ISSN 2186-7437

NII Shonan Meeting Report

No. 2012-5

Parallel Methods for Constraint Solving and Combinatorial Optimization

Philippe Codognet
Kazunori Ueda
Hiroshi Hosobe

May 27–31, 2012

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Parallel Methods for Constraint Solving and Combinatorial Optimization

Organizers:
Philippe Codognet (JFLI-CNRS/UPMC/University of Tokyo)
Kazunori Ueda (Waseda University)
Hiroshi Hosobe (National Institute of Informatics)

May 27–31, 2012

In the last decade, with the development of multi-core workstations, the availability of GPGPU-enhanced systems and the access to Grid platforms and supercomputers worldwide, Parallel Programming reached mainstream programming and appeared as a key issue in order to use in an efficient manner the computing power at hand.

Search methods and combinatorial optimization techniques are not isolated from this phenomenon, as bigger computing power means the ability to attack more complex combinatorial problems.

In the last years some experiments have been done to extend to parallel execution search methods such as Constraint Programming or SAT solving (Boolean satisfiability), and combinatorial optimization methods such as Local Search, Meta-heuristics and Brand & Bound. However these works have mostly been done for shared memory multi-core systems (i.e. with a few cores) or for small PC clusters (a few machines).

The next challenge is to devise efficient techniques and algorithms for massively parallel computers with tens or hundreds of thousands of cores in the form of heterogeneous hybrid systems based on both multi-core processors and GPUs.

We would like to provide a cross-community forum for researchers working on search methods (Constraint Solving, Artificial Intelligence, Logic Programming, SAT solving, etc.), combinatorial optimization methods (metaheuristics, local search, tabu search, evolutionary algorithms, ant colony optimization, particle swarm optimization, memetic algorithms, and other types of algorithms) and High Performance Computing (Grids, large PC clusters, massively parallel computers, GPGPUs) in order to tackle the challenge of efficient implementations on all kinds of parallel hardware: multi-core, GPU-based or heterogeneous massively parallel systems.

This meeting is designed to be a forum for researchers willing to tackle those issues, in order to exchange ideas, theoretical frameworks, design of algorithms and methods, implementation issues, experimental results and further boost this growing area through cross-fertilization.

Overview of Talks

Local Search and Propagation on Hierarchical Multiprocessors (Long)

Salvador Abreu, University of Evora and CENTRIA

Hierarchical Multiprocessors are made up of multiple layers of parallel hardware: multithreading to exploit long pipelines, heterogeneous or homogeneous multicore chips, low-latency system area network clusters, higher latency grids or cloud computing. We discuss some issues related to this diversity and possible approaches on how to address the different requirements of these different parallel architectures in a unified software system.

High-performance Distributed Computing for Optimization Problems (Short)

Kento Aida, National Institute of Informatics

An optimization problem is currently used for modeling many scientific and industrial applications. The real application problem consists of huge computation, and distributed computing systems contribute to solve the problem faster. However, the architecture of the distributed computing systems is complex, and it is not an easy task for application researchers to fully utilize performance of the systems. This talk presents experience to implement applications of optimization problems on distributed computing systems. We implemented applications to solve combinatorial optimization problems on cluster and grid computing environment. The talk presents the results focusing on high-performance computing techniques to utilize distributed computing resources.

Parallel Local Search for SAT (Long)

Alejandro Arbelaez, University of Tokyo / JFLI

Parallel portfolio-based algorithms have become a standard methodology for both complete and incomplete solvers for SAT solving. In this methodology several algorithms explore the search space in parallel, either independently or cooperatively with some communication between the solvers. In this talk, we will study the performance of current state-of-the-art parallel local search solvers for SAT with and without cooperation considering a large degree of parallelism. In the cooperative mechanism solvers exchange the best assignment for the variables found so far with all members of the portfolio to properly craft a new assignment for the variables to restart from.

On Parameterized Multiroute Flows: A tool to manage failures in a network (Long)

Jean-François Baffier, The University of Tokyo - Imai Laboratory/JFLI and Université Paris-Sud - LRI (CNRS)

A K-route channel is a communication path that goes through K edgedisjoint paths, providing a resistance up to $K-1$ edge failures. A K-route-flow (a multiroute flow of K routes) is a sum of K-route channels respecting the network capacities constraint . It provides a lower bound of the max-flow value in a case of $K-1$ edges failures. A max-flow/min-cut theorem in the K-route context has been proved by Kishimoto and Takeuchi. This results can be easily extended to include node failures.

Given a network with one variable edge, we study the impact of the variation of this edge's capacity on the multiroute flow in a general case (previous work already treated the case of 1-route-flow and 2-route-flow). For any integer K, we show that a max-K-flow solution is a piecewise linear function (with at most $K+1$ changes of slope that we call critical points). We provide a polynomial algorithm which gives us all the critical points for a given source/sink pair. We also extend the function analysis to the case with any number of variable edges.

Parallel Constraint-Based Local Search (Short)

Yves Caniou, Université de Lyon, CNRS, INRIA

In this talk, I will present some work done on the parallelization of a constraint based local search algorithm called Adaptive Search.

Wave Search (Long)

Philippe Codognet, JFLI

We will present a novel model for parallel local search named "Wave Search". This is a general model aimed at providing the general control mechanism for guiding the search performed by several cooperating and competitive local search solvers. Indeed as we consider the implementation of local search on massively parallel computers, the issue of controlling the search became utmost important, in particular in the so-called portfolio approach where competing solvers are attacking the same problem and can communicate by means of exchanging their best configurations.

How the declarative multi-threaded implementation of some combinatorial algorithms turned into a garbage collection nightmare (Short)

Bart Demoen, Dept of Computer Science, KU Leuven

We recently started working on parallellizing certain graph algorithms using the multi-threading functionality in hProlog. Efficiency demands that different threads share data instead of communicating it to each other. This lead naturally to a memory model where each thread can access the "local" heap of all other threads, hence the nightmare for garbage collection. Lessons can be learned from other contexts. We try to cover the essence of the different issues.

Resource Allocation for Memory Intensive Parallel Search Algorithms (Long)

Alex Fukunaga, The University of Tokyo

The increasing availability of "utility computing" resources such as clouds, grids, and massively parallel shared clusters can provide practically unlimited processing and memory capacity on demand, at some cost per unit of resource usage. This requires a new perspective in the design and evaluation of parallel search algorithms. Previous work in parallel search implicitly assumed ownership of a cluster with a static amount of CPU cores and RAM, and emphasized wall-clock runtime. With utility computing resources, tradeoffs between performance and monetary costs must be considered. I will discuss some recent theoretical and experimental results on tradeoffs between performance and cost for a class of memory-intensive, parallel search algorithms, including parallel A^* (best-first branch-and-bound).

Combinatorial Optimization using Binary Decision Diagrams (Short)

Hiroyuki Higuchi, FUJITSU LABORATORIES LTD.

Several methods for solving combinatorial optimization problem by using binary decision diagrams are discussed.

Exploring the Power of Soft Constraints (Short)

Hiroshi Hosobe, National Institute of Informatics

Soft constraints are useful for treating over-constrained problems that naturally arise in real-life applications. We have been working on frameworks and algorithms for soft constraints to explore their power. In this talk, we present some of our results on constraint hierarchies, semiring-based CSPs, and probabilistic CSPs.

CSPSAT and Beyond (Short)

Katsumi Inoue, National Institute of Informatics

Recent SAT technologies have been successfully applied to various domains that require high performance reasoning. The CSPSAT project has been conducted by Prof. Naoyuki Tamura from Kobe University, and has attracted several researchers from Japan and other countries since 2008. In the former project, SAT technologies have been applied to constraint satisfaction problems (CSPs), and several award-winning CSP and SAT solvers have been developed by the members. Since 2012, the new project, CSPSAT2, has started to advance the existing SAT-based constraint solving methods in various ways.

In this talk, we will review CSPSAT and CSPSAT2 projects, and discuss how parallelism and distributed reasoning are related with the projects.

Symbolic Execution for Combinatorial Problems (Long)

Joxan Jaffar, National University of Singapore

We will overview the state-of-the-art in symbolic execution, but with emphasis on programs that model combinatorial problems. Symbolic execution can thus be seen as the problem of efficiently searching the space of program traces in search of "solutions" or good solutions. We will then describe some basic optimization methods based on the notions of interpolation and dynamic programming. Finally, we will briefly speculate how the entire methodology can benefit from parallelism.

Two open problems in Parallel Consistency (Long)

Arnaud Lallouet, University of Caen

In his 1990 Artificial Intelligence paper, Simon Kasif proved that arc-consistency is P-complete, thus unlikely to benefit from parallel computation. However, he did not suggest that practical improvements could not be made in the same way NP-complete problems are practically solved. Nevertheless, this field has been a limited subject of research since then. One of the reason is that consistency generates a too fine grained parallelism which is difficult to convert in speedup in a distributed environment.

The purpose of this talk is to make a survey of this field and to introduce two (yet unsolved) open problems on parallel consistency that belong to my research agenda. One is theoretical and the other is practical. It could be interesting also to check whether 20 years later, the increasing size of constraint problem could finally benefit from parallelism in consistency.

Can GPU Speed Up Search in Constraint Satisfaction (Short)

Jimmy Lee, The Chinese University of Hong Kong

This talk shares some of our experience in using GPU computation for constraint satisfaction. We first outline modern GPU architecture, and highlight the memory bottlenecks. Experiments on local search are described. We then explore opportunities on system search.

Parallel Local Search: experiments with a different programming model (Short)

Rui Machado, Fraunhofer ITWM

Local search is a successful approach for solving combinatorial optimization and constraint satisfaction problems. With the progressing move toward multi and many-core systems, GPUs and the quest for Exascale systems, parallelism has become mainstream as the number of cores will continue to increase. New programming models are required, which need to be better understood. Such is the case for local search algorithms when run on hundreds or thousands of processing units. In this talk, we discuss some experiments we've been doing and issues which need to be addressed.

Optimization challenges in the implementation of deployment strategies for protected long-reach PON (Long)

Luis Quesada, Cork Constraint Computation Centre

The mass deployment of fibre access networks is probably the most important network upgrade strategy for operators over the coming decade. Next generation networks, and in particular the Long-Reach Passive Optical Network (LR-PON) solution, aim to increase long term economic viability and sustainability of Fibre-To-The-Premises (FTTP) deployment. The LR-PON solution achieves this by minimizing the number of nodes and the amount of electronic equipment required within the network.

In this talk I will elaborate on two challenging optimization problems that we face when implementing deployment strategies for protected long-reach PON:

(a) How to place metro nodes and exchange sites in such a way that the amount of cable between metro nodes and exchange sites and between exchange sites and customers is minimized, taking into account that (ideally) all customers need to be covered.

(b) Assuming that metro nodes have been associated with exchange sites and exchange sites have been associated with customers, how to route the cable from customers to exchange sites, and from exchange sites to metro nodes in such a way that cable is minimized, the upper bound on the length of the path is respected and reliability is insured.

I will present some of the results that we have obtained following sequential approaches and elaborate on some opportunities that we would like to explore to improve the scalability of our approaches using parallel computation.

Parallel local search for the Costas Array Problem (Long)

Florian Richoux, JFLI, CNRS - University of Tokyo

The Costas Array Problem is a highly combinatorial problem linked to radar applications. We present in this talk its modeling and solving by Adaptive Search, a constraint-based local search method. Experiments have been done on both sequential and parallel hardware up to several hundreds of cores. Performance evaluation of the sequential version shows results outperforming previous implementations, while the parallel version shows nearly linear speedups up to 8,192 cores.

SatX10: An architecture for Capability Constraint Solving in X10 (Long)

Vijay Saraswat, IBM TJ Watson Research Center

Increasing technological advances in multi-core processing, cluster technology and cloud computing mean that the future of advanced computing lies in harnessing hundreds and thousands of cores.

We focus on the problem of capability constraint solving: By using more cores, can we solve problems faster, hence more problems in the same time bound? We believe this is the central question to answer when looking to integrate constraint solving into interactive applications. (We believe that a metric such as speedup is not of much use for SAT because increased communication can fundamentally alter the search space, leading often to non-linear speedups.)

We present SatX10, a communicating, portfolio-based SAT solver that can incorporate many existing state of the art (C++) sequential solvers and can run thousands of them in parallel, using a hub written in the X10 programming language (http://x10-lang.org).

We present results on running SatX10 on problems from the SAT 2011 competition on several hundred cores (upto 2K cores). We show dramatic reductions in run-time for some problems,from thousands of seconds to dozens. We show SatX10 at 1024 cores can solve some problems in hundreds of seconds that were not solvable at all by solvers in the SAT 2011 competition in thousands of seconds.

(Joint work with Ashish Sabhawal, Horst Meinholf, Ben Herta and David Grove of the IBM TJ Watson Research Lab.)

Speculative Computation (Short)

Ken Satoh, National Institute of Informatics

We describe a short introduction of our research on speculative computation which is one of the main mechanism of "kiga-kiku" (which is a Japanese word meaning "doing an appropriate action by predicting other people's intention with small communication") computer.

Parallel algebraic computations for combinatorial problems (Long)

Yosuke Sato, Tokyo University of Science

A combinatorial problem is often describable as a system of Boolean equations. For examples, a four-coloring problem of a graph with n-many nodes can be described with $4n$-many variables, a 9×9 Sudoku puzzle can be described with 729-many variables. Since most of such problems are NP-hard or NP-complete, it is of great importance to reduce the number of variables.

Boolean equations of the above examples are equations over the simplest Boolean ring that is the Galois field GF(2). In [1], we showed that those equations are further translated into more general Boolean equations. For example, a Sudoku puzzle is describable in terms of Boolean equations over the Boolean ring that consists of the power set P(1, ..., 9) of a finite set with 9 elements 1, ..., 9. For such Boolean equations, we need only 81 variables. In [1], we also introduced an efficient algorithm based on Boolean Groebner bases computation([4]) to solve singleton set constraints. Singleton set constraints are a certain type of combinatorial problems which includes the above two examples.

Using the implementation of Boolean Groebner bases([2]) on the computer algebra system Risa/Asir([3]), we developed a parallel program for solving singleton set constraints using the facility OpenXM([5]). Especially, wedeveloped several programs for solving and analyzing Sudoku puzzles. In general, it is not so easy to develop a parallel program for the computation of Groebner bases in a polynomial ring over any field. The simplest Boolean ring GF(2) is not an exception. Our approach is something different. If we want to solve a 9×9 Sudoku puzzle by the naive method, we have to compute a Groebner basis in a polynomial ring over GF(2) with 729 variables. It is not an easy job to compute a Groebner basis for such a huge system even by the fastest Groebner bases computation program in present-day. Instead of working in a polynomial ring over GF(2) with 729 variables, we work in a polynomial ring over P(1, ..., 9) with 81 variables. The algorithm of [1] consists of several independent computations of Groebner bases in it. The implementation of [2] further divides it into the computations of Groebner bases in a polynomial ring over GF(2) with 81 variables. As a result, it is something like that we can transform a Groebner basis computation in a polynomial ring over GF(2) with 729 variables into several independent Groebner bases computations in a polynomial ring over GF(2) with 81 variables. Our implementation not only can solve any Sudoku puzzle but also gives a mathematical criterion for its difficulty. Though the computation time for the first one is within a few seconds for any puzzle, we need much more time for the second one if the puzzle is in higher criterion. Our parallel implementation works well on such a heavy computation. It achieves an almost ideal speedup.

[1] Inoue, S. (2012). Efficient Singleton Set Constraint Solving by Boolean Groebner Bases. Communications of JSSAC, Vol. 1, pp 27-37

[2] Inoue, S. and Nagai, A.(2009). On the implementation of Boolean Groebner bases. The Joint Conference of ASCM 2009 and MACIS 2009, pp 58-62, COE Lect. Note, 22, Kyushu Univ. Fac. Math., Fukuoka, 2009.

[3] Noro, M. et al. (2012). A Computer Algebra System Risa/Asir. http://www.math.kobe-u.ac.jp/Asir/asir.html.

[4] Sato, Y., Inoue, S., Suzuki, A., Nabeshima, K. and Sakai, K.(2011). Boolean Groebner Bases. J. Symbolic Comput. 46 (2011), no. 5, pp 622-632.

[5] Takayama, N. et al. (2012). OpenXM (Open message eXchange for Mathematics) http://www.math.kobe-u.ac.jp/OpenXM/

Parallel Restarted Search and Parallel Dichotomic Search (Long)

Meinolf Sellmann, IBM Research

We consider the problem of parallelizing search in constraint programming (CP). With few exceptions, most commercial and academic CP solvers do not learn no-goods during search. For this case, we develop a simple technique for parallelizing restarted search deterministically and demonstrate experimentally that we can achieve near-linear speedups in practice. We then consider the problem of parallelizing skewed dichotomic search and develop worst-case optimal parallel dichotomic search protocols.

Applying Incremental SAT Solving to Optimization and Enumeration Problems (Short)

Takehide Soh, Kobe University

In this talk, we show an approach to optimization and enumeration problems with technologies for solving propositional satisfiability (SAT). Although SAT technologies were mainly focused on solving decision problems, their recent progress is significant so that methods based on SAT technologies can be potential approaches. In particular, we review a method called incremental SAT solving. It incrementally computes the satisfiability of sub-problems to solve a given original problem. We then show experimental results of applying incremental SAT solving to the two-dimensional strip packing problem and the minimal active pathway finding problem.

Search Combinators for Parallel Search (Long)

Guido Tack, Monash University

The ability to model search in a constraint solver can be an essential asset for solving combinatorial problems. We have recently presented Search Combinators, a lightweight, solver-independent language for expressing search heuristics in a high-level, modular way. The language provides a small set of primitives, which can be combined arbitrarily to express application-tailored heuristics such as branch-and-bound search, restarting strategies, probing search, and many others. The modularity of search combinators on the language level is reflected in a completely modular implementation. In this talk, I will briefly introduce search combinators and then present the opportunities and challenges for parallelisation that arise from this technique.

A compact and efficient SAT encoding of CSP and its evaluation (Short)

Tomoya Tanjo, Transdisciplinary Research Integration Center

In this talk, we describe a new SAT encoding method named compact order encoding. In the compact order encoding, each integer variable is represented by using a numeral system of base $B \geq 2$ and each digit is encoded by using the order encoding. We will show some experimental results of the compact order encoding through some experimental results on some benchmarks such as Open-Shop Scheduling problems with large domain sizes.

Probabilistic Study of Random Restart (Long)

Charlotte Truchet, Université de Nantes

In Local Search (LS), Random Restart (RR) is a widely used technique. For a given number of iterations Max_Iter, it consists in restarting search from scratch every Max_Iter iterations, as long as a solution has not been found. Intuitively, random restart allows the solver to visit a more important part of the search space. We will present a probabilistic model for this mechanism. In this model, we can compute the expectancies of a given LS algorithm solving time, with and without RR. We will then show how to extend the same model to study the mean gain for parallelized LS algorithm.

Development of the LMNtal Parallel Model Checker (Long)

Kazunori Ueda, Waseda University

We have designed and implemented LMNtal, a language based on hierarchical graph rewriting that allows us to encode diverse computational models involving concurrency, mobility and multiset rewriting. The implementation has recently evolved into a model checker that employs LMNtal as the modeling language and PLTL as the specification language. The strengths of our LMNtal model checker include its powerful data structure with an inherent symmetry reduction mechanism, highly nondeterministic computation it can express, and virtually no discrepancy between programming and modeling languages. Since model checking of graph rewriting models calls for time- and space-efficient methods for state management, we have further parallelized our model checker to run on shared-memory multiprocessors and developed various optimization techniques of state management. Those enhancements achieved high scalability and improved both the size of problems the model checker can handle as well as its performance.

Between Optimization and Enumeration (Long)

Takeaki Uno, National Institute of Informatics

Recent “Big Data” increases the necessity of fast computation. However, we would be seeing some limitation of usual approach, especially for enumeration and optimization. In this talk, we will discuss about new research direction; algorithmic modeling. In algorithmic modeling, algorithms are designed for increasing the possibility models that allow fast computation, and models are designed with having first priority on the fast computation. We first discuss the difference between the optimization and enumeration, and clarify why they are difficult to allow fast computation. Then we consider new models locating between optimization and enumeration that allow fast computation, by considering/developing fast algorithms that we can use.

Constraints and Big Data? (Long)

Roland Yap, School of Computing, National University of Singapore

We are now in the age of “Big Data”. The veritable data tsunami gives rise to new challenges in the analysis of very large amounts of data. Cloud computing is also often tightly coupled with Big Data since the data size requires access to scalable compute/storage resources.

The question is whether is Constraint Programming (CP) relevant to Big Data? This talk starts by describing how maximum flow on very large graphs can be computed using the MapReduce framework. After that we will discuss what kind of Big Data problems might exist in a CP context as well as whether constraints can be applied to analysis of Big Data problems.

Participants

Salvador Abreu, University of Evora and CENTRIA
Kento Aida, National Institute of Informatics
Alejandro Arbelaez, University of Tokyo / JFLI
Jean-François Baffier, The University of Tokyo - Imai Laboratory/JFLI and Université Paris-Sud - LRI (CNRS)
Yves Caniou, Université de Lyon, CNRS, INRIA
Philippe Codognet, JFLI
Bart Demoen, Dept of Computer Science, KU Leuven
Alex Fukunaga, The University of Tokyo
Hiroyuki Higuchi, FUJITSU LABORATORIES LTD.
Hiroshi Hosobe, National Institute of Informatics
Katsumi Inoue, National Institute of Informatics
Joxan Jaffar, National University of Singapore
Arnaud Lallouet, University of Caen
Jimmy Lee, The Chinese University of Hong Kong
Rui Machado, Fraunhofer ITWM
Luis Quesada, Cork Constraint Computation Centre
Florian Richoux, JFLI, CNRS - University of Tokyo
Vijay Saraswat, IBM TJ Watson Research Center
Ken Satoh, National Institute of Informatics
Yosuke Sato, Tokyo University of Science
Meinolf Sellmann, IBM Research
Takehide Soh, Kobe University
Guido Tack, Monash University
Tomoya Tanjo, Transdisciplinary Research Integration Center
Charlotte Truchet, Université de Nantes
Kazunori Ueda, Waseda University
Takeaki Uno, National Institute of Informatics
Roland Yap, School of Computing, National University of Singapore