances, and obtain the coefficient or multiplier which separates them. The larger this coefficient is for a given processor, the poorer the degree of scalability.

Looking at those numbers, we see that when comparing Tesla versus FirePro (GPUs running the same OpenCL code), Tesla is 1.5x-2x faster, but FirePro scales better. Also, comparing languages on the same Tesla hardware, CUDA is 1.15x-1.20x faster, but OpenCL scales slightly better. Finally, comparing GPU results with numbers on the CPU, the GPU is faster and scales better: The speed-up factor ranges 9x-15x on four small data sets, 17x-20x on four medium data sets, and finally 21.5x on the large data set.

5 Conclusions

In this paper we presented a comprehensive performance review of different platforms for Ant Colony Optimization, an emerging and fast-growing nature-inspired algorithm.

We discussed the translation of our previous algorithm from CUDA to OpenCL, and highlighted certain issues that may be faced by other practitioners in future. We then performed a performance analysis of three variants of the ACO algorithm, using the Travelling Salesman Problem as a benchmark, and focussed on issues of scalability.

In general, GPUs are superior to CPUs on the high-end segment: they yield twenty times faster execution on large problem instances. The GPU–CPU difference is similar on desktops and laptops, 10-20x in favor of GPUs. At an early stage of its evolution, the APU offers a low-cost platform, without powerful computational units nor swift memory data paths. Our results demonstrate that these two issues have a severe impact on performance.

The growth of heterogeneous systems represents a solid trend in modern systems, and we believe that future work on Ant Colony Optimization in this domain can benefit from the promising insights into scalability demonstrated by our experimental study.

Acknowledgments This work is jointly supported by the Fundación Séneca (Agencia Regional de Ciencia y Tecnología, Región de Murcia) under grant 15290/PI/2010, by the Spanish MEC and European Commission FEDER under grant TIN2012-31345, by the UCAM under grant PMAFI/26/12, by the Junta de Andalucía under Project of Excellence P12-TIC-1741 and by the supercomputing infrastructure of the NLHPC (ECM-02). We also thank NVIDIA for hardware donation under CUDA Teaching Center 2011-14, CUDA Research Center 2012-14 and CUDA Fellow 2012-14 Awards.

References

- Alba E, Luque G, Nesmachnow S (2013) Parallel metaheuristics: recent advances and new trends. Int Trans Oper Res 20(1):1–48. doi:10.1111/j.1475-3995.2012.00862.x
- Brodtkorb AR, Dyken C, Hagen TR, Hjelmervik JM, Storaasli OO (2010) State-of-the-art in heterogeneous computing. Sci Progr 18(1):1–33
- Cecilia JM, Garcia JM, Nisbet A, Amos M, Ujaldón M (2013) Enhancing data parallelism for ant colony optimization on GPUs. J Parallel Distrib Comput 73(1):42–51
- Cecilia JM, Garcia JM, Ujaldon M, Nisbet A, Amos M (2011) Parallelization strategies for ant colony optimisation on GPUs. In: Proceedings of the 2011 IEEE international symposium on parallel and distributed processing. IEEE, pp 339–346
- Chang RSS, Chang JSS, Lin PSS (2009) An ant algorithm for balanced job scheduling in grids. Future Gener Comput Syst 25(1):20–27. doi:10.1016/j.future.2008.06.004
- Chen Y, Miao D, Wang R (2010) A rough set approach to feature selection based on ant colony optimization. Pattern Recognit Lett 31(3):226–233. doi:10.1016/j.patrec.2009.10.013
- Delévacq A, Delisle P, Gravel M, Krajecki M (2013) Parallel ant colony optimization on graphics processing units. J Parallel Distrib Comput 73(1):52–61. doi:10.1016/j.jpdc.2012.01.003
- 8. Dorigo M, Birattari M, Stutzle T (2006) Ant colony optimization. Comput Intell Mag IEEE 1(4):28-39
- 9. Dorigo M, Bonabeau E, Theraulaz G (2000) Ant algorithms and stigmergy. Future Gener Comput Syst 16(8):851–871
- Dorigo M, Di Caro G (1999) Ant colony optimization: a new meta-heuristic. In: Proceedings of the 1999 congress on evolutionary computation (CEC'99). IEEE Press, pp 1470–1477
- Dorigo M, Maniezzo V, Colorni A (1996) Ant system: optimization by a colony of cooperating agents. IEEE Trans Syst Man Cybern B 26(1):29–41
- 12. Dorigo M, Stutzle T (2004) Ant Colony Optimization. Bradford Company
- 13. Dorigo M, Stützle T (2010) Ant colony optimization: overview and recent advances. In: Handbook of metaheuristics. Springer, Berlin, pp 227–263
- Flannery BP, Press WH, Teukolsky SA, Vetterling W (1992) Numerical recipes in c. Press Syndicate of the University of Cambridge, New York

- Garcia MP, Montiel O, Castillo O, Sepúlveda R, Melin P (2009) Path planning for autonomous mobile robot navigation with ant colony optimization and fuzzy cost function evaluation. Appl Soft Comput 9(3):1102–1110. doi:10.1016/j.asoc.2009.02.014
- 16. Goldberg DE (1989) Genetic algorithms in search, optimization, and machine learning. Addison-Wesley Professional, Reading
- He B, Govindaraju NK, Luo Q, Smith B (2007) Efficient gather and scatter operations on graphics processors. In: Proceedings of the 2007 ACM/IEEE conference on supercomputing. ACM, New York, pp 46–57
- Johnson DS, McGeoch LA (1997) The traveling salesman problem: a case study in local optimization. In: Lenstra J, Aarts E (eds) Local search in combinatorial optimization. Wiley, New York, pp 215–310
- Ke BR, Chen MC, Lin CL (2009) Block-layout design using max-min ant system for saving energy on mass rapid transit systems. IEEE Trans Intell Transp Syst 10(2):226–235. doi:10.1109/TITS.2009. 2018324
- Kennedy J, Eberhart R (1995) Particle swarm optimization. In: Proceedings of IEEE international conference on neural networks, vol 4. IEEE, pp 1942–1948
- Komarudin Wong KY (2010) Applying ant system for solving unequal area facility layout problems. Eur J Oper Res 202(3):730–746. doi:10.1016/j.ejor.2009.06.016
- 22. Lee VW, Kim C, Chhugani J, Deisher M, Kim D, Nguyen AD, Satish N, Smelyanskiy M, Chennupaty S, Hammarlund P (2010) Debunking the 100x gpu vs. cpu myth: an evaluation of throughput computing on cpu and gpu. In: ACM international symposium on computer architecture. ACM, pp 451–460
- Manfrin M, Birattari M, Stützle T, Dorigo M (2006) Parallel ant colony optimization for the traveling salesman problem. In: Ant colony optimization and swarm intelligence. Springer, Berlin, pp 224–234
- Nickolls J, Buck I, Garland M, Skadron K (2008) Scalable parallel programming with cuda. Queue 6(2):40–53
- Nvidia (2011) CUDA Toolkit 4.0 CURAND Guide. http://developer.download.nvidia.com/compute/ DevZone/docs/html/CUDALibraries/doc/CURAND_Library.pdf
- Pedemonte M, Nesmachnow S, Cancela H (2011) A survey on parallel ant colony optimization. Appl Soft Comput 11(8):5181–5197. doi:10.1016/j.asoc.2011.05.042
- 27. Reinelt G (1991) TSPLIB—a traveling salesman problem library. ORSA J Comput 3(4):376–384. Library available at http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/
- 28. Rozenberg G, Bäck T, Kok JN (2011) Handbook of natural computing. Springer, Berlin
- Stone JE, Gohara D, Shi G (2010) OpenCL: a parallel programming standard for heterogeneous computing systems. Comput Sci Eng 12(3):66–72. doi:10.1109/MCSE.2010.69
- Stützle T (1998) Parallelization strategies for ant colony optimization. In: Parallel Problem Solving from Nature (PPSN V). Springer, Berlin, pp 722–731
- 31. Stutzle T, Hoos HH (2000) MAX-MIN ant system. Future Gener Comput Syst 16(8):889-914
- Yu B, Yang ZZ, Yao B (2009) An improved ant colony optimization for vehicle routing problem. Eur J Oper Res 196(1):171–176. doi:10.1016/j.ejor.2008.02.028
- Zhu W, Curry J (2009) Parallel ant colony for nonlinear function optimization with graphics hardware acceleration. In: IEEE international conference on systems, man and cybernetics, 2009, SMC 2009. IEEE, pp 1803–1808