Parallel Ant Colony Optimization for the Traveling Salesman Problem

Max MANFRIN, Mauro BIRATTARI, Thomas STÜTZLE, and Marco DORIGO

IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium

{mmanfrin,mbiro,stuetzle,mdorigo}@ulb.ac.be

ULB UNIVERSITÉ LIBRE DE BRUXELLES, UNIVERSITÉ D'EUROPE

IRIDIA Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle

Objectives

- parallelize a **high-performing ant colony optimization (ACO) algorithm** for the traveling salesman problem using **message passing libraries**
- study the **impact of communication** among **multiple homogeneous colonies interconnected** with **various topologies** on the **final solution quality** reached in a **fixed computation time**

The hypothesis is that the **exchange** of the **best-so-far solution** among different colonies **speeds up the search** for **high quality solutions**, having a positive impact on performance of the algorithms

Our contributions

- we use a high-performing ACO algorithms, **MAX-MIN Ant System** (MMAS), as a basis for our implementations
- we test several instances (**10 TSPLIB instances**) of large size (**size ranging from 1002 to 2392**)

Methods

- we choose the **traveling salesman problem**, a **NP-hard** problem, because it played a central role in research on ACO

Interconnection topologies: 8 CPUs (1 colony per CPU)

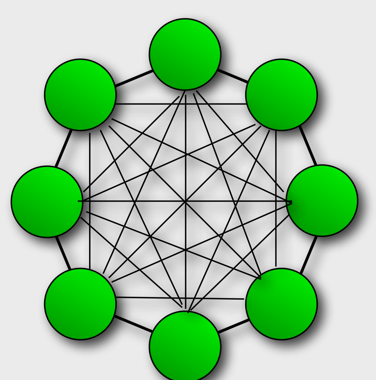

fully-connected

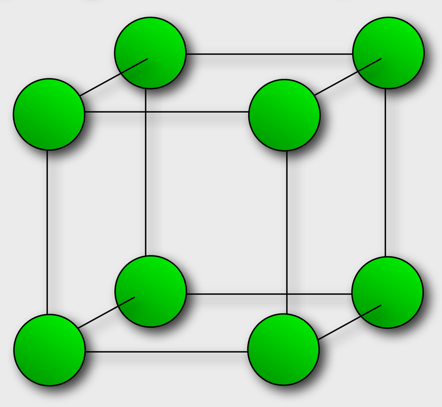

3-D hypercube

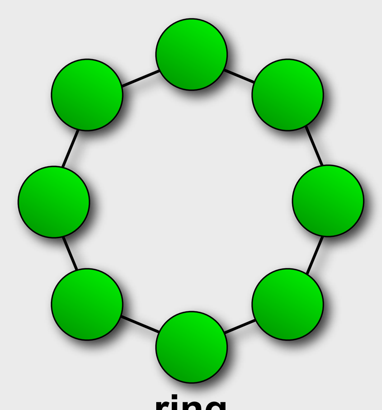

ring

Algorithms

SEQ	sequential MMAS; run-time equal to *8* parallel CPUs
SEQ2	sequential MMAS; run-time equal to 1 parallel CPU
PIR	*8* copies of SEQ2 (each with a different random seeds); chose the best final solution
FC	3*(*8*-1) messages per CPU at each communication step
RW	(*8* + 1) messages per CPU at each communication step
H	*6* messages per CPU at each communication step
R	*2* messages per CPU at each communication step

Implemented algorithms: "interconnection topology" vs "communication"

topology / commun.	fully-connected	hypercube	ring	isolated
synchronous	SFC, SRW	SH	SR	
asynchronous	AFC, ARW	AH	AR	
none				SEQ, SEQ2, PIR

- we extended the ACOTSP code[1] by **quadrant nearest neighbor list**, we **removed** the occasional **pheromone re-initializations**, and we used only a **best-so-far pheromone update**. We use here the **3-opt local search.**
- communication frequency: exchange every 25 iterations, except for the first 100 iterations
- a colony *injects* in his current solution-pool a received best-so-far solution if and only if it is better than its current best-so-far solution

1. http://www.aco-metaheuristic.org/aco-code/public-software.html

Results

Experiments with fixed constant rate communication (10 runs on 10 instances)

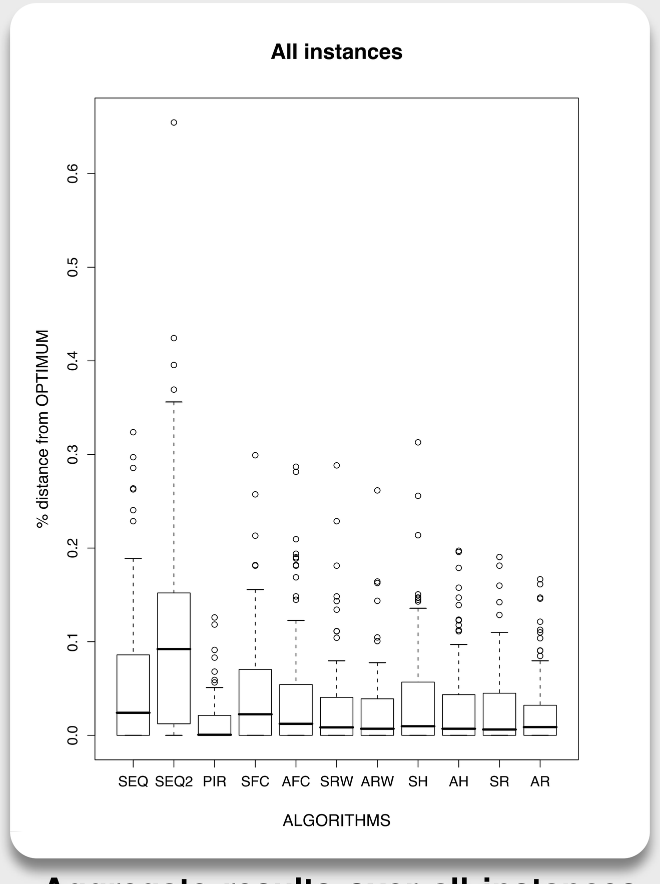

Aggregate results over all instances. Boxplot of normalized results

p-values for the null hypothesis "The distribution of the % distance from optimum of solutions for all instances is the same as PIR". The alternative hypothesis is that "The median of the PIR distribution is lower". The significance level with which we reject the null hypothesis is 0.05

SFC	AFC	SRW	ARW	SH	AH	SR	AR
5e-4	0.01	0.02	0.02	1e-3	0.02	0.02	0.02

p-values for the null hypothesis "The distribution of the % distance from optimum of solutions for all instances are the same". The significance level with which we reject the null hypothesis is 0.05

	FC	RW	H
RW	0.55	-	-
H	1	1	-
R	0.55	1	1

***Wilcoxon rank sum* test with *p-values* adjusted by Holm's method**

Qualified run-time distribution (RTD) to check "stagnation" behavior

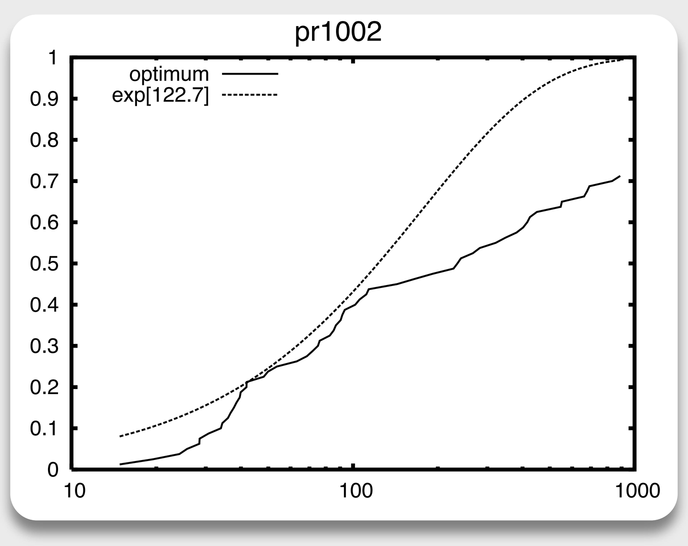

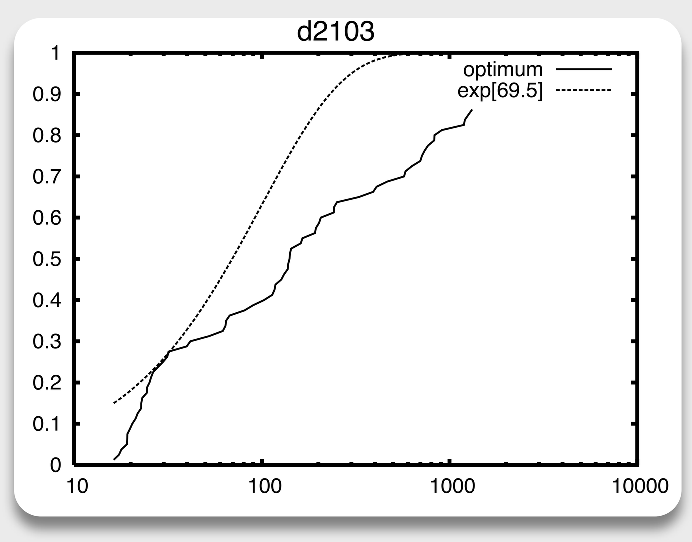

RTD over 80 independent trials of the sequential MMAS algorithm for the instances pr1002 and d2103

Experiments with less frequent constant rate communication (10 runs on 10 instances)

p-values for the null hypothesis "The distribution of the % distance from optimum of solutions for all instances is the same as PIR". The significance level with which we reject the null hypothesis is 0.05

SRW	SR	SRW2	SR2
0.02	0.02	0.3	0.3

***Wilcoxon rank sum* test with *p-values* adjusted by Holm's method**

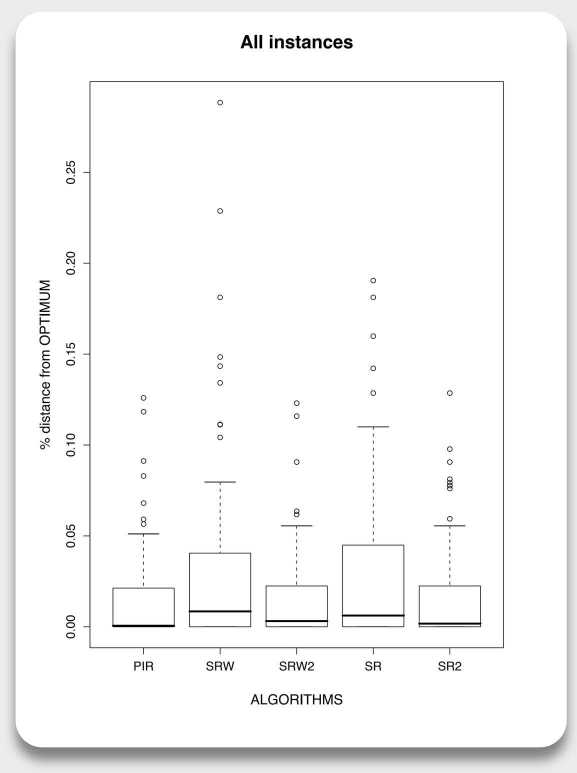

Aggregate results over all instances. Boxplot of normalized results

Conclusions

- the **parallel algorithms** considered **achieve**, on average, **better performance than the equivalent sequential one**
- stagnation behavior observed in run-time distributions explains the good performance of PIR. **In case of long run-times, PIR is apparently a good way of parallelizing**
- **less frequent communication schema** (constant rate dependent on the instance size) **produces better results than** the **more frequent communication schema** initially adopted (fixed constant rate)
- recent experiments on **larger instances** (size 3162) and with **other low-frequency communication schemes** show that with the ring topology **better results than with PIR can be obtained.**

Our conjecture is that PIR becomes less effective for increasing instance size

FNRS MARIE CURIE ACTIONS

Acknowledgments This work is supported by 'COMP²SYS' and by the 'ANTS' project. 'COMP²SYS' is a Marie Curie Early Stage Training Site, funded by the European Community's Sixth Framework Programme under contract number MEST-CT-2004-505079. The 'ANTS' project is an Action de Recherche Concertée funded by the Scientific Research Directorate of the French Community of Belgium. M. Dorigo and T. Stützle acknowledge support from the Belgian National Fund for Scientific Research (FNRS), of which they are a Research Director and a Research Associate, respectively. The information provided is the sole responsibility of the authors and does not reflect the opinion of the sponsors. The European Community and the French Community are not responsible for any use that might be made of data appearing in this publication.