

Nonmonotonic and Paraconsistent Reasoning: From Basic Entailments to Plausible Relations

Ofer Arieli[†] and Arnon Avron[‡]

Abstract

In this paper we develop frameworks for logical systems which are able to reflect not only nonmonotonic patterns of reasoning, but also paraconsistent reasoning. For this we consider a sequence of generalizations of the pioneering works of Gabbay, Kraus, Lehmann, Magidor and Makinson. Our sequence of frameworks culminates in what we call plausible, nonmonotonic, multiple-conclusion consequence relations (which are based on a given monotonic one). Our study yields intuitive justifications for conditions that have been proposed in previous frameworks, and also clarifies the connections among some of these systems. In addition, we present a general method for constructing plausible nonmonotonic relations. This method is based on a multiple-valued semantics, and on Shoham's idea of preferential models.

Presented at the 5th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ecsquaru'99). London, UK, July 5-9 1999. Lecture Notes in Computer Science No. 1638, pages 11-22, Springer, 1999.

[†]*Supported by the Deutsch Institute*

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Sparse Universal Graphs for General Max-degree 3 Graphs

Noga Alon[†] Vera Asodi[‡]

Abstract

For every n , we describe an explicit construction of a graph on n vertices with at most $O(n^{2-\epsilon})$ edges, for $\epsilon = 0.133\dots$, that contains every graph on n vertices with maximum degree 3 as a subgraph. The construction is explicit, but the proof of its properties is based on probabilistic arguments. It is easy to see that each such graph has $\Omega(n^{\frac{4}{3}})$ edges. The study of this problem is motivated by questions in VLSI circuit design.

Vera Asodi was awarded the Deutsch Prize for the year 2000 in recognition of her research.

[†]Department of Mathematics, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel. Research supported in part by a USA-Israeli BSF grant, by the Israel Science Foundation and by the Hermann Minkowski Minerva Center for Geometry at Tel Aviv University. Email: noga@math.tau.ac.il.

[‡] Department of Computer Science, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel. Email: vera@math.tau.ac.il. Research supported by the Deutsch Institute.

Individual Sequence Prediction - Upper Bounds and Application for Complexity

Chamy Allenberg^{†‡}

Abstract

The loss version of the multi-armed bandit problem is carried out in T iterations. At the beginning of any iteration an adversary assigns losses from $[0, 1]$ to each of the K options (also called arms). Then, without knowing the adversary's assignments, we are required to select one out of the K arms, and suffer the loss that was assigned to it.

Here we consider the *loss game* which is the *adversarial* version of the loss version of the multi-armed bandit problem. In this version *no* stochastic assumption is made, thus the results hold for any possible assignment of $K \times T$ losses.

We compete against L_{opt} , the *optimal loss*, which is the minimal total loss of any consistent choice of an arm in this game, i.e., the performance of the best arm. Our goal is to minimize the *regret*, the maximization over all possible assignments of losses, of the difference between our expected total loss and L_{opt} .

In a previous work Auer, Cesa-Bianchi, Freund and Schapire showed that the *regret* in the loss game has an upper bound of $O(T^{1/2})$ and a lower bound of $\Omega(L_{\text{opt}}^{1/2})$. Since the losses in the loss game are normalized to the $[0,1]$ range, a loss of 1 is the upper bound on the loss possible in any one iteration. Thus T , the number of iterations, can even be higher than the total loss of the worst consistent choice of an arm (i.e., the performance of the worst arm).

In this work an upper bound of $O(L_{\text{opt}}^{2/3})$ on the *regret* is presented.

Presented at the Twelfth Annual Conference on Computational Learning Theory 1999 California

[†] Supported by the Deutsch Institute

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Segment Intersection Searching Problems in General Settings

Vladlen Koltun^{†§}

Abstract

We consider segment intersection searching amidst (possibly intersecting) algebraic arcs in the plane. We show how to preprocess n arcs in time $O(n^{2+\varepsilon})$ into a data structure of size $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, such that the k arcs intersecting a query segment can be counted in time $O(\log n)$ or reported in time $O(\log n + k)$. This problem was extensively studied in restricted settings (e.g., amidst segments, circles or circular arcs), but no solution with comparable performance was previously presented for the general case of possibly intersecting algebraic arcs. Our data structure for the general case matches or improves (sometimes by an order of magnitude) the size of the best previously presented solutions for the special cases.

As an immediate application of this result, we obtain an efficient data structure for the triangular windowing problem, which is a generalization of triangular range searching. As another application, the first substantially sub-quadratic algorithm for a red-blue intersection counting problem is derived. We also describe simple data structures for segment intersection searching among disjoint arcs, and ray shooting among algebraic arcs.

Presented at the 17th ACM Symposium on Computational Geometry, Boston, USA, June 3-5, 2001.
Vladlen Koltun was awarded the Deutsch Prize for the year 2001 in recognition of his research.

**Supported by the Deutsch Institute*

[†]School of Computer Science, Tel Aviv University, Tel Aviv, Israel

An Improved Lower Bound for Approximating CVP

Irit Dinur^{†‡} Guy Kindler^{§¶} Shmuel Safra^{||}

Abstract

This paper shows the problem of finding the closest vector in an n -dimensional lattice to be NP-hard to approximate to within factor $n^{c/\log \log n}$ for some constant $c > 0$.

Presented at the 39th IEEE Symposium on Foundations of Computer Science, Palo Alto, California, USA, November 8-11, 1998.

[†]Tel-Aviv University, Israel

[‡]*Supported by the Deutsch Institute*

[§]Tel-Aviv University, Israel

[¶]*Supported by the Deutsch Institute*

^{||}Tel-Aviv University, Israel

PCP Characterizations of NP: Towards a Polynomially-Small Error-Probability

Irit Dinur^{†‡} Eldar Fischer[§] Guy Kindler^{¶||} Ran Raz^{**} Shmuel Safra^{††}

Abstract

This paper strengthens the low-error PCP characterization of NP, coming closer to the upper limit of the BGLR conjecture. Namely, we prove that witnesses for membership in any NP language can be verified with a constant number of accesses, and with an error probability exponentially small in the number of bits accessed, where this number is as high as $\log^\beta n$, for *any* constant $\beta < 1$. (The BGLR conjecture claims the same for any $\beta \leq 1$).

Our results are in fact stronger, implying the Gap- Quadratic-Solvability problem to be NP-hard even if the equations are restricted to having a constant number of variables. That is, given a system of quadratic-equations over a field \mathcal{F} (of size up to $2^{\log^\beta n}$), where each equation depends on a constant number of variables, it is NP-hard to decide between the case where there is a common solution for all of the equations, and the case where any assignment satisfies no more than a $\frac{2}{|\mathcal{F}|}$ fraction of them.

At the same time, our proof presents a *direct* construction of a low-degree-test whose error-probability is exponentially small in the number of bits accessed. Such a result was previously known only relying on recursive applications of the entire PCP theorem.

Presented at the Thirty-First Annual ACM Symposium on Theory of Computing, May 1-4, 1999.

[†]Tel-Aviv University, Israel

[‡]*Supported by the Deutsch Institute*

[§]Tel-Aviv University, Israel

[¶]Tel-Aviv University, Israel

^{||}*Supported by the Deutsch Institute*

^{**}Weizmann Institute, Israel

^{††}Tel-Aviv University, Israel

Predicting User Intentions In Graphical User Interfaces Using Implicit Disambiguation

David Noy^{†‡}

Abstract

We address the problem of predicting user intentions in cases of pointing ambiguities in graphical user interfaces. We argue that it is possible to heuristically resolve pointing ambiguities using implicit information that resides in natural pointing gestures, thus eliminating the need for explicit interaction methods and encouraging natural human-computer interaction. We present two speed-accuracy measures for predicting the size of the intended target object. These two measures are tested empirically and shown to be valid and robust. Additionally, we demonstrate the use of exact mouse location for disambiguation and the use of estimated movement continuation for predicting intended target objects at early stages of the pointing gesture.

Presented at CHI2001, Seattle, USA, March 2001.

[†]Supported by the Deutsch Institute

[‡]Tel-Aviv University, School of Computer Science

Lower Bounds for On-line Scheduling with Precedence Constraints on Identical Machines

Leah Epstein^{†‡}

Abstract

We consider the on-line scheduling problem of jobs with precedence constraints on m parallel identical machines. Each job has a time processing requirement, and may depend on other jobs (has to be processed after them). A job arrives only after its predecessors have been completed. The cost of an algorithm is the time that the last job is completed. We show lower bounds on the competitive ratio of on-line algorithms for this problem in several versions.

We prove a lower bound of $2 - 1/m$ on the competitive ratio of any deterministic algorithm (with or without preemption) and a lower bound of $2 - 2/(m+1)$ on the competitive ratio of any randomized algorithm (with or without preemption).

The lower bounds for the cases that preemption is allowed require arbitrarily long sequences. If we use only sequences of length $O(m^2)$, we can show a lower bound of $2 - 2/(m+1)$ on the competitive ratio of deterministic algorithms with preemption, and a lower bound of $2 - O(1/m)$ on the competitive ratio of any randomized algorithm with preemption. All the lower bounds hold even for sequences of unit jobs only.

The best algorithm that is known for this problem is the well known List Scheduling algorithm of Graham. The algorithm is deterministic and does not use preemption. The competitive ratio of this algorithm is $2 - 1/m$. Our randomized lower bounds are very close to this bound (a difference of $O(1/m)$) and our deterministic lower bounds match this bound.

Presented at the 1st Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX98), Aalborg, Denmark, July 1998.

[†] Supported by the Deutsch Institute

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

On-line Machine Covering

Yossi Azar[†] Leah Epstein^{‡§}

Abstract

We consider the problem of scheduling a sequence of jobs on m parallel identical machines so as to maximize the minimum load over the machines. This situation corresponds to a case that a system which consists of the m machines is alive (i.e. productive) only when all the machines are alive, and the system should be maintained alive as long as possible. It is well known that any on-line deterministic algorithm for identical machines has a competitive ratio of at least m and that greedy is an m competitive algorithm. In contrast we design an on-line randomized algorithm which is $O(\sqrt{m} \log m)$ competitive and a lower bound of $\Omega(\sqrt{m})$ for any on-line randomized algorithm. In the case where the weights of the jobs are polynomially related we design an optimal $O(\log m)$ competitive randomized algorithm and a matching tight lower bound for any on-line randomized algorithm. In fact, if F is the ratio between the weights of largest job and the smallest job then our randomized algorithm is $O(\log F)$ competitive.

A sub-problem that we solve which is interesting in its own right is the problem where the value of the optimal algorithm is known in advance. Here we show a deterministic (constant) $2 - \frac{1}{m}$ competitive algorithm. We also show that our algorithm is optimal for two, three and four machines and that no on-line deterministic algorithm can achieve a better competitive ratio than 1.75 for $m \geq 4$ machines.

Presented at the 5th Annual European Symposium on Algorithms, Graz, Austria, September 1997.

[†]Dept. of Computer Science,

Tel Aviv University, Tel Aviv, Israel

[‡] Supported by the Deutsch Institute

[§]Dept. of Computer Science,

Tel Aviv University, Tel Aviv, Israel

Learning Rates for Q-Learning

Eyal Even-dar^{†‡} Yishay Mansour[§]

Abstract

In this paper we derive convergence rates for Q-learning. We show an interesting relationship between the convergence rate and the learning rate used in the Q-learning. For a polynomial learning rate, one which is $1/t^\omega$ at time t where $\omega \in (1/2, 1)$, we show that the convergence rate is polynomial in $1/(1 - \gamma)$, where γ is the discount factor. In contrast we show that for a linear learning rate, one which is $1/t$ at time t , the convergence rate has an exponential dependence on $1/(1 - \gamma)$. In addition we show a simple example that proves that this exponential behavior is inherent for a linear learning rate.

Presented at the 14th Annual Conference on Computational Learning Theory, Amsterdam, the Nedarlands, July 16-19, 2001.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

The Design and Implementation of Planar Maps in CGAL

Eyal Flato[†] Dan Halperin Iddo Hanniel Oren Nechushtan
Department of Computer Science
Tel Aviv University, Tel Aviv 69978, Israel

Abstract

Planar maps are fundamental structures in computational geometry. They are used to represent the subdivision of the plane into regions and have numerous applications. We describe the planar map package of CGAL — the Computational Geometry Algorithms Library. We discuss problems that arose in the design and implementation of the package and report the solutions we have found for them. In particular we introduce the two main classes of the design—*planar maps* and *topological maps* that enable the convenient separation between geometry and topology. We also describe the *geometric traits* which make our package flexible by enabling to use it with any family of curves as long as the user supplies a small set of operations for the family. Finally, we present the algorithms we implemented for point location in the map, together with experimental results that compare their performance.

Presented by Eyal Flato at the 3rd International Workshop on Algorithm Engineering , London, UK, July 19-21, 1999.

[†]Department of Computer Science, Tel Aviv University, supported by the Deutsch Institute.

Deep Compression for Streaming Texture Intensive Animations

Daniel Cohen-Or[†] Yair Mann[‡] Shachar Fleishman[§]

Abstract

This paper presents a streaming technique for synthetic texture intensive 3D animation sequences. There is a short latency time while downloading the animation, until an initial fraction of the compressed data is read by the client. As the animation is played, the remainder of the data is streamed online seamlessly to the client. The technique exploits frame-to-frame coherence for transmitting geometric and texture streams. Instead of using the original textures of the model, the texture stream consists of view-dependent textures which are generated by rendering offline nearby views. These textures have a strong temporal coherency and can thus be well compressed. As a consequence, the bandwidth of the stream of the view-dependent textures is narrow enough to be transmitted together with the geometry stream over a low bandwidth network. These two streams maintain a small online cache of geometry and view-dependent textures from which the client renders the walkthrough sequence in real-time. The overall data transmitted over the network is an order of magnitude smaller than an MPEG post-rendered sequence with an equivalent image quality.

Presented at SIGGRAPH 99' August 8-13, 1999, in Los Angeles, California.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]Webglide Ltd.

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel, supported by the Deutsch Institute

Multiple Structural Alignment and Core Detection by Geometric Hashing

Nathaniel Leibowitz[†] Zipora Y. Fligelman[‡] Ruth Nussinov^{§¶} Haim J. Wolfson^{||}

Abstract

A Multiple Structural Alignment algorithm is presented in this paper. The algorithm accepts an ensemble of protein structures and finds the largest substructure (core) of C_α atoms whose geometric configuration appear in all the molecules of the ensemble (core). Both the detection of this core and the resulting structural alignment are done simultaneously. Other large enough multi-structural superimpositions are detected as well. Our method is based on the Geometric Hashing paradigm and a superimposition clustering technique which represents superimpositions by sets of matching atoms. The algorithm proved to be efficient on real data in a series of experiments. The same method can be applied to any ensemble of molecules (not necessarily proteins) since our basic technique is sequence order independent.

Keywords : Multiple structural alignment; Geometric Hashing; invariants; structural core; transformation clustering.

Presented at the International Conference on Intelligent System for Molecular Biology, Heidleberg, Germany, August 6-10, 1999.

[†]*Supported by the Deutsch Institute*

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Loss-Bounded Analysis for Differentiated Services

Alexander Kesselman ^{†‡} Yishay Mansour [§]

Abstract

We consider a network providing Differentiated Services (Diffserv) which allow network service providers to offer different levels of Quality of Service (QoS) to different traffic streams. We focus on loss and first show that only trivial bounds could be obtained by means of traditional competitive analysis. Then we introduce a new approach for estimating loss of an online policy called *loss-bounded* analysis. In loss-bounded analysis the loss of an online policy are bounded by the loss of an optimal offline policy plus a constant fraction of the benefit of an optimal offline policy. We derive tight upper and lower bounds for various settings of Diffserv parameters using the new loss-bounded model. We believe that loss-bounded analysis is an important technique that may complement traditional competitive analysis and provide new insight and interesting results.

Presented at the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), January 7-9, 2001, Washington DC.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

A Formula-Preferential Base for Paraconsistent and Plausible Reasoning Systems

Arnon Avron[†] Iddo Lev^{†‡}

Abstract

We provide a general framework for constructing natural consequence relations for paraconsistent and plausible nonmonotonic reasoning. The framework is based on preferential systems whose preferences are based on the satisfaction of formulas in models. We show that these natural preferential systems that were originally designed for paraconsistent reasoning satisfy a key condition (stopperedness or smoothness) from the theoretical research of nonmonotonic reasoning. Consequently, the nonmonotonic consequence relations that they induce satisfy the desired conditions of plausible consequence relations. Hence our framework encompasses different types of preferential systems that were developed from different motivations of paraconsistent reasoning and non-monotonic reasoning, and reveals an important link between them.

Presented at the 17th International Joint Conference on Artificial Intelligence, Seattle, Washington, USA, August 4–10, 2001.

[†]School of Computer Science, Tel-Aviv University

[‡]*Supported by the Deutsch Institute*

TVLA: A System for Implementing Static Analyses

Tal Lev-Ami^{†‡} Mooly Sagiv[§]

Abstract

We present TVLA (Three-Valued-Logic Analyzer). TVLA is a “YACC”-like framework for automatically constructing static-analysis algorithms from an operational semantics, where the operational semantics is specified using logical formulae. TVLA has been implemented in Java and was successfully used to perform shape analysis on programs manipulating linked data structures (singly and doubly linked lists), to prove safety properties of Mobile Ambients, and to verify the partial correctness of several sorting programs.

Presented at the International Static Analysis Symposium (SAS2000), Santa Barbara, USA, June 29–July 1, 2000.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University

Truly Online Paging with Locality of Reference

Amos Fiat[†] Manor Mendel^{‡§}

Abstract

The competitive analysis fails to model locality of reference in the online paging problem. To deal with it, Borodin *et. al.* introduced the access graph model for the paging problem, which attempts to capture the locality of reference. However, the access graph model has a number of troubling aspects. The access graph has to be known in advance to the paging algorithm and the memory required to represent the access graph itself may be very large.

In this paper we present truly online strongly competitive paging algorithms in the access graph model that do not have any prior information on the access sequence. We present both deterministic and randomized algorithms. The algorithms need only $O(k \log n)$ bits of memory, where k is the number of page slots available and n is the size of the virtual address space. *I.e.*, asymptotically no more memory than needed to store the virtual address translation table. In fact, it can be reduced to $O(k \log k)$ bits using appropriate probabilistic data structures.

Presented at the 38th Annual Symposium on Foundations of Computer Science, Miami Beach, FL, U.S.A. October, 1997.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Better Algorithms for Unfair Metrical Task Systems and Applications

Amos Fiat[†] Manor Mendel^{‡§}

Abstract

Unfair metrical task systems are a generalization of online metrical task systems. In this paper we introduce new techniques to combine algorithms for unfair metrical task systems and apply these techniques to obtain the following results:

1. Better randomized algorithms for unfair metrical task systems on the uniform metric space.
2. Better randomized algorithms for metrical task systems on general metric spaces, $O(\log^2 n(\log \log n)^2)$ competitive, improving on the best previous result of $O(\log^5 n \log \log n)$.
3. A tight randomized competitive ratio for the k -weighted caching problem on $k + 1$ points, $O(\log k)$, improving on the best previous result of $O(\log^2 k)$.

Presented at the thirty-second annual ACM Symposium on Theory of Computing, Portland, OR, USA, May, 2000

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

A Ramsey-type Theorem for Metric Spaces and its Applications for Metrical Task Systems and Related Problems

Yair Bartal[†] Béla Bollobás[‡] Manor Mendel^{§¶}

Abstract

This paper gives a nearly logarithmic lower bound on the randomized competitive ratio for the Metrical Task Systems model [BLS92]. This implies a similar lower bound for the extensively studied K -server problem. Our proof is based on proving a Ramsey-type theorem for metric spaces. In particular we prove that in every metric space there exists a large subspace which is approximately a “hierarchically well-separated tree” (HST) [Bar96]. This theorem may be of independent interest.

Presented at the The 42nd Annual Symposium on Foundations of Computer Science, Las Vegas, Nevada, USA, October 14-17, 2001.

[†]Hebrew University Jerusalem, Israel

[‡]The University of Memphis, Memphis, TN 38152

[§]*Supported by the Deutsch Institute*

[¶]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

An Efficient Data Structure for Feature Extraction in a Foveated Environment

Efri Nattel^{†‡} Yehezkel Yeshurun^{§¶}

Abstract

Foveated sampling and representation of images is a powerful tool for various vision applications. However, there are many inherent difficulties in implementing it. We present a simple and efficient mechanism to manipulate image analysis operators directly on the foveated image; A single typed table-based structure is used to represent various known operators. Using the Complex Log as our foveation method, we show how several operators such as edge detection and Hough transform could be efficiently computed almost at frame rate, and discuss the complexity of our approach.

Presented at the IEEE International Workshop on Biologically Motivated Computer Vision (BMCV2000), Seoul, Korea, May 15–17, 2000.

[†]Dept. of Computer Science, Tel-Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[¶]*Supported by the Minerva Minkowski center for Geometry, and by a grant from the Israel Academy of Science for geometric Computing*

Detecting Memory Errors via Static Pointer Analysis

Nurit Dor[†] Michael Rodeh[‡] Mooly Sagiv[§]

Abstract

Programs which manipulate pointers are hard to debug. Pointer analysis algorithms (originally aimed at optimizing compilers) may provide some remedy by identifying potential errors such as dereferencing `NULL` pointers by statically analyzing the behavior of programs on all their input data.

Our goal is to identify the “core program analysis techniques” that can be used when developing realistic tools which do not generate too many false alarms. It is an open question if exists a conservative technique that will yield only a modest number of false alarms, and if so, if it will scale for large programs. Our preliminary experience indicates that the following techniques are necessary: (i) finding aliases between pointers, (ii) flow sensitive techniques that account for the program control flow constructs, (iii) partial interpretation of conditional statements, (iv) analysis of the relationships between pointers, and sometimes (v) analysis of the underlying data structures manipulated by the C program.

We show that a combination of these techniques yields better results than those achieved by state of art tools.

Presented in the 1998 ACM SIGPLAN Workshop on Program Analysis for Software Tools and Engineering Montreal, CA, 14 June 1998

[†]*Supported by the Deutsch Institute*, Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]Faculty of Computer Science, The Technion, Haifa, Israel

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Advanced course on the Principles of Program Analysis

Nurit Dor^{†‡}

Abstract

The course gave an overview of a number of key approaches to program analysis, all of which have a quite extensive literature, and shows that there is a large amount of commonality among the approaches. This helps in increasing the ability to choose the right approach for the task at hand and in exploiting insights developed in one approach to enhance the power of other approaches.

More concretely, the course presented the foundations of the four approaches:

- Data Flow Analysis,
- Control Flow Analysis,
- Abstract Interpretation, and
- Type and Effect Systems

held in Schloss Dagstuhl, Germany on 9-13 November 1998. Lectures by:

- Flemming Nielson, Aarhus University, Denmark
- Hanne Riis Nielson, Aarhus University, Denmark
- Chris Hankin, Imperial College, UK

[†] Supported by the Deutsch Institute

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

On the Complexity of Positional Sequencing by Hybridization

Amir Ben-Dor[†] Itsik Pe'er^{‡§} Ron Shamir[‡] Roded Sharan[‡]

Abstract

In sequencing by hybridization (SBH), one has to reconstruct a sequence from its k -long substrings. SBH was proposed as a promising alternative to gel-based DNA sequencing approaches, but in its original form the method is not competitive. Positional SBH is a recently proposed enhancement of SBH in which one has additional information about the possible positions of each substring along the target sequence. We give a linear time algorithm for solving the positional SBH problem when each substring has at most two possible positions. On the other hand, we prove that the problem is NP-complete if each substring has at most three possible positions.

Presented at the 10th Annual Symposium on Combinatorial Pattern Matching, Warwick, UK, July 22–24, 1999.

[†]Department of Computer Science and Engineering, University of Washington, Washington, USA.

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[§]*Supported by the Deutsch Institute*

An Algorithm Combining Discrete and Continuous Methods for Optical Mapping

Richard M. Karp[†] Itsik Pe'er^{‡§} Ron Shamir[‡]

Abstract

Optical mapping is a novel technique for generating the restriction map of a DNA molecule by observing many single, partially digested copies of it, using fluorescence microscopy. The real-life problem is complicated by numerous factors: false positive and false negative cut observations, inaccurate location measurements, unknown orientations and faulty molecules. We present an algorithm for solving the real-life problem. The algorithm combines continuous optimization and combinatorial algorithms, applied to a non-uniform discretization of the data. We present encouraging results on real experimental data, and on simulated data.

Presented at the 7th international conference on Intelligent Systems for Molecular Biology, Heidelberg, Germany, August 6–10, 1999.

[†]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, California, USA.

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[§]*Supported by the Deutsch Institute*

Automatic Removal of Array Memory Leaks in Java

Ran Shaham^{†‡} Elliot K. Kolodner[§] Mooly Sagiv[¶]

Abstract

Current garbage collection (*GC*) techniques do not (and in general cannot) collect all the garbage that a program produces. This may lead to a performance slowdown and to programs running out of memory space.

In this paper, we present a practical algorithm for statically detecting memory leaks occurring in arrays of objects in a garbage collected environment. No previous algorithm exists. The algorithm is conservative, i.e., it never detects a leak on a piece of memory that is subsequently used by the program, although it may fail to identify some leaks. The presence of the detected leaks is exposed to the garbage collector, thus allowing GC to collect more storage.

We have instrumented the *Java* virtual machine to measure the effect of memory leaks in arrays. Our initial experiments indicate that this problem occurs in many Java applications. Our measurements of heap size show improvement on some example programs.

Presented at the 9th International Conference on Compiler Construction, Berlin, Germany, March/April 2000.

[†]*Supported by the Deutsch Institute*

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[§]IBM Haifa Research Laboratory, Haifa, Israel

[¶]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Minimizing the Flow Time without Migration

Baruch Awerbuch[†] Yossi Azar[‡] Stefano Leonardi[§] Oded Regev^{¶||}

Abstract

We consider the classical problem of scheduling jobs in a multiprocessor setting in order to minimize the flow time (total time in the system). The performance of the algorithm, both in offline and online settings, can be significantly improved if we allow preemption: i.e., interrupt a job and later continue its execution, perhaps migrating it to a different machine. Preemption is inherent to make a scheduling algorithm efficient. While in case of a single processor, most operating systems can easily handle preemptions, migrating a job to a different machine results in a huge overhead. Thus, it is not commonly used in most multiprocessor operating systems. The natural question is whether migration is an inherent component for an efficient scheduling algorithm, in either online or offline setting.

Leonardi and Raz (STOC'97) showed that the well known algorithm, *shortest remaining processing time* (SRPT), performs within a logarithmic factor of the optimal algorithm. Note that SRPT must use *both* preemption and migration to schedule the jobs. It is not known if better approximation factors can be reached. In fact, in the on-line setting, Leonardi and Raz showed that no algorithm can achieve a better bound.

Without migration, no (offline or online) approximations are known. This paper introduces a new algorithm that does not use migration, works *online*, and is just as effective (in terms of approximation ratio) as the best known *offline* algorithm (SRPT) that uses migration.

Presented at the Symposium on the Theory of Computing '99.

[†]Johns Hopkins University, Baltimore, MD 21218, and MIT Lab. for Computer Science

[‡]Department of Computer Science, Tel Aviv University, Israel

[§]Dipartimento di Informatica Sistemistica, Università di Roma “La Sapienza”, via Salaria 113, 00198-Roma, Italia

[¶]Department of Computer Science, Tel Aviv University, Israel

^{||}*Supported by the Deutsch Institute*

Strongly Polynomial Algorithms for the Unsplittable Flow Problem

Yossi Azar Oded Regev

Abstract

We provide the first strongly polynomial algorithms with the best approximation ratio for all three variants of the unsplittable flow problem (*UFP*). In this problem we are given a (possibly directed) capacitated graph with n vertices and m edges, and a set of terminal pairs each with its own demand and profit. The objective is to connect a subset of the terminal pairs each by a single flow path as to maximize the total profit of the satisfied terminal pairs subject to the capacity constraints. Classical *UFP*, in which demands must be lower than edge capacities, is known to have an $O(\sqrt{m})$ approximation algorithm. We provide the same result with a strongly polynomial combinatorial algorithm. The extended *UFP* case is when some demands might be higher than edge capacities. For that case we both improve the current best approximation ratio and use strongly polynomial algorithms. We also use a lower bound to show that the extended case is provably harder than the classical case. The last variant is the bounded *UFP* where demands are at most $\frac{1}{K}$ of the minimum edge capacity. Using strongly polynomial algorithms here as well, we improve the currently best known algorithms. Specifically, for $K = 2$ our results are better than the lower bound for classical *UFP* thereby separating the two problems.

Accepted to the Eighth Conference on Integer Programming and Combinatorial Optimization (IPCO VIII), Utrecht, The Netherlands, June 13-15, 2001

A Polynomial Approximation Algorithm for the Minimum Fill-In Problem

Assaf Natanzon[†] Ron Shamir[†] Roded Sharan^{†‡}

Abstract

In the *minimum fill-in* problem, one wishes to find a set of edges of smallest size, whose addition to a given graph will make it chordal. The problem has important applications in numerical algebra and has been studied intensively since the 1970s. We give the first polynomial approximation algorithm for the problem. Our algorithm constructs a triangulation whose size is at most eight times the optimum size squared. The algorithm builds on the recent parameterized algorithm of Kaplan, Shamir and Tarjan for the same problem.

For bounded degree graphs we give a polynomial approximation algorithm with a polylogarithmic approximation ratio. We also improve the parameterized algorithm.

Presented at the Thirtieth Annual ACM Symposium on Theory of Computing, Dallas, Texas, May 23–26, 1998.

[†]Department of Computer Science, Tel Aviv University, Tel Aviv, Israel.
{natanzon,shamir,roded}@math.tau.ac.il.

[‡]Supported by the *Deutsch Institute*.

Complexity Classification of Some Edge Modification Problems

Assaf Natanzon[†] Ron Shamir[†] Roded Sharan^{†‡}

Abstract

In an edge modification problem one has to change the edge set of a given graph as little as possible so as to satisfy a certain property. We prove in this paper the NP-hardness of a variety of edge modification problems with respect to some well-studied classes of graphs. These include perfect, chordal, chain, comparability, split and asteroidal triple free. We show that some of these problems become polynomial when the input graph has bounded degree. We also give a general constant factor approximation algorithm for deletion and editing problems on bounded degree graphs with respect to properties that can be characterized by a finite set of forbidden induced subgraphs.

Presented at the Twenty-Fifth International Workshop on Graph-Theoretic Concepts in Computer Science, Ascona, Switzerland, June 17–19, 1999.

[†]Department of Computer Science, Tel Aviv University, Tel Aviv, Israel.
{natanzon,shamir,roded}@math.tau.ac.il.

[‡]*Supported by the Deutsch Institute.*

An Improved Bound for k -Sets in Three Dimensions

Micha Sharir[†] Shakhar Smorodinsky[‡] Gábor Tardos[§]

Abstract

Let S be a set of n points in R^d . A k -set of S is a subset $S' \subset S$ such that $S' = S \cap H$ for some halfspace H and $|S'| = k$. The problem of determining tight asymptotic bounds on the maximum number of k -sets is one of the most intriguing open problems in combinatorial geometry. Due to its importance in analyzing geometric algorithms, the problem has caught the attention of computational geometers as well. A close to optimal solution for the problem remains elusive even in the plane. The best asymptotic upper and lower bounds in the plane are $O(nk^{1/3})$ and $n \cdot 2^{\Omega(\sqrt{\log k})}$, respectively. In this paper we obtain the following result: **Theorem:** The number of k -sets in a set of n points in R^3 is $O(nk^{3/2})$.

This result improves the previous best known asymptotic upper bound of $O(nk^{5/3})$ (see Dey and Edelsbrunner and Agarwal et al.). The best known asymptotic lower bound for the number of k -sets in three dimensions is $nk \cdot 2^{\Omega(\sqrt{\log k})}$.

Presented at the 16th Annual ACM Symposium on Computational Geometry 2000 Hong-Kong.

[†]Tel-Aviv University

[‡]Tel-Aviv University

[§]Rényi Institute of the Hungarian Academy of Sciences, H-1364 Budapest, POB 127, Hungary;

A Model for Visual Camouflage Breaking

Ariel Tankus^{†‡} Yehezkel Yeshurun[‡]

Abstract

Some animals use counter-shading in order to prevent their detection by predators. Counter-shading means that the albedo of the animal is such that its image has a flat intensity function rather than a convex intensity function. This implies that there might exist predators who can detect 3D objects based on the convexity of the intensity function. In this paper, we suggest a mathematical model which describes a possible explanation of this detection ability. We demonstrate the effectiveness of convexity based camouflage breaking using an operator (“ D_{arg} ”) for detection of 3D convex or concave graylevels. Its high robustness and the biological motivation make D_{arg} particularly suitable for camouflage breaking. As will be demonstrated, the operator is able to break very strong camouflage, which might delude even human viewers. Being non-edge-based, the performance of the operator is juxtaposed with that of a representative edge-based operator in the task of camouflage breaking. Better performance is achieved by D_{arg} for both animal and military camouflage breaking.

Presented at the IEEE International Workshop on Biologically Motivated Computer Vision (BMCV2000), Seoul, Korea, May 15–17, 2000.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]Supported by the Deutsch Institute

Detection of Regions of Interest and Camouflage Breaking by Direct Convexity Estimation

Ariel Tankus^{†‡} Yehezkel Yeshurun[‡]

Abstract

Detection of Regions of Interest is usually based on edge maps. We suggest a novel non-edge-based mechanism for detection of regions of interest, which extracts 3D information from the image. Our operator detects smooth 3D convex and concave objects based on direct processing of intensity values. Invariance to a large family of functions is mathematically proved. It follows that our operator is robust to variation in illumination, orientation, and scale, in contrast with most other attentional operators. The operator is also demonstrated to efficiently detect 3D objects camouflaged in noisy areas. An extensive comparison with edge-based attentional operators is delineated.

Presented at the IEEE International Workshop on Visual Surveillance '98 (VS98), Bombay, India, January 2, 1998 (In conjunction with ICCV'98).

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]Supported by the Deutsch Institute

Warped Textures for UV Mapping Encoding

Olga Sorkine^{†‡} Daniel Cohen-Or[§]

Abstract

This paper introduces an implicit representation of the u, v texture mapping. Instead of using the traditional explicit u, v mapping coordinates, a non-distorted piecewise embedding of the triangular mesh is created, on which the original texture is remapped, yielding warped textures. This creates an effective atlas of the mapped triangles and provides a compact encoding of the texture mapping.

Presented at EUROGRAPHICS '01, Manchester, England, September 4-7, 2001

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Verifying Safety Properties of Concurrent Java Programs using 3-Valued Logic

Eran Yahav[†]

Abstract

We provide a parametric framework for verifying safety properties of concurrent Java programs. The framework combines thread-scheduling information with information about the shape of the heap. This leads to error-detection algorithms that are more precise than existing techniques. The framework also provides the most precise shape-analysis algorithm for concurrent programs. In contrast to existing verification techniques, we do not put a bound on the number of allocated objects. The framework even produces interesting results when analyzing Java programs with an unbounded number of threads. The framework is applied to successfully verify the following properties of a concurrent program:

- Concurrent manipulation of linked-list based ADT preserves the ADT datatype invariant.
- The program does not perform inconsistent updates due to interference.
- The program does not reach a deadlock.
- The program does not produce run-time errors due to illegal thread interactions.

We also find bugs in erroneous versions of such implementations.

A prototype of our framework has been implemented.

Presented at the 28th international conference of Principles of Programming Languages (POPL), London, UK, January 17–19, 2001.

[†]School of Computer Science, Tel-Aviv University, Tel-Aviv, supported by the Deutsch Institute.

A Centralized Dynamic Access Probability Protocol for Next Generation Wireless Networks

Zohar Naor[†] Hanoch Levy[‡]

Abstract

A multiple access protocol that is particularly suitable for cellular Internet access and satellite-based networks with on-board processing is developed in this paper. The basic idea is that when a user wishes to send a message, it transmits with a probability p_{access} that depends on the load on the channel. Under conditions of low load, the probability p_{access} approaches 1, while at high load p_{access} is relatively low. This media access control protocol guarantees high channel utilization at high load, as well as low delay at low load periods. Using the statistical usage of the shared channel, the load is estimated with certain uncertainty. Our analysis shows that using the statistical usage of the shared channel, the optimal access probability can be well estimated for a broad class of load distribution patterns. In addition, we propose to use a central station to broadcast the value of p_{access} in networks with poor collision detection capability, or long feedback delay. The proposed method is particularly suitable for shared channels with poor collision detection capability, under conditions of bursty traffic and a large number of users. Examples for such channels are the reservation channel in satellite-based networks with on-board processing, and the control channel in cellular networks. Hence, the proposed method can be used for cellular Internet access and for accessing public satellite-based networks. *The broadcast mechanism that already exists in such networks can be used to inform the users the dynamic access probability.*

Presented at INFOCOM 2001, ALASKA, USA, April 22-26, 2001.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel
[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel
Supported by the Deutsch Institute

A Server-based Interactive Remote Walkthrough

Daniel Cohen-Or[†] Yuval Noimark^{‡§¶} Tali Zvi^{||}

Abstract

This paper presents a server-based remote walkthrough system. The client is assumed to be a thin client, like a handset or a mobile device, with no strong processor but with some embedded video chip. The server holds the large environment, generates the frames, encodes and transmits them to the client. The encoded frames are transmitted as a video stream to the client, which then decodes the stream and displays it. We show how the computer generated frames can be efficiently encoded using layering techniques to yield a lighter stream, which enables its transmission over narrow bandwidth channels and minimizes the communication latency. To enable the interactivity of the system, the rendering engine generates the frames in real-time according to the client input, and feeds the frames to an accelerated video encoder based on the available optical flow.

Presented at the 6th Eurographics Workshop on Multimedia, Manchester, United Kindom, September 8–9, 2001.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[§]*Supported by the Deutsch Institute*

[¶]IBM Research Lab in Haifa

^{||}Enbaya Ltd.

Detection of Pseudo Periodic Functional Patterns using Partial Acquisition of Magnetic Resonance Images

Oren Boiman^{†‡} Yehezkel Yeshurun[§] Sharon Peled[¶] Talma Hendler^{||}

Abstract

Pseudo-periodic patterns are frequently encountered in the cerebral cortex due to its columnar functional organization (best exemplified by orientation columns and ocular dominance columns of the Visual Cortex). This work presents a novel partial acquisition strategy and reconstruction algorithm, suitable for detection of these pseudo-periodic patterns. We present a new Magnetic Resonance Imaging (MRI) research methodology, in which we seek an activity pattern with a spatial scale below the observable resolution, and a pattern-specific experiment is devised to detect it. Such specialized experiments extend the limits of conventional MRI experiments by substantially reducing the scan time (saving up to 90%). Using the fact that pseudo-periodic patterns are localized in the Fourier domain, we present an optimality criterion for partial acquisition of the MR signal and a strategy for obtaining the optimal discrete Fourier transform coefficients. A by-product of this strategy is an optimal linear extrapolation estimate. We also present a non-linear spectral extrapolation algorithm, based on POCS, used to perform the actual reconstruction. The proposed strategy was tested and analyzed on simulated signals and in MRI Phantom experiments.

Presented at the Minimum MR Data Acquisition Methods Workshop, Marco Island, Florida ,USA , October 20–21, 2001.

[†]School of Computer Science, Tel-Aviv University

[‡]*Supported by the Deutsch Institute*

[§]School of Computer Science, Tel-Aviv University

[¶]Wohl Institute for Advanced Imaging. Tel-Aviv Sourasky Medical Center

^{||}Wohl Institute for Advanced Imaging. Tel-Aviv Sourasky Medical Center

On Neighbours in Geometric Permutations

Micha Sharir Shakhar Smorodinsky[†]

Abstract

We introduce a new notion of ‘neighbors’ in geometric permutations. We conjecture that the maximum number of neighbors in a set of n pairwise disjoint convex bodies in \mathcal{R}^d is $O(n)$, and we settle this conjecture for $d = 2$. We show that if the set of pairs of neighbors in a set S is of size N , then S admits at most $O(N^{d-1})$ geometric permutations. Hence we obtain an alternative proof of a linear upper bound on the number of geometric permutations for any finite family of pairwise disjoint convex bodies in the plane.

Presented at the Eighth Scandinavian Workshop on Algorithmic Theory, Turku, Finland, July 3–5, 2002.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel
supported by the *Deutsch Institute*

Online Subpath Profiling

David Oren^{*†} Yossi Matias^{‡§} Mooly Sagiv^{¶||}

Abstract

We present an efficient online subpath profiling algorithm, OSP, that reports hot subpaths executed by a program in a given run. The hot subpaths can start at arbitrary basic block boundaries, and their identification is important for code optimization; e.g., to locate program traces in which optimizations could be most fruitful, and to help programmers in identifying performance bottlenecks.

The OSP algorithm is online in the sense that it reports at any point during execution the hot subpaths as observed so far. It has very low memory and runtime overheads, and exhibits high accuracy in reports for benchmarks such as **JLex** and **FFT**. These features make the OSP algorithm potentially attractive for use in just-in-time (JIT) optimizing compilers, in which profiling performance is crucial and it is useful to locate hot subpaths as early as possible.

The OSP algorithm is based on an adaptive sampling technique that makes effective utilization of memory with small overhead. Both memory and runtime overheads can be controlled, and the OSP algorithm can therefore be used for arbitrarily large applications, realizing a tradeoff between report accuracy and performance.

We have implemented a Java prototype of the OSP algorithm for Java programs. The implementation was tested on programs from the Java Grande benchmark suite and exhibited a low average runtime overhead.

Presented at the International Conference on Compiler Construction, Grenoble, France, April 8–12, 2002.

^{*}School of Computer Science, Tel-Aviv University

[†]*Supported by the Deutsch Institute*

[‡]School of Computer Science, Tel-Aviv University

[§]Research supported in part by an Alon Fellowship and by the Israel Science Foundation founded by The Academy of Sciences and Humanities

[¶]School of Computer Science, Tel-Aviv University

^{||}Research supported in part by the Israel Science Foundation founded by The Academy of Science and Humanities

Polynomial Curves in Parallel Coordinates :

Results and Constructive Algorithm

Tsur Izhakian^{†‡} Alfred Inselberg[§]

Abstract

An application, based on parallel coordinates (abbr. \parallel -coords) on "approximated planes" was presented at this conference in 2000 by Matskewich. With Parallel coordinates, objects in R^n can be represented, without loss of information, by planar patterns for arbitrary n . In R^n , embedded in the projective plane, parallel coordinates induce a *point* \leftrightarrow *line* and other dualities which generalize nicely to R^n . In 1981 it was shown that conics are mapped into conics in 6 different ways. Later this was generalized to bounded and unbounded convex sets and eventually applied to higher dimensions. Since then the question of "what is the dual image of general polynomial curves" has not been answered. Here we show that the dual image in \parallel -coords of an algebraic curve of degree n is also algebraic of degree $n(n - 1)$ in absence of singular points. Further an algorithm for the construction of the dual even in the presence of singularities is found and presented here. The result is of interest in its right and opens the prospects for extending the multi-dimensional applications.

Presented at the 5th International Conference on Curves and Surfaces, Saint-Malo - France, June 27 - July 3, 2002

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Obvious or Not? Regulating Architectural Decisions Using Aspect-Oriented Programming

Mati Shomrat^{†‡} Amiram Yehudai[§]

Abstract

The construction of complex, evolving software systems requires a high-level design model. However, this model tends not to be enforced on the system, leaving room for the implementors to diverge from it, thus differentiating the designed system from the actual implemented one. The essence of the problem of enforcing such models lies in their globality. The principles and guidelines conveyed by these models cannot be localized in a single module, they must be observed everywhere in the system. A mechanism for enforcement needs to have a global view of the system and to report breaches in the model at the time they occur. Aspect-Oriented Programming has been proposed as a new software engineering approach. Unlike contemporary software engineering methods, which are module centered, Aspect Oriented Programming provides mechanisms for the definition of cross-module interactions. We explore the possibility of using Aspect-Oriented Programming in general and the AspectJ programming language in particular for the enforcement of design models.

Presented at the 1st International Conference on Aspect-Oriented Software Development, Enschede, The Netherlands, April 22–26, 2002.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel. Supported by the Deutsch Institute.

[‡]*Supported by the Deutsch Institute*

[§]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel. Currently at The Academic College of Tel-Aviv-Yaffo.

Roundtrip Spanners and Roundtrip Routing in Directed Graphs

Liam Roditty^{†‡} Mikkel Thorup[§] Uri Zwick[†]

Abstract

We introduce the notion of *roundtrip-spanners* of weighted *directed* graphs and describe efficient algorithms for their construction. For every integer $k \geq 1$ and any $\epsilon > 0$, we show that any directed graph on n vertices with edge weights in the range $[1, W]$ has a $(2k + \epsilon)$ -roundtrip-spanner with $O(\frac{k^2}{\epsilon} n^{1+1/k} \log(nW))$ edges. We then extend these constructions and obtain compact roundtrip routing schemes. For every integer $k \geq 1$ and every $\epsilon > 0$, we describe a roundtrip routing scheme that has stretch $4k + \epsilon$, and uses at each vertex a routing table of size $\tilde{O}(\frac{k^2}{\epsilon} n^{1/k} \log(nW))$. We also show that any weighted directed graph with *arbitrary* positive edge weights has a 3-roundtrip-spanner with $O(n^{3/2})$ edges. This result is optimal. Finally, we present a stretch 3 roundtrip routing scheme that uses local routing tables of size $\tilde{O}(n^{1/2})$. This routing scheme is essentially optimal. The roundtrip-spanner constructions and the roundtrip routing schemes for directed graphs that we describe are only slightly worse than the best available spanners and routing schemes for undirected graphs. Our roundtrip routing schemes substantially improve previous results of Cowen and Wagner. Our results are obtained by combining ideas of Cohen, Cowen and Wagner, Thorup and Zwick, with some new ideas.

Presented at the 13th ACM-SIAM Symposium on discrete Algorithms, San-Francisco, CA, USA, January 6–8, 2002.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel.

[‡]*Supported by the Deutsch Institute*

[§]AT&T Labs - Research, 180 Park Avenue, Florham Park, NJ 07932, USA.

Improved Algorithms for the Random Cluster Graph Model

Ron Shamir[†] Dekel Tsur^{† ‡}

Abstract

The following probabilistic process models the generation of noisy clustering data: Clusters correspond to disjoint sets of vertices in a graph. Each two vertices from the same set are connected by an edge with probability p , and each two vertices from different sets are connected by an edge with probability $r < p$. The goal of the clustering problem is to reconstruct the clusters from the graph. We give algorithms that solve this problem with high probability. Compared to previous studies, our algorithms have lower time complexity and wider parameter range of applicability. In particular, our algorithms can handle $O(\sqrt{n}/\log n)$ clusters in an n -vertex graph, while all previous algorithms require that the number of clusters is constant.

Presented at the 8th Scandinavian Workshop on Algorithm Theory,
Turku, Finland, July 3–5, 2002.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

All-norm Approximation Algorithms

Yossi Azar[†] Leah Epstein[‡] Yossi Richter^{§¶} Gerhard Woeginger^{||}

Abstract

A major drawback in optimization problems and in particular in scheduling problems is that for every measure there may be a different optimal solution. In many cases the various measures are different ℓ_p norms. We address this problem by introducing the concept of an *All-norm ρ -approximation algorithm*, which supplies one solution that guarantees ρ -approximation to all ℓ_p norms simultaneously. Specifically, we consider the problem of scheduling in the restricted assignment model, where there are m machines and n jobs, each is associated with a subset of the machines and should be assigned to one of them. Previous work considered approximation algorithms for each norm separately. Lenstra *et al.* (LST90) showed a 2-approximation algorithm for the problem with respect to the ℓ_∞ norm. For any fixed ℓ_p norm the previously known approximation algorithm has a performance of $\theta(p)$. We provide an all-norm 2-approximation polynomial algorithm for the restricted assignment problem. On the other hand, we show that for any given ℓ_p norm ($p > 1$) there is no PTAS unless $P=NP$ by showing an APX-hardness result. We also show for any given ℓ_p norm a FPTAS for any fixed number of machines.

Presented at the 8th Scandinavian Workshop on Algorithm Theory, Turku, Finland, July 3–5, 2002.

[†]School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel

[‡]School of Computer and Media Sciences, The Interdisciplinary Center, Herzliya, Israel

[§]School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel

[¶]*Supported by the Deutsch Institute*

^{||}Department of Mathematics, University of Twente, Enschede, The Netherlands

Language Diversity: Evidence For Language Fitness

Zach Solan^{†‡} Shimon Edelman[§] David Horn[¶] Eytan Ruppin^{||}

Abstract

In the last few years the study of language evolution has received extensive attention. Many Studies have explored the dynamics and the conditions in which a structured coherent language can emerge. These efforts have been mainly guided by the notion of natural selection, which assumes that language contributes to human fitness, and that successful communication leads to an increase in individual survival probability. Although this assumption is quite intuitive, evidence for its presence is less obvious. In this work we will present some new observations, which are based on empirical data and simulations that address this issue by studying the evolutionary dynamics of language diversity.

While looking at the distribution of languages around the world, one can notice an interesting phenomenon: languages are not uniformly distributed across the globe. Hundred of languages have evolved in Africa and Oceania while very few did in North America or Europe. While most of these languages are almost extinct, these relics from ancient times can shed light on the evolutionary forces that were involved in generating languages.

One of the aims of genetic research of human populations is to explore the reasons for their diversity. Several studies in this domain revealed the point in time in which population expansion occurred. Excoffier and Schneider [1999] examined potential signals of population expansions by analyzing the mitochondrial DNA diversity of 62 human population samples. Their results suggest that population expansions occurred in different regions at different times. Comparing Excoffier and Schneider's estimates of the age of population expansion in 32 different regions with the number of languages evolved in these regions reveals a correlation of 0.5 ($p > 0.001$). In contrast, the number of languages per region is neither correlated with population size, nor with region size. Thus, we can conclude that the ancient societies have generated much more languages than the younger ones. This conclusion, and the notion that most of the languages die out as time passes (95% of the population on earth uses only 100 of the existing 6000 languages), reveals a picture in which most of the languages that have been categorized are actually products of the early stage when languages have just started to evolve.

What kind of evolutionary dynamics can account for this mass of languages, which appeared in the early stage of evolution? In order to investigate this question we introduce a model (based on a shared lexicon model of Nowak 1999) that simulates evolvement of shared lexicons in isolated populations with controlled migration between them. We tested two hypotheses, one that uses a linear fitness function of natural selection and one that is neutral and is based on genetic drift solely. In a situation where no fitness function is present, it takes much longer to reach the steady state of one common language. However, even a small amount of migration (1% of the population) suffices for one language to be dominant. In contradistinction, when natural selection is introduced (i.e., the fitness of individuals is assumed to be proportional to their level of shared communication) the subpopulations stabilize quite fast to form several distinct languages. Only relatively high migration rates will now force the system to move into the regime of one dominant language. This phenomenon has an analogy in solid-state physics, where when certain materials are cooled rapidly they crystallize heterogenically, while if the same materials are cooled slowly they crystallize homogenously. Hence, these findings support the hypothesis that language incurs increased human fitness, and speak against the possibility that a shared lexicon evolved as a result of neutral drift, unless migration rates were very small. We conclude with an analytic estimate of an upper bound on migration rate that allows for the

diversity of languages observed today.

[†]School of Physics Tel-Aviv University, Tel-Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]Department of Psychology, Cornell University Ithaca, NY

[¶]School of Physics Tel-Aviv University, Tel-Aviv, Israel

^{||}School of Computer Science, Tel Aviv University, Tel Aviv, Israel

Scene-Consistent Detection of Feature Points in Video Sequences

Ariel Tankus^{†‡} Yehezkel Yeshurun[‡]

Abstract

Detection of feature points in images is an important preprocessing stage for many algorithms in Computer Vision. We address the problem of detection of feature points in video sequences of 3D scenes, which could be mainly used for obtaining scene correspondence. The main feature we use is the zero crossing of the intensity gradient argument. We analytically show that this local feature corresponds to specific constraints on the local 3D geometry of the scene, thus ensuring that the detected points are based on real 3D features. We present a robust algorithm that tracks the detected points along a video sequence, and suggest some criteria for quantitative evaluation of such algorithms. These criteria serve in a comparison of the suggested operator with two other feature trackers. The suggested criteria are generic and could serve other researchers as well for performance evaluation of stable point detectors.

Presented at the Computer Vision and Pattern Recognition conference (CVPR 2001), Kauai, Hawaii, USA, Dec. 11–13, 2001.

[†]School of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

Union-find with deletions

Haim Kaplan [†] Nira Shafir ^{‡§} Robert E. Tarjan [¶]

Abstract

In the classical union-find problem we maintain a partition of a universe of n elements into disjoint sets subject to the operations union and find. The operation $union(A, B, C)$ replaces sets A and B in the partition by their union, given the name C . The operation $find(x)$ returns the name of the set containing the element x . In this paper we revisit the union-find problem in a context where the underlying partitioned universe is not fixed. Specifically, we allow a $delete(x)$ operation which removes the element x from the set containing it. We consider both worst-case performance and amortized performance. In both settings the challenge is to dynamically keep the size of the structure representing each set proportional to the number of elements in the set which may now decrease as a result of deletions.

For any fixed k , we describe a data structure that supports find and delete in $O(\log_k n)$ worst-case time and union in $O(k)$ worst-case time. This matches the best possible worst-case bounds for find and union in the classical setting. Furthermore, using an incremental global rebuilding technique we obtain a reduction converting any union-find data structure to a union-find with deletions data structure. Our reduction is such that the time bounds for find and union change only by a constant factor. The time it takes to delete an element x is the same as the time it takes to find the set containing x plus the time it takes to unite a singleton set with this set.

In an amortized setting a classical data structure of Tarjan supports a sequence of m finds and at most n unions on a universe of n elements in $O(n + m\alpha(m+n, n, \log n))$ time where $\alpha(m, n, l) = \min\{k | A_k(\lfloor \frac{m}{n} \rfloor) > l\}$ and $A_i(j)$ is Ackermann's function. data structure and show that in fact the cost of each find is proportional to the size of the corresponding set. Specifically, we show that one can pay for a sequence of union and find operations by charging a constant to each participating element and $O(\alpha(m, n, \log(l)))$ for a find of an element in a set of size l . We also show how keep these amortized costs for each find and each participating element while allowing deletions. The amortized cost of deleting an element from a set of l elements is the same as the amortized cost of finding the element; namely, $O(\alpha(m, n, \log(l)))$.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Tel Aviv, Israel.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Tel Aviv, Israel.

[§]*Supported by the Deutsch Institute*

[¶]Department of Computer Science, Princeton University, Princeton, NJ 08544, and InterTrust Technologies Corporation 4750 Patrick Henry Drive, Santa Clara, 95054-1851.

A comparison of labeling schemes for ancestor queries

Haim Kaplan [†] Tova Milo [‡] Ronen Shabo ^{§¶}

Abstract

Motivated by a recent application in XML search engines we study the problem of labeling the nodes of a tree (XML file) such that given the labels of two nodes one can determine whether one node is an ancestor of the other. We describe several new prefix-based labeling schemes, where an ancestor query roughly amounts to testing whether one label is a prefix of the other. We compare our new schemes to a simple interval-based scheme currently used by search engines, as well as, to schemes with the best theoretical guarantee on the maximum label length. We performed our experimental evaluation on real XML data and on some families of random trees.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Tel Aviv, Israel.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Tel Aviv, Israel.

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Tel Aviv, Israel.

[¶]*Supported by the Deutsch Institute*

Efficient Construction of the Union of Geometric Objects [†]

Eti Ezra^{‡§} Dan Halperin[¶] Micha Sharir^{||}

Abstract

We present a new incremental algorithm for constructing the union of n triangles in the plane. In our experiments, the new algorithm, which we call the Disjoint-Cover (DC) algorithm, performs significantly better than the standard randomized incremental construction (RIC) of the union. Our algorithm is rather hard to analyze rigorously, but we provide an initial such analysis, which yields an upper bound on its performance that is expressed in terms of the expected cost of the RIC algorithm. Our approach and analysis generalize verbatim to the construction of the union of other objects in the plane, and, with slight modifications, to three dimensions. We present experiments with a software implementation of our algorithm using the CGAL library of geometric algorithms.

Presented at the 18th European Workshop on Computational Geometry , Warszawa, Poland, April 10–12, 2002.

[†]Work reported in this paper has been supported in part by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-2000-26473 (ECG - Effective Computational Geometry for Curves and Surfaces), by The Israel Science Foundation founded by the Israel Academy of Sciences and Humanities (Center for Geometric Computing and its Applications), and by the Hermann Minkowski – Minerva Center for Geometry at Tel Aviv University. Micha Sharir has also been supported by NSF Grants CCR-97-32101 and CCR-00-98246, and by a grant from the U.S.-Israeli Binational Science Foundation.

[‡]*Supported by the Deutsch Institute*

[§]School of Computer Science, Tel Aviv University

[¶]School of Computer Science, Tel Aviv University

^{||}School of Computer Science, Tel Aviv University

Harmonic Buffer Management Policy for Shared Memory Switches

Alexander Kesselman^{†‡} Yishay Mansour[§]

Abstract

We introduce a new general scheme for shared memory non-preemptive scheduling policies. Our scheme utilizes a system of inequalities and thresholds and accepts a packet if it does not violate any of the inequalities. We demonstrate that many of the existing policies can be described using our scheme, thus validating its generality.

We propose a new scheduling policy, based on our general scheme, which we call the Harmonic policy. Our simulations show that the Harmonic policy both achieves high throughput and easily adapts to changing load conditions. We also perform a theoretical analysis of the Harmonic policy and demonstrate that its throughput competitive ratio is almost optimal. _____

Presented at the INFOCOM'02.

[†]School of Computer Science, Tel Aviv University, Tel Aviv, Israel

[‡]*Supported by the Deutsch Institute*

[§]School of Computer Science, Tel Aviv University, Tel Aviv, Israel

A Ramsey-type Theorem for Metric Spaces and its Applications for Metrical Task Systems and Related Problems

Yair Bartal[†] Béla Bollobás[‡] Manor Mendel^{§¶}

Abstract

A nearly logarithmic lower bound on the randomized competitive ratio for the *metrical task systems* problem is presented. This implies a similar lower bound for the extensively studied K -server problem. The proof is based on Ramsey-type theorems for metric spaces, that state that every metric space contains a large subspace which is approximately a “hierarchically well-separated tree” (HST) (and in particular an ultrametric). These Ramsey-type theorems may be of independent interest.

Presented at the 42nd annual Symposium on Foundations of Computer Science, 2001.

[†]The Hebrew University, Jerusalem, Israel

[‡]University of Memphis, Memphis, TN

[§]*Supported by the Deutsch Institute*

[¶]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Estimating the Impact of Heap Liveness Information on Space Consumption

Ran Shaham^{†‡} Elliot K. Kolodner[§] Mooly Sagiv[¶]

Abstract

We study the potential impact of different kinds of liveness information on the space consumption of a program in a garbage collected environment, specifically for Java. The idea is to measure the time difference between the actual time an object is collected by the garbage collector (GC) and the potential earliest time an object could be collected assuming liveness information were available. We focus on the following kinds of liveness information: (i) stack reference liveness (local reference variable liveness in Java), (ii) global reference liveness (static reference variable liveness in Java), (iii) heap reference liveness (instance reference variable liveness or array reference liveness in Java), and (vi) any combination of (i)-(iii). We also provide some insights on the kind of interface between a compiler and GC that could achieve these potential savings.

The Java Virtual Machine (JVM) was instrumented to measure (dynamic) liveness information. Experimental results are given for 10 benchmarks, including 5 of the SPEC-jvm98 benchmark suite. We show that in general stack reference liveness may yield small benefits, global reference liveness combined with stack reference liveness may yield medium benefits, and heap reference liveness yields the largest potential benefit.

Presented at the 2002 International Symposium on Memory Management, Berlin, Germany, June 21–22, 2002.

[†]*Supported by the Deutsch Institute*

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

[§]IBM Haifa Research Laboratory, Haifa, Israel

[¶]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Approximating MIN k -SAT

Adi Avidor^{†‡} Uri Zwick^{§¶}

Abstract

We obtain substantially improved approximation algorithms for the MIN k -SAT problem, for $k = 2, 3$. More specifically, we obtain a 1.1037-approximation algorithm for the MIN 2-SAT problem, improving a previous 1.5-approximation algorithm, and a 1.2136-approximation algorithm for the MIN 3-SAT problem, improving a previous 1.75-approximation algorithm for the problem. These results are obtained by adapting techniques that were previously used to obtain approximation algorithms for the MAX k -SAT problem. We also obtain some hardness of approximation results.

Presented at the 13th Annual International Symposium on Algorithms and Computation Vancouver, Canada, November 20-23, 2002.

[†]This research was supported by the ISRAEL SCIENCE FOUNDATION (grant no. 246/01)

[‡]School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel

[§]This research was supported by the ISRAEL SCIENCE FOUNDATION (grant no. 246/01)

[¶]School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel

[†] Supported by the Deutsch Institute

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Cache Satellite Distribution Systems: Modeling and Analysis

Aner Armon Hanoch Levy

Abstract

Web caches have become an integral component contributing to the improvement of the performance observed by Web clients. Content Distribution Networks (CDN) and Cache Satellite Distribution Systems (CSDS) have emerged as technologies for feeding the caches with the information clients are expected to request, ahead of time. In a Cache Satellite Distribution System (CSDS), the proxies participating in the CSDS periodically report to a central station about the requests they are receiving from their clients. The central station processes this information and selects a collection of Web documents (or Web pages), which it then "pushes" via a satellite broadcast to all, or some, of the participating proxies, hoping most of them will request most documents in the near future. The result is that upon such request, the documents will reside in the local cache, and will not need to be fetched.

In this paper we aim at addressing the issues of how to operate the CSDS, how to design it, and how to estimate its effect. Questions of interest are 1) What classes of Web documents should be transmitted by the central station, and how they are characterized, and 2) What is the benefit of adding a particular proxy into a CSDS. We offer a model of this system that accounts for the request streams addressed to the proxies and which captures the intricate interaction between the proxy caches. Unlike models that are based only on the access frequency of the various documents, this model captures both their frequency and their locality of reference. We provide an analysis of this system that is based on the stochastic properties of the traffic streams that can be derived from HTTP logs. The model and analysis can serve as a basis for the design and efficient operation of the system.

Presented at IEEE INFOCOM'03 , San Francisco, USA, March 30–April 3, 2003.

Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel *Supported by the Deutsch Institute and by MAGNET, Chief Scientist Office, Ministry of Trade and Commerce, Israel.*

On the structure and application of BGP policy Atoms

Yehuda Afek Omer Ben-Shalom Anat Bremler-Barr

Abstract

The notion of Internet *Policy Atoms* has been recently introduced by Andre Broido and kc claffy from CAIDA as groups of prefixes sharing a common BGP AS path at any Internet backbone router. In this paper we further research these 'Atoms'. First we offer a new method for computing the Internet *policy atoms*, and use the RIPE RIS database to derive their structure. Second, we show that atoms remain stable with only about 2-3% of prefixes changing their atom membership in eight hour periods. We support the 'Atomic' nature of the *policy atoms* by showing BGP update and withdraw notifications carry updates for complete atoms in over 70% of updates, while the complete set of prefixes in an AS is carried in only 21% of updates. We track the locations where atoms are created (first different AS in the AS path going back from the common origin AS ⁴) showing 86% are split between the origin AS and it's peers thus supporting the assumption that they are created by policies. Finally applying atoms to "real life" applications we achieve a modest savings in BGP updates due to the low average prefix count in the atoms.

Presented at the IMW 2002 conference, Marseilles France, Nov 6-8 2002

[¶]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel. Supported in part by the Deutsch Institute

⁴The origin AS is the last AS in the AS path and is the AS that has the prefix allocated to it. In this article we also refer to it as the 'owning AS' to avoid possible confusion with an AS sending traffic to a prefix

Verifying Temporal Heap Properties Specified via Evolution Logic

Eran Yahav Thomas Reps Mooly Sagiv Reinhard Wilhelm

Abstract

This paper addresses the problem of establishing temporal properties of programs written in languages, such as Java, that make extensive use of the heap to allocate—and deallocate—new objects and threads. Establishing liveness properties is a particularly hard challenge. One of the crucial obstacles is that heap locations have no static names and the number of heap locations is unbounded. The paper presents a framework for the verification of Java-like programs. Unlike classical model checking, which uses propositional temporal logic, we use first-order temporal logic to specify temporal properties of heap evolutions; this logic allows domain changes to be expressed, which permits allocation and deallocation to be modelled naturally. The paper also presents an abstract-interpretation algorithm that automatically verifies temporal properties expressed using the logic.

To be presented at the European Symposium on Programming, Warsaw, Poland, April 5–13, 2003.

Supported by the Deutsch Institute, School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel, {yahave,msagiv}@post.tau.ac.il Computer Science Dept., University of Wisconsin, Madison, WI 53706, USA, reps@cs.wisc.edu Informatik, University des Saarlandes, Saarbrücken, Germany, wilhelm@cs.uni-sb.de

2D Arrangements in CGAL: Recent Developments

Efi Fogel

Abstract

Given a collection of curves in the plane C , the arrangement of C is the subdivision of the plane into vertices, edges, and facets induced by the curves of C . Constructing arrangements of curves in the plane is a basic problem in computational geometry. Applications relying on arrangements arise in fields such as robotics, computer vision, and computer graphics. Many algorithms for constructing and maintaining arrangements under various conditions have been published. However, not many (general) arrangement packages are publicly available. The CGAL library contains a package that provides the generic and robust construction and manipulation of arrangements of curves.

We provide an overview of the recent developments that have taken place recently within the arrangement package in CGAL. The changes are divided into two categories. One is improvements in the efficiency of the various operations. The other is improvements in the interface. A simple benchmark system is used to monitor all changes and reject changes that cause regressions.

Presented at the ECG Workshop on Robustness and Efficiency Issues in Implementing Arrangements of Curves and Surfaces, Inria, France, December 18-19, 2002.

^{||}Computer Science, Tel-Aviv University, Tel-Aviv, Israel. Supported by the Deutsch Institute.

Testing Juntas

Eldar Fischer Guy Kindler Dana Ron Muli Safra Alex Samorodnitsky

Abstract

We show that a Boolean function over n Boolean variables can be tested for the property of depending on only k of them, using a number of queries that depends only on k and the approximation parameter ϵ . We present two tests, both non-adaptive, that require a number of queries that is polynomial k and linear in ϵ^{-1} . The first test is stronger in that it has a 1-sided error, while the second test has a more compact analysis. We also present an adaptive version and a 2-sided error version of the first test, that have a somewhat better query complexity than the other algorithms.

We then provide a lower bound of $\tilde{\Omega}(\sqrt{k})$ on the number of queries required for the non-adaptive testing of the above property; a lower bound of $\Omega(\log(k+1))$ for adaptive algorithms naturally follows from this. In providing this we also prove a result about random walks on the group Z_2^q that may be interesting in its own right. We show that for some $t(q) = \tilde{O}(q^2)$, the distributions of the random walk at times t and $t+2$ are close to each other, independently of the step distribution of the walk.

Presented at the 43rd Annual IEEE Symposium on Foundations of Computer Science, Vancouver, Canada, November 16–19, 2002.

**Faculty of Computer Science, The Technion, Haifa, Israel.

Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel, *Supported by the Deutsch Institute*. Dept. of Electrical Engineering - Systems, Tel-Aviv University, Tel Aviv, Israel. Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel. School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel

Improved dynamic reachability algorithms for directed graphs

Liam Roditty^{†‡} Uri Zwick^{†‡}

Abstract

We obtain several new dynamic algorithms for maintaining the transitive closure of a directed graph, and several other algorithms for answering reachability queries without explicitly maintaining a transitive closure matrix. Among our algorithms are:

- (i) A decremental algorithm for maintaining the transitive closure of a directed graph, through an arbitrary sequence of edge deletions, in $O(mn)$ *total* expected time, essentially the time needed for computing the transitive closure of the initial graph. Such a result was previously known only for *acyclic* graphs.
- (ii) Two fully dynamic algorithms for answering reachability queries. The first is deterministic and has an amortized insert/delete time of $O(m\sqrt{n})$, and worst-case query time of $O(\sqrt{n})$. The second is randomized and has an amortized insert/delete time of $O(m^{0.58}n)$ and worst-case query time of $O(m^{0.43})$. This significantly improves the query times of algorithms with similar update times.
- (iii) A fully dynamic algorithm for maintaining the transitive closure of an acyclic graph. The algorithm is deterministic and has a worst-case insert time of $O(m)$, *constant* amortized delete time of $O(1)$, and a worst-case query time of $O(n/\log n)$.

Our algorithms are obtained by combining several new ideas, one of which is a simple *sampling* idea used for detecting decompositions of strongly connected components, with techniques of Even and Shiloach , Italiano , Henzinger and King , and Frigioni et al. . We also adapt results of Cohen on estimating the *size* of the transitive closure to the dynamic setting.

Presented at the 43rd Annual IEEE Symposium on Foundations of Computer Science, Vancouver, Canada, November 16-19, 2002.

^{†‡}School of Computer Science, Tel Aviv University, Tel Aviv, Israel. Supported by the Deutsch Institute

Objects Based Change Detection in a Pair of Gray Images

Amir Averbuch Ofer Miller Arie Pikaz

School of Computer Science
Tel Aviv University
Tel Aviv 69978, Israel

The goal of the presented change detection algorithm is to extract the objects that appear only in one of two registered images. A typical application is surveillance, where a scene is sampled at different times. In this paper we assume a significant illumination difference between the two images. For example, one image may be captured during daylight while the other image may be captured at night with infrared device. By analyzing the connectivity along gray-levels, all the blobs that are candidates to be classified as 'change' are extracted from both images. Then, the candidate blobs from both images are analyzed. A Blob from one image that has no matched blob in the other image is considered as a 'change'. The algorithm was found to be reliable, fast, accurate, and robust even under significant changes in illumination. The performance of the algorithm is demonstrated using real world images. The worst-case time complexity of the algorithm is almost linear in the image size. Therefore, it is suitable for real-time applications.

Supported by the Deutsch Institute.

Hybrid Motion Planning: Coordinating Two Discs Moving Among Polygonal Obstacles in the Plane

Shai Hirsch Dan Halperin

Abstract

The basic motion-planning problem is to plan a collision-free motion for an object moving among obstacles between free initial and goal positions, or to determine that no such motion exists. The basic problem as well as numerous variants of it have been intensively studied over the past two decades yielding a wealth of results and techniques, both theoretical and practical. In this paper, we propose a novel approach to motion planning, *hybrid motion planning*, in which we integrate complete solutions along with probabilistic roadmap (PRM) methods in order to combine their strengths and offset their weaknesses. We incorporate robust tools, that have not been available before, in order to implement the complete solutions. We exemplify our approach in the case of two discs moving among polygonal obstacles in the plane. The planner we present easily solves problems where a narrow passage in the workspace can be arbitrarily small. Our planner is also capable of providing correct nontrivial “no” answers, namely it can, for some queries, detect the situation where no solution exists. We envision our planner not as a total solution but rather as a new tool that cooperates with existing planners. We demonstrate the advantages and shortcomings of our planner with experimental results.

Presented at the Fifth International Workshop on Algorithmic Foundations of Robotics, Nice, France, December 15–17, 2002.

School of Computer Science, Tel Aviv University, Tel Aviv, Israel. Supported by the Deutsch Institute.

Automatic acquisition and efficient representation of syntactic structures

Zach Solan Shimon Edelman David Horn Eytan Ruppin

Abstract

The principle of complementary distributions, according to which morphemes that occur in identical contexts belong, in some sense, to the same category, has been advanced as a means for extracting syntactic structures from corpus data. We extend this principle by applying it recursively, and by using mutual information for estimating category coherence. The resulting model learns, in an unsupervised fashion, highly structured, distributed representations of syntactic knowledge from corpora. It also exhibits promising behavior in tasks usually thought to require representations anchored in a grammar, such as systematicity.

School of Physics Tel-Aviv University, Tel-Aviv, Israel. Department of Psychology, Cornell University Ithaca, NY. School of Physics Tel-Aviv University, Tel-Aviv, Israel. School of Computer Science, Tel Aviv University, Tel Aviv, Israel. Supported by the Deutsch Institute.

Bounded-distortion Piecewise Mesh Parameterization

Olga Sorkine Daniel Cohen-Or Rony Goldenthal Dani Lischinski

Abstract

Many computer graphics operations, such as texture mapping, 3D painting, remeshing, mesh compression, and digital geometry processing, require finding a low-distortion parameterization for irregular connectivity triangulations of arbitrary genus 2-manifolds. This paper presents a simple and fast method for computing parameterizations with strictly bounded distortion. The new method operates by flattening the mesh onto a region of the 2D plane. To comply with the distortion bound, the mesh is automatically cut and partitioned on-the-fly. The method guarantees avoiding global and local self-intersections, while attempting to minimize the total length of the introduced seams.

To our knowledge, this is the first method to compute the mesh partitioning and the parameterization simultaneously and entirely automatically, while providing guaranteed distortion bounds. Our results on a variety of objects demonstrate that the method is fast enough to work with large complex irregular meshes in interactive applications.

Presented at IEEE Visualization 2002, Boston, Massachusetts, USA, October 27 – November 1, 2002.

Supported by the Deutsch Institute. School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel.
School of Engineering and Computer Science, The Hebrew University of Jerusalem, Jerusalem, Israel.

Seeing People in the Dark: Face recognition in Infrared Images

Gil Friedrich[†] Yehezkel Yeshurun[‡]

Abstract

An IR image of the human face presents its unique heat-signature and can be used for recognition. The characteristics of IR images maintain advantages over visible light images, and can be used to improve algorithms of human face recognition in several aspects. IR images are obviously invariant under extreme lighting conditions (including complete darkness). The main findings of this research are that IR face images are less effected by changes of pose or facial expression and enable a simple method for detection of facial features. In this paper we explore several aspects of face recognition in IR images. First, we compare the effect of varying environment conditions over IR and visible light images through a case study. Finally, we propose a method for automatic face recognition in IR images, through which we use a preprocessing algorithm for detecting facial elements, and show the applicability of commonly used face recognition methods in the visible light domain.

Presented at the 2nd International Workshop of Biologically Motivated Computer Vision, Teubingen, Germany, November 22-24, 2002.

[†]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel. Research supported by the Deutsch Institute.

[‡]Dept. of Computer Science, Tel Aviv University, Tel Aviv, Israel

Testing Subgraphs in Directed Graphs

Noga Alon[†] Asaf Shapira[‡]

Abstract

Let H be a fixed directed graph on h vertices, let G be a directed graph on n vertices and suppose that at least ϵn^2 edges have to be deleted from it to make it H -free. We show that in this case G contains at least $f(\epsilon, H)n^h$ copies of H . This is proved by establishing a directed version of Szemerédi's regularity lemma, and implies that for every H there is a *one-sided error* property tester whose query complexity is bounded by a function of ϵ only for testing the property P_H of being H -free.

As is common with applications of the undirected regularity lemma, here too the function $1/f(\epsilon, H)$ is an extremely fast growing function in ϵ . We therefore further prove the following precise characterization of all the digraphs H , for which $f(\epsilon, H)$ has a polynomial dependency on ϵ : a *homomorphism* $\varphi : V(H) \mapsto V(K)$, from a digraph H to K , is a function that satisfies $(u, v) \in E(H) \Rightarrow (\varphi(u), \varphi(v)) \in E(K)$. The *core* of a digraph H is the smallest subgraph K of H , for which there is a homomorphism from H to K . We show that for a connected H , $f(\epsilon, H)$ has a polynomial dependency on $1/\epsilon$, **if and only if** the core of H is either an oriented tree or a directed cycle of length 2.

Presented at the 35th Annual ACM Symposium on Theory of Computing (STOC), San Diego, CA, USA, June 9-11, 2003.

[†]Schools of Mathematics and Computer Science, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel. Email: noga@math.tau.ac.il.

[‡]School of Computer Science, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel. Email: asafico@math.tau.ac.il. Research supported by the Deutsch Institute

Adaptive AIMD Congestion Control

Alex Kesselman[†] Yishay Mansour

Abstract

The main objectives of a congestion control algorithm are high bandwidth utilization, fairness and responsiveness in changing environment. However, these objectives are contradicting in particular situations since the algorithm has to constantly probe available bandwidth, which may affect its stability. This paper proposes a novel congestion control algorithm that achieves high bandwidth utilization providing fairness among competing connections and, on the other hand, is sufficiently responsive to changes of available bandwidth. The main idea of the algorithm is to use *adaptive* setting for the additive increase/multiplicative decrease (AIMD) congestion control scheme, where parameters may change dynamically, with respect to the current network conditions.

Presented at Twenty-Second ACM Symposium on Principles of Distributed Computing, Boston, Massachusetts, July 13-16, 2003.

[†]School of Computer Science, Tel Aviv University, Tel Aviv, Israel. Research supported by the Deutsch Institute.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv, Israel

Scheduling Policies for CIOQ Switches

Alex Kesselman[†] Adi Rosén[‡]

Abstract

Combined input and output queued (CIOQ) architectures with a moderate fabric *speedup* $S > 1$ have come to play a major role in the design of high performance switches. The switch policy that controls such switches must consist of two components. A buffer management policy that controls admission to buffers, and a scheduling policy that schedules the transfer of packets from input buffers to output buffers. The goal of the switch policy is to maximize the throughput of the switch. When all packets have a uniform value (or importance), this corresponds to the number of packets sent from the switch. When packets have variable values, this corresponds to the total value of the packets sent.

We mainly consider switches with virtual output queuing (VOQ) at the inputs. For the case of packets with uniform values we present a switch policy that is 3-competitive for any speedup. For the case of packets with variable values we propose two preemptive switch policies. One achieves a competitive ratio of $4S$, and the other achieves a competitive ratio of $8 \min(k, 2 \log \alpha)$, where k is the number of distinct packet values and α is the ratio between the largest and smallest values.

Presented at Fifteenth ACM Symposium on Parallelism in Algorithms and Architectures, San Diego, California, USA, June 7–9, 2003.

[†]School of Computer Science, Tel Aviv University, Tel Aviv, Israel. Research supported by the Deutsch Institute

[‡]Faculty of Computer Science, The Technion, Haifa, Israel