BRICS Research Series, Abstracts, 2001

August 5, 2003

This document is also available as *PostScript* and *DVI*.

References

[RS-01-55] *PostScript*, *PDF*, *DVI*.
Daniel Damian and Olivier Danvy.
A Simple CPS Transformation of Control-Flow Information.
December 2001.
18 pp. Appears in *Logic Journal of the IGPL*, 10(2):501–515, 2002.
Abstract: We build on Danvy and Nielsen's first-order program transformation into continuation-passing style (CPS) to design a new CPS transformation of flow information that is simpler and more efficient than what has been presented in previous work. The key to simplicity and efficiency is that our CPS transformation constructs the flow information in one go, instead of first computing an intermediate result and then exploiting it to construct the flow information.

More precisely, we show how to compute control-flow information for CPS-transformed programs from control-flow information for direct-style programs and vice-versa. As a corollary, we confirm that CPS transformation has no effect on the control-flow information obtained by constraint-based control-flow analysis. The transformation has immediate applications in assessing the effect of the CPS transformation over other analyses such as, for instance, binding-time analysis.

[RS-01-54] *PostScript*, *PDF*, *DVI*.
Daniel Damian and Olivier Danvy.
Syntactic Accidents in Program Analysis: On the Impact of the CPS Transformation.
December 2001.
41 pp. To appear in the *Journal of Functional Programming*. This report supersedes the earlier BRICS report RS-00-15.
Abstract: We show that a non-duplicating transformation into continuation-passing style (CPS) has no effect on control-flow analysis, a positive effect on binding-time analysis for traditional partial evaluation, and no effect on binding-time analysis for continuation-based partial evaluation: a monovariant control-flow analysis yields equivalent results on a direct-style program

and on its CPS counterpart, a monovariant binding-time analysis yields less precise results on a direct-style program than on its CPS counterpart, and an enhanced monovariant binding-time analysis yields equivalent results on a direct-style program and on its CPS counterpart. Our proof technique amounts to constructing the CPS counterpart of flow information and of binding times.

Our results formalize and confirm a folklore theorem about traditional binding-time analysis, namely that CPS has a positive effect on binding times. What may be more surprising is that the benefit does not arise from a standard refinement of program analysis, as, for instance, duplicating continuations.

The present study is symptomatic of an unsettling property of program analyses: their quality is unpredictably vulnerable to syntactic accidents in source programs, i.e., to the way these programs are written. More reliable program analyses require a better understanding of the effect of syntactic change.

[RS-01-53] *PostScript*, *PDF*, *DVI*.
Zoltn sik and Masami Ito.
Temporal Logic with Cyclic Counting and the Degree of Aperiodicity of Finite Automata.
December 2001.
31 pp.
Abstract: We define the degree of aperiodicity of finite automata and show that for every set M of positive integers, the class $\mathbf{QA}_M$ of finite automata whose degree of aperiodicity belongs to the division ideal generated by M is closed with respect to direct products, disjoint unions, subautomata, homomorphic images and renamings. These closure conditions define q-varieties of finite automata. We show that q-varieties are in a one-to-one correspondence with literal varieties of regular languages. We also characterize $\mathbf{QA}_M$ as the cascade product of a variety of counters with the variety of aperiodic (or counter-free) automata. We then use the notion of degree of aperiodicity to characterize the expressive power of first-order logic and temporal logic with cyclic counting with respect to any given set M of moduli. It follows that when M is finite, then it is decidable whether a regular language is definable in first-order or temporal logic with cyclic counting with respect to moduli in M.

[RS-01-52] *PostScript*, *PDF*, *DVI*.
Jens Groth.
Extracting Witnesses from Proofs of Knowledge in the Random Oracle Model.
December 2001.
23 pp.
Abstract: We prove that a 3-move interactive proof system with the special soundness property made non-interactive by applying the Fiat-Shamir heuristic is almost a non-interactive proof of knowledge in the random oracle model. In an application of

the result we demonstrate that the Damgrd-Jurik voting scheme based on homomorphic threshold encryption is secure against a nonadaptive adversary according to Canetti's definition of multiparty computation security.

[RS-01-51] *PostScript*, *PDF*, *DVI.*
Ulrich Kohlenbach.
On Weak Markov's Principle.
December 2001.
10 pp. Appears in *Mathematical Logic Quarterly*, 48(suppl. 1):59–65, 2002.

Abstract: We show that the so-called weak Markov's principle (WMP) which states that every pseudo-positive real number is positive is underivable in $\mathcal{T}^\omega :=$E-HA$^\omega$+AC. Since $\mathcal{T}^\omega$ allows to formalize (at least large parts of) Bishop's constructive mathematics this makes it unlikely that WMP can be proved within the framework of Bishop-style mathematics (which has been open for about 20 years). The underivability even holds if the ineffective schema of full comprehension (in all types) for negated formulas (in particular for $\exists$-free formulas) is added which allows to derive the law of excluded middle for such formulas.

[RS-01-50] *PostScript*, *PDF*, *DVI.*
Jiř Srba.
Note on the Tableau Technique for Commutative Transition Systems.
December 2001.
19 pp. Appears in Nielsen and Engberg, editors, *Foundations of Software Science and Computation Structures*, FoSSaCS '02 Proceedings, LNCS 2303, 2002, pages 387–401.

Abstract: We define a class of transition systems called effective commutative transition systems (ECTS) and show, by generalising a tableau-based proof for BPP, that bisimilarity between any two states of such a transition system is decidable. It gives a general technique for extending decidability borders of strong bisimilarity for a wide class of infinite-state transition systems. This is demonstrated for several process formalisms, namely BPP process algebra, lossy BPP processes, BPP systems with interrupt and timed-arc BPP nets.

[RS-01-49] *PostScript*, *PDF*, *DVI.*
Olivier Danvy and Lasse R. Nielsen.
A First-Order One-Pass CPS Transformation.
December 2001.
21 pp. To appear in *Theoretical Computer Science*. Extended version of a paper appearing in Nielsen and Engberg, editors, *Foundations of Software Science and Computation Structures*, FoSSaCS '02 Proceedings, LNCS 2303, 2002, pages 98–113.

Abstract: We present a new transformation of call-by-value lambda-terms into continuation-passing style (CPS). This transformation operates in one pass and is both compositional and first-order. Because it operates in one pass, it directly yields compact CPS programs that are comparable to what one would write by hand. Because it is compositional, it allows proofs by structural induction. Because it is first-order, reasoning about it does not require the use of a logical relation.

This new CPS transformation connects two separate lines of research. It has already been used to state a new and simpler correctness proof of a direct-style transformation, and to develop a new and simpler CPS transformation of control-flow information.

[RS-01-48] *PostScript*, *PDF*, *DVI*.
Mogens Nielsen and Frank D. Valencia.
Temporal Concurrent Constraint Programming: Applications and Behavior.
December 2001.
36 pp. Appears in Brauer, Ehrig, Karhumki and Salomaa, editors, *Formal and Natural Computing*, LNCS 2300, 2001, pages 298–321.

Abstract: The *ntcc* calculus is a model of non-deterministic temporal concurrent constraint programming. In this paper we study behavioral notions for this calculus. In the underlying computational model, concurrent constraint processes are executed in discrete time intervals. The behavioral notions studied reflect the reactive interactions between concurrent constraint processes and their environment, as well as internal interactions between individual processes. Relationships between the suggested notions are studied, and they are all proved to be decidable for a substantial fragment of the calculus. Furthermore, the expressive power of this fragment is illustrated by examples.

[RS-01-47] *PostScript*, *PDF*, *DVI*.
Jesper Buus Nielsen.
Non-Committing Encryption is Too Easy in the Random Oracle Model.
December 2001.
20 pp.

Abstract: The non-committing encryption problem arises in the setting of adaptively secure cryptographic protocols, as the task of implementing secure channels. We prove that in the random oracle model, where the parties have oracle access to a uniformly random function, non-committing encryption can be implemented efficiently using any trapdoor permutation.

We also prove that no matter how the oracle is instantiated in practice the resulting scheme will never be non-committing, and we give a short discussion of the random oracle model in light of this.

[RS-01-46] *PostScript*, *PDF*, *DVI*.
Lars Kristiansen.
The Implicit Computational Complexity of Imperative Programming Languages.
November 2001.
46 pp.

Abstract: During the last decade Cook, Bellantoni, Leivant and others have developed the theory of implicit computational complexity, i.e. the theory of predicative recursion, tiered definition schemes, etcetera. We extend and modify this theory to work in a context of imperative programming languages, and characterise complexity classes like P, LINSPACE, PSPACE and the classes in the Grzegorczyk hiearchy. Our theoretical framework seems promising with respect to applications in engineering.

[RS-01-45] *PostScript*, *PDF*.
Ivan B. Damgrd and Gudmund Skovbjerg Frandsen.
An Extended Quadratic Frobenius Primality Test with Average Case Error Estimates.
November 2001.
43 pp.
Abstract: We present an Extended Quadratic Frobenius Primality Test (EQFT), which is related to the Miller-Rabin test and the Quadratic Frobenius test (QFT) by Grantham. EQFT is well-suited for generating large, random prime numbers since on a random input number, it takes time about equivalent to 2 Miller-Rabin tests, but has much smaller error probability. EQFT extends QFT by verifying additional algebraic properties related to the existence of elements of order 3 and 4. We obtain a simple closed expression that upper bounds the probability of acceptance for any input number. This in turn allows us to give strong bounds on the average-case behaviour of the test: consider the algorithm that repeatedly chooses random odd k bit numbers, subjects them to t iterations of our test and outputs the first one found that passes all tests. We obtain numeric upper bounds for the error probability of this algorithm as well as a general closed expression bounding the error. For instance, it is at most 2^{-143} for $k = 500, t = 2$. Compared to earlier similar results for the Miller-Rabin test, the results indicates that our test in the average case has the effect of 9 Miller-Rabin tests, while only taking time equivalent to about 2 such tests. We also give bounds for the error in case a prime is sought by incremental search from a random starting point. While EQFT is slower than the average case on a small set of inputs, we present a variant that is always fast, i.e. takes time about 2 Miller-Rabin tests. The variant has slightly larger worst case error probability than EQFT, but still improves on previous proposed tests.

[RS-01-44] *PostScript*, *PDF*.
M. Oliver Mller, Harald Rue, and Maria Sorea.
Predicate Abstraction for Dense Real-Time Systems.
November 2001.
27 pp. Appears in Asarin, Maler and Yovine, editors, *Workshop on Theory and Practice of Timed Systems*, TPTS '02 Proceedings, ENTCS 65(6), 2002.
Abstract: We propose predicate abstraction as a means for verifying a rich class of safety and liveness properties for dense real-

time systems. First, we define a restricted semantics of timed systems which is observationally equivalent to the standard semantics in that it validates the same set of μ-calculus formulas without a next-step operator. Then, we recast the model checking problem $\mathcal{S} \models \varphi$ for a timed automaton $\mathcal{S}$ and a μ-calculus formula φ in terms of predicate abstraction. Whenever a set of abstraction predicates forms a so-called *basis*, the resulting abstraction is strongly preserving in the sense that $\mathcal{S}$ validates φ iff the corresponding finite abstraction validates this formula φ. Now, the abstracted system can be checked using familiar μ-calculus model checking.

Like the region graph construction for timed automata, the predicate abstraction algorithm for timed automata usually is prohibitively expensive. In many cases it suffices to compute an approximation of a finite bisimulation by using only a subset of the basis of abstraction predicates. Starting with some coarse abstraction, we define a finite sequence of refined abstractions that converges to a strongly preserving abstraction. In each step, new abstraction predicates are selected nondeterministically from a finite basis. Counterexamples from failed μ-calculus model checking attempts can be used to heuristically choose a small set of new abstraction predicates for refining the abstraction.

[RS-01-43] *PostScript*, *PDF*, *DVI*.
Ivan B. Damgrd and Jesper Buus Nielsen.
From Known-Plaintext Security to Chosen-Plaintext Security.
November 2001.
18 pp.

Abstract: We present a new encryption mode for block ciphers. The mode is efficient and is secure against chosen-plaintext attack (CPA) already if the underlying symmetric cipher is secure against known-plaintext attack (KPA). We prove that known (and widely used) encryption modes as CBC mode and counter mode do not have this property. In particular, we prove that CBC mode using a KPA secure cipher is KPA secure, but need not be CPA secure, and we prove that counter mode using a KPA secure cipher need not be even KPA secure. The analysis is done in a concrete security framework.

[RS-01-42] *PostScript*, *PDF*, *DVI*.
Zoltn sik and Werner Kuich.
Rationally Additive Semirings.
November 2001.
11 pp. Appears in *Journal of Universal Computer Science*, 8(2):173–183, 2002.

Abstract: We define rationally additive semirings that are a generalization of (ω)complete and (ω-)continuous semirings. We prove that every rationally additive semiring is an iteration semiring. Moreover, we characterize the semirings of rational power series with coefficients in N_∞, the semiring of natural numbers equipped with a top element, as the free rationally additive semirings.

[RS-01-41] *PostScript*, *PDF*, *DVI*.

Ivan B. Damgrd and Jesper Buus Nielsen.

Perfect Hiding and Perfect Binding Universally Composable Commitment Schemes with Constant Expansion Factor.

October 2001.

43 pp. Appears in Yung, editor, *Advances in Cryptology: 22nd Annual International Cryptology Conference*, CRYPTO '02 Proceedings, LNCS 2442, 2002, pages 581–596.

Abstract: Canetti and Fischlin have recently proposed the security notion *universal composability* for commitment schemes and provided two examples. This new notion is very strong. It guarantees that security is maintained even when an unbounded number of copies of the scheme are running concurrently, also it guarantees non-malleability, resilience to selective decommitment, and security against adaptive adversaries. Both of their schemes uses $\Theta(k)$ bits to commit to one bit and can be based on the existence of trapdoor commitments and non-malleable encryption.

We present new universally composable commitment schemes based on the Paillier cryptosystem and the Okamoto-Uchiyama cryptosystem. The schemes are efficient: to commit to k bits, they use a constant number of modular exponentiations and communicates $O(k)$ bits. Further more the scheme can be instantiated in either perfectly hiding or perfectly binding versions. These are the first schemes to show that constant expansion factor, perfect hiding, and perfect binding can be obtained for universally composable commitments.

We also show how the schemes can be applied to do efficient zero-knowledge proofs of knowledge that are universally composable.

[RS-01-40] *PostScript*, *PDF*, *DVI*.

Daniel Damian and Olivier Danvy.

CPS Transformation of Flow Information, Part II: Administrative Reductions.

October 2001.

9 pp. To appear in the *Journal of Functional Programming*. Superseeded by the BRICS report RS-02-36.

Abstract: We characterize the impact of a linear beta-reduction on the result of a control-flow analysis. (By "a linear beta-reduction" we mean the beta-reduction of a linear lambda-abstraction, i.e., of a lambda-abstraction whose parameter occurs exactly once in its body.)

As a corollary, we consider the administrative reductions of a Plotkin-style transformation into continuation-passing style (CPS), and how they affect the result of a constraint-based control-flow analysis and in particular the least element in the space of solutions. We show that administrative reductions preserve the least solution. Since we know how to construct least solutions, preservation of least solutions solves a problem that was left open in Palsberg and Wand's paper "CPS Transformation of Flow Information."

Therefore, together, Palsberg and Wand's article "CPS Transformation of Flow Information" and the present article show how to map, in linear time, the least solution of the flow constraints of a program into the least solution of the flow constraints of the CPS counterpart of this program, after administrative reductions. Furthermore, we show how to CPS transform control-flow information in one pass.

[RS-01-39] *PostScript*, *PDF*, *DVI*.
Olivier Danvy and Mayer Goldberg.
There and Back Again.
October 2001.
14 pp. To appear in the proceedings of the 2002 ACM SIGPLAN International Conference on Functional Programming.

Abstract: We illustrate a variety of programming problems that seemingly require two separate list traversals, but that can be efficiently solved in one recursive descent, without any other auxiliary storage but what can be expected from a control stack. The idea is to perform the second traversal when returning from the first.

This programming technique yields new solutions to traditional problems. For example, given a list of length $2n$ or $2n+1$, where n is unknown, we can detect whether this list is a palindrome in $n+1$ recursive calls and no heap allocation.

[RS-01-38] *PostScript*, *PDF*, *DVI*.
Zoltn sik.
Free De Morgan Bisemigroups and Bisemilattices.
October 2001.
13 pp. To appear in the journal *Algebra Colloquium*.

Abstract: We give a geometric representation of free De Morgan bisemigroups, free commutative De Morgan bisemigroups and free De Morgan bisemilattices, using labeled graphs.

[RS-01-37] *PostScript*, *PDF*, *DVI*.
Ronald Cramer and Victor Shoup.
Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption.
October 2001.
34 pp.

Abstract: We present several new and fairly practical public-key encryption schemes and prove them secure against adaptive chosen ciphertext attack. One scheme is based on Paillier's Decision Composite Residuosity (DCR) assumption, while another is based in the classical Quadratic Residuosity (QR) assumption. The analysis is in the standard cryptographic model, i.e., the security of our schemes does not rely on the Random Oracle model.

We also introduce the notion of a universal hash proof system. Essentially, this is a special kind of non-interactive zero-knowledge proof system for an NP language. We do not show that universal hash proof systems exist for all NP languages, but we do show how to construct very efficient universal hash proof systems for a general class of group-theoretic language membership problems.

Given an efficient universal hash proof system for a language with certain natural cryptographic indistinguishability properties, we show how to construct an efficient public-key encryption schemes secure against adaptive chosen ciphertext attack in the standard model. Our construction only uses the universal hash proof system as a primitive: no other primitives are required, although even more efficient encryption schemes can be obtained by using hash functions with appropriate collision-resistance properties. We show how to construct efficient universal hash proof systems for languages related to the DCR and QR assumptions. From these we get corresponding public-key encryption schemes that are secure under these assumptions. We also show that the Cramer-Shoup encryption scheme (which up until now was the only practical encryption scheme that could be proved secure against adaptive chosen ciphertext attack under a reasonable assumption, namely, the Decision Diffie-Hellman assumption) is also a special case of our general theory.

[RS-01-36] *PostScript*, *PDF*.
Gerth Stlting Brodal, Rolf Fagerberg, and Riko Jacob.
Cache Oblivious Search Trees via Binary Trees of Small Height.
October 2001.
20 pp. Appears in Eppstein, editor, *The Thirteenths Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '02 Proceedings, 2002, pages 39–48.

Abstract: We propose a version of cache oblivious search trees which is simpler than the previous proposal of Bender, Demaine and Farach-Colton and has the same complexity bounds. In particular, our data structure avoids the use of weight balanced B-trees, and can be implemented as just a single array of data elements, without the use of pointers. The structure also improves space utilization.

For storing n elements, our proposal uses $(1 + \varepsilon)n$ times the element size of memory, and performs searches in worst case $O(\log_B n)$ memory transfers, updates in amortized $O((\log^2 n)/(\varepsilon B))$ memory transfers, and range queries in worst case $O(\log_B n + k/B)$ memory transfers, where k is the size of the output.

The basic idea of our data structure is to maintain a dynamic binary tree of height $\log n + O(1)$ using existing methods, embed this tree in a static binary tree, which in turn is embedded in an array in a cache oblivious fashion, using the van Emde Boas layout of Prokop.

We also investigate the practicality of cache obliviousness in the area of search trees, by providing an empirical comparison of different methods for laying out a search tree in memory.

The source code of the programs, our scripts and tools, and the data we present here are available online under ftp.brics.dk/RS/01/36/Experiments/.

[RS-01-35] *PostScript*, *PDF*, *DVI*.
Mayer Goldberg.
A General Schema for Constructing One-Point Bases in the Lambda Calculus.
September 2001.
8 pp.
Abstract: In this paper, we present a schema for constructing one-point bases for recursively enumerable sets of lambda terms. The novelty of the approach is that we make no assumptions about the terms for which the one-point basis is constructed: They need not be combinators and they may contain constants and free variables. The significance of the construction is twofold: In the context of the lambda calculus, it characterises one-point bases as ways of "packaging" sets of terms into a single term; And in the context of realistic programming languages, it implies that we can define a single procedure that generates any given recursively enumerable set of procedures, constants and free variables in a given programming language.

[RS-01-34] *PostScript*, *PDF*.
Flemming Friche Rodler and Rasmus Pagh.
Fast Random Access to Wavelet Compressed Volumetric Data Using Hashing.
August 2001.
31 pp. To appear in *ACM Transactions on Graphics*.
Abstract: We present a new approach to lossy storage of the coefficients of wavelet transformed data. While it is common to store the coefficients of largest magnitude (and let all other coefficients be zero), we allow a slightly different set of coefficients to be stored. This brings into play a recently proposed hashing technique that allows space efficient storage and very efficient retrieval of coefficients. Our approach is applied to compression of volumetric data sets. For the "Visible Man" volume we obtain up to 80% improvement in compression ratio over previously suggested schemes. Further, the time for accessing a random voxel is quite competitive.

[RS-01-33] *PostScript*, *PDF*.
Rasmus Pagh and Flemming Friche Rodler.
Lossy Dictionaries.
August 2001.
14 pp. Short version appears in Meyer auf der Heide, editor, *9th Annual European Symposium on Algorithms*, ESA '01 Proceedings, LNCS 2161, 2001, pages 300–311.
Abstract: Bloom filtering is an important technique for space efficient storage of a conservative approximation of a set S. The set stored may have up to some specified number of "false positive" members, but all elements of S are included. In this paper we

consider *lossy dictionaries* that are also allowed to have "false negatives", i.e., leave out elements of S. The aim is to maximize the weight of included keys within a given space constraint. This relaxation allows a very fast and simple data structure making almost optimal use of memory. Being more time efficient than Bloom filters, we believe our data structure to be well suited for replacing Bloom filters in some applications. Also, the fact that our data structure supports information associated to keys paves the way for new uses, as illustrated by an application in lossy image compression.

[RS-01-32] *PostScript*, *PDF*.
Rasmus Pagh and Flemming Friche Rodler.
Cuckoo Hashing.
August 2001.
21 pp. Short version appears in Meyer auf der Heide, editor, *9th Annual European Symposium on Algorithms*, ESA '01 Proceedings, LNCS 2161, 2001, pages 121–133.

Abstract: We present a simple and efficient dictionary with worst case constant lookup time, equaling the theoretical performance of the classic dynamic perfect hashing scheme of Dietzfelbinger et al. (*Dynamic perfect hashing: Upper and lower bounds. SIAM J. Comput., 23(4):738–761, 1994*). The space usage is similar to that of binary search trees, i.e., three words per key on average. The practicality of the scheme is backed by extensive experiments and comparisons with known methods, showing it to be quite competitive also in the average case.

[RS-01-31] *PostScript*, *PDF*, *DVI*.
Olivier Danvy and Lasse R. Nielsen.
Syntactic Theories in Practice.
July 2001.
37 pp. Extended version of an article to appear in the informal proceedings of the *Second International Workshop on Rule-Based Programming*, RULE 2001 (Firenze, Italy, September 4, 2001). Superseeded by the BRICS report RS-02-4.

Abstract: The evaluation function of a syntactic theory is canonically defined as the transitive closure of (1) decomposing a program into an evaluation context and a redex, (2) contracting this redex, and (3) plugging the result in the context. Directly implementing this evaluation function therefore yields an interpreter with a quadratic time factor over its input. We present sufficient conditions over a syntactic theory to circumvent this quadratic factor, and illustrate the method with two programming-language interpreters and a transformation into continuation-passing style (CPS). In particular, we mechanically change the time complexity of this CPS transformation from quadratic to linear.

[RS-01-30] *PostScript*, *PDF*, *DVI*.
Lasse R. Nielsen.
A Selective CPS Transformation.

July 2001.

24 pp. Appears in Brookes and Mislove, editors, *27th Annual Conference on the Mathematical Foundations of Programming Semantics*, MFPS '01 Proceedings, ENTCS 45, 2001. A preliminary version appeared in Brookes and Mislove, editors, *17th Annual Conference on Mathematical Foundations of Programming Semantics*, MFPS '01 Preliminary Proceedings, BRICS Notes Series NS-01-2, 2001, pages 201–222.

Abstract: The CPS transformation makes all functions continuation-passing, uniformly. Not all functions, however, need continuations: they only do if their evaluation includes computational effects. In this paper we focus on control operations, in particular "call with current continuation" and "throw". We characterize this involvement as a control effect and we present a selective CPS transformation that makes functions and expressions continuation-passing if they have a control effect, and that leaves the rest of the program in direct style. We formalize this selective CPS transformation with an operational semantics and a simulation theorem la Plotkin.

[RS-01-29] *PostScript*, *PDF*, *DVI*.

Olivier Danvy, Bernd Grobauer, and Morten Rhiger.

A Unifying Approach to Goal-Directed Evaluation.

July 2001.

23 pp. Appears in *New Generation Computing*, 20(1):53–73, November 2001. A preliminary version appeared in Taha, editor, *2nd International Workshop on Semantics, Applications, and Implementation of Program Generation*, SAIG '01 Proceedings, LNCS 2196, 2001, pages 108–125.

Abstract: Goal-directed evaluation, as embodied in Icon and Snobol, is built on the notions of backtracking and of generating successive results, and therefore it has always been something of a challenge to specify and implement. In this article, we address this challenge using computational monads and partial evaluation.

We consider a subset of Icon and we specify it with a monadic semantics and a list monad. We then consider a spectrum of monads that also fit the bill, and we relate them to each other. For example, we derive a continuation monad as a Church encoding of the list monad. The resulting semantics coincides with Gudeman's continuation semantics of Icon.

We then compile Icon programs by specializing their interpreter (i.e., by using the first Futamura projection), using type-directed partial evaluation. Through various back ends, including a run-time code generator, we generate ML code, C code, and OCaml byte code. Binding-time analysis and partial evaluation of the continuation-based interpreter automatically give rise to C programs that coincide with the result of Proebsting's optimized compiler.

[RS-01-28] *PostScript*, *PDF*, *DVI*.

Luca Aceto, Zoltn sik, and Anna Inglfsdttir.
A Fully Equational Proof of Parikh's Theorem.
June 2001.
28 pp. Appears in *RAIRO, Theoretical Informatics and Applications*, 36(2):129–153, 2002, special issue devoted to selected papers from FICS 2001 (A. Labella ed.).
Abstract: We show that the validity of Parikh's theorem for context-free languages depends only on a few equational properties of least pre-fixed points. Moreover, we exhibit an infinite basis of μ-term equations of continuous commutative idempotent semirings.

[RS-01-27] *PostScript*, *PDF*, *DVI.*
Mario Jos Cccamo and Glynn Winskel.
A Higher-Order Calculus for Categories.
June 2001.
24 pp. Appears in Boulton and Jackson, editors, *Theorem Proving in Higher Order Logics: 14th International Conference*, TPHOLs '01 Proceedings, LNCS 2152, 2001, pages 136–153.
Abstract: A calculus for a fragment of category theory is presented. The types in the language denote categories and the expressions functors. The judgements of the calculus systematise categorical arguments such as: an expression is functorial in its free variables; two expressions are naturally isomorphic in their free variables. There are special binders for limits and more general ends. The rules for limits and ends support an algebraic manipulation of universal constructions as opposed to a more traditional diagrammatic approach. Duality within the calculus and applications in proving continuity are discussed with examples. The calculus gives a basis for mechanising a theory of categories in a generic theorem prover like Isabelle.

[RS-01-26] *PostScript*, *PDF*, *DVI.*
Ulrik Frendrup and Jesper Nyholm Jensen.
A Complete Axiomatization of Simulation for Regular CCS Expressions.
June 2001.
18 pp.
Abstract: This paper gives axiomatizations of strong and weak simulation over regular CCS expressions. The proof of completeness of the axiomatization of strong simulation is inspired by Milner's proof of completeness of his axiomatization for strong equivalence over regular CCS expressions. Soundness and completeness of the axiomatization for weak simulation is easily deduced from the corresponding result for the axiomatization of strong simulation over regular CCS expressions.

[RS-01-25] *PostScript*, *PDF*, *DVI.*
Bernd Grobauer.
Cost Recurrences for DML Programs.
June 2001.

51 pp. Extended version of a paper appearing in Leroy, editor, *Proceedings of the 6th ACM SIGPLAN International Conference on Functional Programming*, 2001, pages 253–264.

Abstract: A cost recurrence describes an upper bound for the running time of a program in terms of the size of its input. Finding cost recurrences is a frequent intermediate step in complexity analysis, and this step requires an abstraction from data to data size. In this article, we use information contained in dependent types to achieve such an abstraction: Dependent ML (DML), a conservative extension of ML, provides dependent types that can be used to associate data with size information, thus describing a possible abstraction. We systematically extract cost recurrences from first-order DML programs, guiding the abstraction from data to data size with information contained in DML type derivations.

[RS-01-24] *PostScript*, *PDF*, *DVI*.
Zoltn sik and Zoltn L. Nmeth.
Automata on Series-Parallel Biposets.
June 2001.
15 pp. Appears in Kuich, Rozenberg and Salomaa, editors, *5th International Conference, Developments in Language Theory*, DLT '01 Revised Papers, LNCS 2295, 2002, pages 217–227. Superseded by the later BRICS report RS-02-44.

Abstract: We provide the basics of a 2-dimensional theory of automata on series-parallel biposets. We define recognizable, regular and rational sets of series-parallel biposets and study their relationship. Moreover, we relate these classes to languages of series-parallel biposets definable in monadic second-order logic.

[RS-01-23] *PostScript*, *PDF*, *DVI*.
Olivier Danvy and Lasse R. Nielsen.
Defunctionalization at Work.
June 2001.
45 pp. Extended version of an article appearing in Sndergaard, editor, *3rd International Conference on Principles and Practice of Declarative Programming*, PPDP '01 Proceedings, 2001, pages 162–174.

Abstract: We study practical applications of Reynolds's defunctionalization technique, which is a whole-program transformation from higher-order to first-order functional programs. This study leads us to discover new connections between seemingly unrelated higher-order and first-order specifications and their correctness proofs. We thus perceive defunctionalization both as a springboard and as a bridge: as a springboard for discovering new connections between the first-order world and the higher-order world; and as a bridge for transferring existing results between first-order and higher-order settings.

[RS-01-22] *PostScript*, *PDF*, *DVI*.
Zoltn sik.

The Equational Theory of Fixed Points with Applications to Generalized Language Theory.
June 2001.
21 pp. Appears in Kuich, Rozenberg and Salomaa, editors, *5th International Conference, Developments in Language Theory*, DLT '01 Revised Papers, LNCS 2295, 2002,pages 21–36.
Abstract: We review the rudiments of the equational logic of (least) fixed points and provide some of its applications for axiomatization problems with respect to regular languages, tree languages, and synchronization trees.

[RS-01-21] *PostScript*, *PDF*, *DVI*.
Luca Aceto, Zoltn sik, and Anna Inglfsdttir.
Equational Theories of Tropical Semirings.
June 2001.
52 pp. To appear in *Theoretical Computer Science*. Extended abstracts of parts of this paper have appeared in Honsell and Miculan, editors, *Foundations of Software Science and Computation Structures*, FoSSaCS '01 Proceedings, LNCS 2030, 2001, pages 42–56 and in Gaubert and Loiseau, editors, *Workshop on Max-Plus Algebras and Their Applications to Discrete-Event Systems, Theoretical Computer Science, and Optimization*, MAX-PLUS '01 Proceedings, IFAC (International Federation of Automatic Control) IFAC Publications, 2001.
Abstract: This paper studies the equational theory of various exotic semirings presented in the literature. Exotic semirings are semirings whose underlying carrier set is some subset of the set of real numbers equipped with binary operations of minimum or maximum as sum, and addition as product. Two prime examples of such structures are the $(\max, +)$ *semiring* and the *tropical semiring*. It is shown that none of the exotic semirings commonly considered in the literature has a finite basis for its equations, and that similar results hold for the commutative idempotent weak semirings that underlie them. For each of these commutative idempotent weak semirings, the paper offers characterizations of the equations that hold in them, decidability results for their equational theories, explicit descriptions of the free algebras in the varieties they generate, and relative axiomatization results.

[RS-01-20] *PostScript*, *PDF*, *DVI*.
Catuscia Palamidessi and Frank D. Valencia.
A Temporal Concurrent Constraint Programming Calculus.
June 2001.
31 pp. Appears in Walsh, editor, *7th International Conference on Principles and Practice of Constraint Programming*, CP '01 Proceedings, LNCS 2239, 2001, pages 302–316.
Abstract: The tcc model is a formalism for reactive concurrent constraint programming. In this paper we propose a model of temporal concurrent constraint programming which adds to tcc the capability of modeling asynchronous and non-deterministic

timed behavior. We call this tcc extension the ntcc calculus. The expressiveness of ntcc is illustrated by modeling cells and asynchronous bounded broadcasting, by specifying temporal requirements such as response and invariance, and by modeling timed systems such as RCX controllers. We present a denotational semantics for modeling the strongest-postcondition behavior of ntcc processes, and, based on this semantics, we develop a proof system for proving linear temporal properties of these processes.

[RS-01-19] *PostScript*, *PDF*, *DVI.*
Jiř Srba.
On the Power of Labels in Transition Systems.
June 2001.
23 pp. Full and extended version of Larsen and Nielsen, editors, *Concurrency Theory: 12th International Conference*, CONCUR '01 Proceedings, LNCS 2154, 2001, pages 277–291.

Abstract: In this paper we discuss the role of labels in transition systems with regard to bisimilarity and model checking problems. We suggest a general reduction from labelled transition systems to unlabelled ones, preserving bisimilarity and satisfiability of mu-calculus formulas. We apply the reduction to the class of transition systems generated by Petri nets and pushdown automata, and obtain several decidability/complexity corollaries for unlabelled systems. Probably the most interesting result is undecidability of strong bisimilarity for unlabelled Petri nets.

[RS-01-18] *PostScript*, *PDF.*
Katalin M. Hangos, Zsolt Tuza, and Anders Yeo.
Some Complexity Problems on Single Input Double Output Controllers.
May 2001.
27 pp.

Abstract: We investigate the time complexity of constructing single input double output state feedback controller structures, given the directed structure graph G of a system. Such a controller structure defines a restricted type of P_3-partition of the graph G. A necessary condition $(*)$ has been found and two classes of graphs have been identified where the search problem of finding a feasible P_3-partition is polynomially solvable and, in addition, $(*)$ is not only necessary but also sufficient for the existence of a P_3-partition. It is shown further that the decision problem on two particular graph classes — defined in terms of forbidden subgraphs — remains *NP*-complete, but is polynomially solvable on the intersection of those two classes. Moreover, for every natural number m, a stabilizing structure with Single Input m-Output controllers can be found in polynomial time for the system in question, if it admits one.

[RS-01-17] *PostScript*, *PDF.*
Claus Brabrand, Anders Mller, Steffan Olesen, and Michael I. Schwartzbach.

Language-Based Caching of Dynamically Generated HTML.
May 2001.
18 pp. Appears in *World Wide Web Journal*, 5(4):305–323, 2002.

Abstract: Increasingly, HTML documents are dynamically generated by interactive Web services. To ensure that the client is presented with the newest versions of such documents it is customary to disable client caching causing a seemingly inevitable performance penalty. In the `<bigwig>` system, dynamic HTML documents are composed of higher-order templates that are plugged together to construct complete documents. We show how to exploit this feature to provide an automatic fine-grained caching of document templates, based on the service source code. A `<bigwig>` service transmits not the full HTML document but instead a compact JavaScript recipe for a client-side construction of the document based on a static collection of fragments that can be cached by the browser in the usual manner. We compare our approach with related techniques and demonstrate on a number of realistic benchmarks that the size of the transmitted data and the latency may be reduced significantly.

[RS-01-16] *PostScript*, *PDF*, *DVI*.
Olivier Danvy, Morten Rhiger, and Kristoffer H. Rose.
Normalization by Evaluation with Typed Abstract Syntax.
May 2001.
9 pp. Appears in *Journal of Functional Programming*, 11(6):673–68, November 2000.

Abstract: We present a simple way to implement typed abstract syntax for the lambda calculus in Haskell, using phantom types, and we specify normalization by evaluation (i.e., type-directed partial evaluation) to yield this typed abstract syntax. Proving that normalization by evaluation preserves types and yields normal forms then reduces to type-checking the specification.

[RS-01-15] *PostScript*, *PDF*, *DVI*.
Luigi Santocanale.
A Calculus of Circular Proofs and its Categorical Semantics.
May 2001.
30 pp. Appears in Nielsen and Engberg, editors, *Foundations of Software Science and Computation Structures*, FoSSaCS '02 Proceedings, LNCS 2303, 2002, pages 129–143.

Abstract: We present a calculus of proofs, the intended models of which are categories with finite products and coproducts, initial algebras and final coalgebras of functors that are recursively constructible out of these operations, that is, μ-bicomplete categories. The calculus satisfies the cut elimination and its main characteristic is that the underlying graph of a proof is allowed to contain a certain amount of cycles. To each proof of the calculus we associate a system of equations which has a meaning in every μ-bicomplete category. We prove that this system admits always a unique solution, and by means of this theorem we define the semantics of the calculus.

[RS-01-14] *PostScript*, *PDF*, *DVI*.
Ulrich Kohlenbach and Paulo B. Oliva.
Effective Bounds on Strong Unicity in L_1-Approximation.
May 2001.
38 pp. A revised and numerically much improved version appeared in *Annals of Pure and Applied Logic*, 121:1–38, 2003. Extended extended abstract appears in Hofmann, editor, *3rd International Workshop on Implicit Computational Complexity*, ICC '01 Proceedings, BRICS Notes Series NS-01-3, 2001, pages 117–122.

Abstract: In this paper we present another case study in the general project of Proof Mining which means the logical analysis of prima facie non-effective proofs with the aim of extracting new computationally relevant data. We use techniques based on monotone functional interpretation (developed by the first author) to analyze Cheney's simplification from 1965 of Jackson's original proof from 1921 of the uniqueness of the best L_1-approximation of continuous functions $f \in C[0,1]$ by polynomials $p \in P_n$ of degree $\leq n$. Cheney's proof is non-effective in the sense that it is based on classical logic and on the non-computational principle WKL (binary Knig Lemma). The result of our analysis provides the first effective (in all parameters f, n and ε) uniform modulus of uniqueness (a concept which generalizes 'strong uniqueness' studied extensively in approximation theory). Moreover, the extracted modulus has the optimal ε-dependency as follows from Kroo 78. The paper also describes how the uniform modulus of uniqueness can be used to compute the best L_1-approximations of a fixed $f \in C[0,1]$ with arbitrary precision, and includes some remarks on the case of best Chebycheff approximation.

[RS-01-13] *PostScript*, *PDF*, *DVI*.
Federico Crazzolara and Glynn Winskel.
Events in Security Protocols.
April 2001.
30 pp. Appears in *8th ACM conference on Computer and Communications Security*, CCS '01 Proceedings, 2001, pages 96–105.

Abstract: The events of a security protocol and their causal dependency can play an important role in the analysis of security properties. This insight underlies both strand spaces and the inductive method. But neither of these approaches builds up the events of a protocol in a compositional way, so that there is an informal spring from the protocol to its model. By broadening the models to certain kinds of Petri nets, a restricted form of contextual nets, a compositional event-based semantics is given to an economical, but expressive, language for describing security protocols; so the events and dependency of a wide range of protocols are determined once and for all. The net semantics is formally related to a transition semantics, strand spaces and inductive rules, as well as trace languages and event structures, so unifying a range of approaches, as well as providing conditions under which particular, more limited, models are adequate for the analysis of protocols. The net semantics allows the derivation of general properties and proof principles which are demonstrated in establishing an authentication property, following a diagrammatic style of proof.

[RS-01-12] *PostScript*, *PDF*, *DVI*.

Torben Amtoft, Charles Consel, Olivier Danvy, and Karoline Malmkjr.

The Abstraction and Instantiation of String-Matching Programs.

April 2001.

37 pp. Appears in Mogensen, Schmidt and Sudborough, editors, *The Essence of Computation: Complexity, Analysis*, Transformation, LNCS 2556, 2002, pages 332–357.

Abstract: We consider a naive, quadratic string matcher testing whether a pattern occurs in a text, we equip it with a cache mediating its access to the text, and we abstract the traversal policy of the pattern, the cache, and the text. We then specialize this abstracted program with respect to a pattern, using the off-the-shelf partial evaluator Similix.

Instantiating the abstracted program with a left-to-right traversal policy yields the linear-time behavior of Knuth, Morris and Pratt's string matcher. Instantiating it with a right-to-left policy yields the linear-time behavior of Boyer and Moore's string matcher.

[RS-01-11] *PostScript*, *PDF*.

Alexandre David and M. Oliver Mller.

From HUPPAAL *to* UPPAAL*: A Translation from Hierarchical Timed Automata to Flat Timed Automata*.

March 2001.

40 pp. Appears in Kutsche and Weber, editors, *Fundamental Approaches to Software Engineering: Fifth International Conference*, FASE '02 Proceedings, LNCS 2306, 2002, pages 218–232, with the title *Formal Verification of UML Statecharts with Real-Time Extensions*.

Abstract: We present a hierarchical version of timed automata, equipped with data types, hand-shake synchronization, and local variables. We describe the formal semantics of this *hierarchical timed automata* (HTA) formalism in terms of a transition system.

We report on the implementation of a flattening algorithm, that translates our formalism to a network of UPPAAL timed automata. We establish a correspondence between symbolic states of an HTA and its translations, and thus are able to make use of UPPAAL's simulator and model checking engine.

This technique is exemplified with a cardiac pacemaker model. Here, the overhead introduced by the translation is tolerable. We give run-time data for deadlock checking, timed reachability, and timed response analysis.

[RS-01-10] *PostScript*, *PDF*, *DVI*.
Daniel Fridlender and Mia Indrika.
Do we Need Dependent Types?
March 2001.
6 pp. Appears in *Journal of Functional Programming*, 10(4):409–415, 2000. Superseeds BRICS Report RS-98-38.

Abstract: Inspired by Danvy, we describe a technique for defining, within the Hindley-Milner type system, some functions which seem to require a language with dependent types. We illustrate this by giving a general definition of `zipWith` for which the Haskell library provides a family of functions, each member of the family having a different type and arity. Our technique consists in introducing ad hoc codings for natural numbers which resemble numerals in λ-calculus.

[RS-01-9] *PostScript*, *PDF*.
Claus Brabrand, Anders Mller, and Michael I. Schwartzbach.
Static Validation of Dynamically Generated HTML.
February 2001.
18 pp. Appears in Field and Snelting, editors, *ACM SIGPLAN-SIGSOFT Workshop on Program Analysis For Software Tools and Engineering*, PASTE '01 Proceedings, 2001, pages 38–45.

Abstract: We describe a static analysis of `<bigwig>` programs that efficiently decides if all dynamically computed XHTML documents presented to the client will validate according to the official DTD. We employ two interprocedural flow analyses to construct a graph summarizing the possible documents. This graph is subsequently analyzed to determine validity of those documents.

[RS-01-8] *PostScript*, *PDF*.
Ulrik Frendrup and Jesper Nyholm Jensen.
Checking for Open Bisimilarity in the π-Calculus.
February 2001.
61 pp.

Abstract: This paper deals with algorithmic checking of open bisimilarity in the π-calculus. Most bisimulation checking algorithms are based on the partition refinement approach. Unfortunately the definition of open bisimulation does not permit us to use a partition refinement approach for open bisimulation checking directly, but in the paper *A Partition Refinement Algorithm for the π-Calculus* Marco Pistore and Davide Sangiogi present an iterative method that makes it possible to check for open bisimilarity using partition refinement. We have implemented the algorithm presented by Marco Pistore and Davide Sangiorgi. Furthermore, we have optimized this algorithm and implemented this optimized algorithm. The time-complexity of this algorithm is the same as the time-complexity for the first algorithm, but performance tests have shown that in many cases the running time of the optimized algorithm is shorter than the running time of the first algorithm.

Our implementation of the optimized open bisimulation checker

algorithm and a user interface have been integrated in a system called the OBC Workbench. The source code and a manual for it is available from www.cs.auc.dk/research/FS/ny/PR-pi/.

[RS-01-7] *PostScript*, *PDF*, *DVI*.
Gregory Gutin, Khee Meng Koh, Eng Guan Tay, and Anders Yeo.
On the Number of Quasi-Kernels in Digraphs.
January 2001.
11 pp.

Abstract: A vertex set X of a digraph $D = (V, A)$ is a *kernel* if X is independent (i.e., all pairs of distinct vertices of X are non-adjacent) and for every $v \in V - X$ there exists $x \in X$ such that $vx \in A$. A vertex set X of a digraph $D = (V, A)$ is a *quasi-kernel* if X is independent and for every $v \in V - X$ there exist $w \in V - X, x \in X$ such that either $vx \in A$ or $vw, wx \in A$. In 1994, Chvtal and Lovsz proved that every digraph has a quasi-kernel. In 1996, Jacob and Meyniel proved that, if a digraph D has no kernel, then D contains at least three quasi-kernels. We characterize digraphs with exactly one and two quasi-kernels, and, thus, provide necessary and sufficient conditions for a digraph to have at least three quasi-kernels. In particular, we prove that every strong digraph of order at least three, which is not a 4-cycle, has at least three quasi-kernels. We conjecture that every digraph with no sink has a pair of disjoint quasi-kernels and provide some support to this conjecture.

[RS-01-6] *PostScript*, *PDF*, *DVI*.
Gregory Gutin, Anders Yeo, and Alexey Zverovich.
Traveling Salesman Should not be Greedy: Domination Analysis of Greedy-Type Heuristics for the TSP.
January 2001.
7 pp.

Abstract: Computational experiments show that the greedy algorithm (GR) and the nearest neighbor algorithm (NN), popular choices for tour construction heuristics, work at acceptable level for the Euclidean TSP, but produce very poor results for the general Symmetric and Asymmetric TSP (STSP and ATSP). We prove that for every $n \geq 2$ there is an instance of ATSP (STSP) on n vertices for which GR finds the worst tour. The same result holds for NN. We also analyze the repetitive NN (RNN) that starts NN from every vertex and chooses the best tour obtained. We prove that, for the ATSP, RNN always produces a tour, which is not worse than at least $n/2 - 1$ other tours, but for some instance it finds a tour, which is not worse than at most $n - 2$ other tours, $n \geq 4$. We also show that, for some instance of the STSP on $n \geq 4$ vertices, RNN produces a tour not worse than at most 2^{n-3} tours. These results are in sharp contrast to earlier results by G. Gutin and A. Yeo, and A. Punnen and S. Kabadi, who proved that, for the ATSP, there are tour construction heuristics, including some popular ones, that always build a tour not worse than at least $(n - 2)!$ tours.

[RS-01-5] *PostScript*, *PDF*.
Thomas S. Hune, Judi Romijn, Marille Stoelinga, and Frits W. Vaandrager.
Linear Parametric Model Checking of Timed Automata.
January 2001.
44 pp. Appears in Margaria and Yi, editors, *Tools and Algorithms for The Construction and Analysis of Systems: 7th International Conference*, TACAS '01 Proceedings, LNCS 2031, 2001, pages 174–188.
Abstract: We present an extension of the model checker UPPAAL capable of synthesize linear parameter constraints for the correctness of parametric timed automata. The symbolic representation of the (parametric) state-space is shown to be correct. A second contribution of this paper is the identification of a subclass of parametric timed automata (L/U automata), for which the emptiness problem is decidable, contrary to the full class where it is know to be undecidable. Also we present a number of lemmas enabling the verification effort to be reduced for L/U automata in some cases. We illustrate our approach by deriving linear parameter constraints for a number of well-known case studies from the literature (exhibiting a flaw in a published paper).

[RS-01-4] *PostScript*, *PDF*.
Gerd Behrmann, Ansgar Fehnker, Thomas S. Hune, Kim G. Larsen, Paul Pettersson, and Judi Romijn.
Efficient Guiding Towards Cost-Optimality in UPPAAL.
January 2001.
21 pp. Appears in Margaria and Yi, editors, *Tools and Algorithms for The Construction and Analysis of Systems: 7th International Conference*, TACAS '01 Proceedings, LNCS 2031, 2001, pages 174–188.
Abstract: In this paper we present an algorithm for efficiently computing the minimum cost of reaching a goal state in the model of Uniformly Priced Timed Automata (UPTA). This model can be seen as a submodel of the recently suggested model of linearly priced timed automata, which extends timed automata with prices on both locations and transitions. The presented algorithm is based on a symbolic semantics of UTPA, and an efficient representation and operations based on difference bound matrices. In analogy with Dijkstra's shortest path algorithm, we show that the search order of the algorithm can be chosen such that the number of symbolic states explored by the algorithm is optimal, in the sense that the number of explored states can not be reduced by any other search order. We also present a number of techniques inspired by branch-and-bound algorithms which can be used for limiting the search space and for quickly finding near-optimal solutions.

The algorithm has been implemented in the verification tool Uppaal. When applied on a number of experiments the presented techniques reduced the explored state-space with up to 90%.

[RS-01-3] *PostScript*, *PDF*.
Gerd Behrmann, Ansgar Fehnker, Thomas S. Hune, Kim G. Larsen, Paul Pettersson, Judi Romijn, and Frits W. Vaandrager.
Minimum-Cost Reachability for Priced Timed Automata.
January 2001.
22 pp. Appears in Di Benedetto and Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control, Fourth International Workshop*, HSCC '01 Proceedings, LNCS 2034, 2001, 147–161.
Abstract: This paper introduces the model of linearly priced timed automata as an extension of timed automata, with prices on both transitions and locations. For this model we consider the minimum-cost reachability problem: i.e. given a linearly priced timed automaton and a target state, determine the minimum cost of executions from the initial state to the target state. This problem generalizes the minimum-time reachability problem for ordinary timed automata. We prove decidability of this problem by offering an algorithmic solution, which is based on a combination of branch-and-bound techniques and a new notion of priced regions. The latter allows symbolic representation and manipulation of reachable states together with the cost of reaching them.

[RS-01-2] *PostScript*, *PDF*, *DVI*.
Rasmus Pagh and Jakob Pagter.
Optimal Time-Space Trade-Offs for Non-Comparison-Based Sorting.
January 2001.
ii+20 pp. Appears in Eppstein, editor, *The Thirteenths Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '02 Proceedings, 2002, pages 9–18.
Abstract: We study the fundamental problem of sorting n integers of w bits on a unit-cost RAM with word size w, and in particular consider the time-space trade-off (product of time and space in bits) for this problem. For comparison-based algorithms, the time-space complexity is known to be $\Theta(n^2)$. A result of Beame shows that the lower bound also holds for non-comparison-based algorithms, but no algorithm has met this for time below the comparison-based $\Omega(n \lg n)$ lower bound.

We show that if sorting within some time bound $\tilde{T}$ is possible, then time $T = O(\tilde{T} + n \lg^* n)$ can be achieved with high probability using space $S = O(n^2/T + w)$, which is optimal. Given a deterministic priority queue using amortized time $t(n)$ per operation and space $n^{O(1)}$, we provide a deterministic algorithm sorting in time $T = O(n\,(t(n) + \lg^* n))$ with $S = O(n^2/T + w)$. Both results require that $w \leq n^{1-\Omega(1)}$. Using existing priority queues and sorting algorithms, this implies that we can deterministically sort time-space optimally in time $\Theta(T)$ for $T \geq n\,(\lg\lg n)^2$, and with high probability for $T \geq n \lg\lg n$.

Our results imply that recent lower bounds for deciding element distinctness in $o(n \lg n)$ time are nearly tight.

[RS-01-1] *PostScript*, *PDF*, *DVI*.
Gerth Stlting Brodal, Rolf Fagerberg, Christian N. S. Pedersen, and

Anna stlin.

The Complexity of Constructing Evolutionary Trees Using Experiments.

January 2001.

28 pp. Appears in Orejas, Spirakis and Leeuwen, editors, *28th International Colloquium on Automata, Languages, and Programming*, ICALP '01 Proceedings, LNCS 2076, 2001, pages 140–151.

Abstract: We present a tight upper and lower bound for constructing evolutionary trees using experiments. We present an algorithm which constructs an evolutionary tree of n species using experiments in time $O(n \log_d n)$, where d is the degree of the constructed tree. We show by an explicit adversary argument that any algorithm for constructing an evolutionary tree of n species using experiments must perform $\Omega(n \log_d n)$ experiments, where d is the degree of the constructed tree.