# BRICS Research Series, Abstracts, 2005

March 6, 2006

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

## References

[RS-05-38] *PostScript*, *PDF*, *DVI*.
Małgorzata Biernacka and Olivier Danvy.
*A Syntactic Correspondence between Context-Sensitive Calculi and Abstract Machines.*
December 2005.
iii+39 pp. Revised version of BRICS RS-05-22.

**Abstract:** We present a systematic construction of environment-based abstract machines from context-sensitive calculi of explicit substitutions, and we illustrate it with ten calculi and machines for applicative order with an abort operation, normal order with generalized reduction and call/cc, the lambda-mu-calculus, delimited continuations, stack inspection, proper tail-recursion, and lazy evaluation. Most of the machines already exist but have been obtained independently and are only indirectly related to the corresponding calculi. All of the calculi are new and they make it possible to directly reason about the execution of the corresponding machines.

In connection with the functional correspondence between evaluation functions and abstract machines initiated by Reynolds, the present syntactic correspondence makes it possible to construct reduction-free normalization functions out of reduction-based ones, which was an open problem in the area of normalization by evaluation.

[RS-05-37] *PostScript*, *PDF*, *DVI*.
Gerth Stølting Brodal, Kanela Kaligosi, Irit Katriel, and Martin Kutz.
*Faster Algorithms for Computing Longest Common Increasing Subsequences.*
December 2005.
16 pp.

**Abstract:** We present algorithms for finding a longest common increasing subsequence of two or more input sequences. For two sequences of lengths $m$ and $n$, where $m \geq n$, we present an algorithm with an output-dependent expected running time of

$O((m + n\ell) \log\log \sigma + Sort)$ and $O(m)$ space, where $\ell$ is the length of a LCIS, $\sigma$ is the size of the alphabet, and *Sort* is the time to sort each input sequence.

For $k \geq 3$ length-$n$ sequences we present an algorithm running time $O(\min\{kr^2, r \log^{k-1} r\} + Sort)$, which improves the previous best bound by more than a factor $k$ for many inputs. In both cases, our algorithms are conceptually quite simple but rely on existing sophisticated data structures.

Finally, we introduce the problem of longest common weakly-increasing (or non-decreasing) subsequences (LCWIS), for which we present an $O(m + n \log n)$ time algorithm for the 3-letter alphabet case. For the extensively studied Longest Common Subsequence problem, comparable speedups have not been achieved for small alphabets.

[RS-05-36] *PostScript*, *PDF*, *DVI*.
Dariusz Biernacki, Olivier Danvy, and Chung-chieh Shan.
*On the Static and Dynamic Extents of Delimited Continuations.*
December 2005.
ii+33 pp. To appear in the journal *Science of Computer Programming*. Supersedes BRICS RS-05-13.

**Abstract:** We show that breadth-first traversal exploits the difference between the static delimited-control operator shift (alias S) and the dynamic delimited-control operator control (alias F). For the last 15 years, this difference has been repeatedly mentioned in the literature but it has only been illustrated with one-line toy examples. Breadth-first traversal fills this vacuum.

We also point out where static delimited continuations naturally give rise to the notion of control stack whereas dynamic delimited continuations can be made to account for a notion of 'control queue.'.

[RS-05-35] *PostScript*, *PDF*, *DVI*.
Kristian Støvring.
*Extending the Extensional Lambda Calculus with Surjective Pairing is Conservative.*
November 2005.
19 pp. To appear in *Logical Methods in Computer Science*.

**Abstract:** We answer Klop and de Vrijer's question whether adding surjective-pairing axioms to the extensional lambda calculus yields a conservative extension. The answer is positive. As a byproduct we obtain the first "syntactic" proof that the extensional lambda calculus with surjective pairing is consistent.

[RS-05-34] *PostScript*, *PDF*, *DVI*.
Henning Korsholm Rohde.
*Formal Aspects of Polyvariant Specialization.*
November 2005.
27 pp.

**Abstract:** We present the first formal correctness proof of an offline polyvariant specialization algorithm for first-order recursive equations. As a corollary, we show that the specialization algorithm generates a program implementing the search phase of the Knuth-Morris- Pratt algorithm from an inefficient but binding-time-improved string matcher.

[RS-05-33] *PostScript*, *PDF*, *DVI*.
Luca Aceto, Willem Jan Fokkink, Anna Ingólfsdóttir, and Sumit Nain.
*Bisimilarity is not Finitely Based over BPA with Interrupt.*
October 2005.
33 pp. This paper supersedes BRICS Report RS-04-24. An extended abstract of this paper appeared in *Algebra and Coalgebra in Computer Science, 1st Conference, CALCO 2005*, Swansea, Wales, 3–6 September 2005, Lecture Notes in Computer Science 3629, pp. 54–68, Springer-Verlag, 2005.
**Abstract:** This paper shows that bisimulation equivalence does not afford a finite equational axiomatization over the language obtained by enriching Bergstra and Klop's Basic Process Algebra with the interrupt operator. Moreover, it is shown that the collection of closed equations over this language is also not finitely based. In sharp contrast to these results, the collection of closed equations over the language BPA enriched with the disrupt operator is proven to be finitely based.

[RS-05-32] *PostScript*, *PDF*.
Anders Møller, Mads Østerby Olesen, and Michael I. Schwartzbach.
*Static Validation of XSL Transformations.*
October 2005.
50 pp.
**Abstract:** XSL Transformations (XSLT) is a programming language for defining transformations between XML languages. The structure of these languages is formally described by schemas, for example using DTD, which allows individual documents to be validated. However, existing XSLT tools offer no static guarantees that, under the assumption that the input is valid relative to the input schema, the output of the transformation is valid relative to the output schema.

We present a validation technique for XSLT based on the summary graph formalism introduced in the static analysis of JWIG Web services. Being able to provide static guarantees, we can detect a large class of errors in an XSLT stylesheet at the time it is written instead of later when it has been deployed, and thereby provide benefits similar to those of static type checkers for modern programming languages.

Our analysis takes a pragmatic approach that focuses its precision on the essential language features but still handles the entire XSLT 1.0 language. We evaluate the analysis precision on a range of real stylesheets and demonstrate how it may be useful in practice.

[RS-05-31] *PostScript*, *PDF*.

Christian Kirkegaard and Anders Møller.
*Type Checking with XML Schema in* XACT.
September 2005.
20 pp.
**Abstract:** We show how to extend the program analysis technique used in the XACT system to support XML Schema as type formalism. Moreover, we introduce optional type annotations to improve modularity of the type checking. The resulting system supports a flexible style of programming XML transformations and provides static guarantees of validity of the generated XML data.

[RS-05-30] *PostScript*, *PDF*.
Karl Krukow.
*An Operational Semantics for Trust Policies.*
September 2005.
38 pp.
**Abstract:** In the trust-structure model of trust management, principals specify their trusting relationships with other principals in terms of trust policies. In their paper on trust structures, Carbone et al. present a language for trust policies, and provide a suitable denotational semantics. The semantics ensures that for any collection of trust policies, there is always a unique global trust-state, compatible with all the policies, specifying everyone's degree of trust in everyone else. However, as the authors themselves point out, the language lacks an operational model: the global trust-state is a well-defined mathematical object, but it is not clear how principals can actually compute it. This becomes even more apparent when one considers the intended application environment: vast numbers of autonomous principals, distributed and possibly mobile. We provide a compositional operational semantics for a language of trust policies. The operational semantics is given in terms of a composition of I/O automata. We prove that this semantics is faithful to its corresponding denotational semantics, in the sense that any run of the I/O automaton "converges to" the denotational semantics of the policies. Furthermore, as I/O automata are a natural model of asynchronous distributed computation, the semantics leads to an algorithm for distributedly computing the trust-state, which is suitable in the application environment.

[RS-05-29] *PostScript*, *PDF*, *DVI*.
Olivier Danvy and Henning Korsholm Rohde.
*On Obtaining the Boyer-Moore String-Matching Algorithm by Partial Evaluation.*
September 2005.
ii+9 pp. To appear in *Information Processing Letters*. This version supersedes BRICS RS-05-14.
**Abstract:** We present the first derivation of the search phase of the Boyer-Moore string-matching algorithm by partial evaluation of an inefficient string matcher. The derivation hinges on identifying

the *bad-character-shift* heuristic as a binding-time improvement, bounded static variation. An inefficient string matcher incorporating this binding-time improvement specializes into the search phase of the Horspool algorithm, which is a simplified variant of the Boyer-Moore algorithm. Combining the *bad-character-shift* binding-time improvement with our previous results yields a new binding-time-improved string matcher that specializes into the search phase of the Boyer-Moore algorithm.

[RS-05-28] *PostScript*, *PDF*, *DVI*.
Jiří Srba.
*On Counting the Number of Consistent Genotype Assignments for Pedigrees.*
September 2005.
15 pp. To appear in FSTTCS '05.

**Abstract:** Consistency checking of genotype information in pedigrees plays an important role in genetic analysis and for complex pedigrees the computational complexity is critical. We present here a detailed complexity analysis for the problem of counting the number of complete consistent genotype assignments. Our main result is a polynomial time algorithm for counting the number of complete consistent assignments for non-looping pedigrees. We further classify pedigrees according to a number of natural parameters like the number of generations, the number of children per individual and the cardinality of the set of alleles. We show that even if we assume all these parameters as bounded by reasonably small constants, the counting problem becomes computationally hard (#P-complete) for looping pedigrees. The border line for counting problems computable in polynomial time (i.e. belonging to the class FP) and #P-hard problems is completed by showing that even for general pedigrees with unlimited number of generations and alleles but with at most one child per individual and for pedigrees with at most two generations and two children per individual the counting problem is in FP.

[RS-05-27] *PostScript*, *PDF*, *DVI*.
Pascal Zimmer.
*A Calculus for Context-Awareness.*
August 2005.
21 pp.

**Abstract:** In order to answer the challenge of pervasive computing, we propose a new process calculus, whose aim is to describe dynamic systems composed of agents able to move and react differently depending on their location. This *Context-Aware Calculus* features a hierarchical structure similar to mobile ambients, and a generic multi-agent synchronization mechanism, inspired from the join-calculus. After general ideas and introduction, we review the full calculus' syntax and semantics, as well as some motivating examples, study its expressiveness, and show how the notion of computation itself can be made context-dependent.

[RS-05-26] *PostScript*, *PDF*.
Henning Korsholm Rohde.
*Measuring the Propagation of Information in Partial Evaluation.*
August 2005.
39 pp.

**Abstract:** We present the first measurement-based analysis of the information propagated by a partial evaluator. Our analysis is based on measuring implementations of string-matching algorithms, based on the observation that the sequence of character comparisons accurately reflects maintained information. Notably, we can easily prove matchers to be different and we show that they display more variety and finesse than previously believed. As a consequence, we are able to pinpoint differences and inaccuracies in many results previously considered equivalent.

Our analysis includes a framework that lets us obtain string matchers – notably the family of Boyer-Moore algorithms – in a systematic formalism-independent way from a few information-propagation primitives. By leveraging the existing research in string matching, we show that the landscape of information propagation is non-trivial in the sense that small changes in information propagation may dramatically change the properties of the resulting string matchers. We thus expect that this work will prove useful as a test and feedback mechanism for information propagation in the development of advanced program transformations, such as GPC or Supercompilation.

[RS-05-25] *PostScript*, *PDF*, *DVI*.
Dariusz Biernacki and Olivier Danvy.
*A Simple Proof of a Folklore Theorem about Delimited Control.*
August 2005.
ii+11 pp. To appear in *Journal of Functional Programming*. This version supersedes BRICS RS-05-10.

**Abstract:** We formalize and prove the folklore theorem that the static delimited-control operators shift and reset can be simulated in terms of the dynamic delimited-control operators control and prompt. The proof is based on small-step operational semantics.

[RS-05-24] *PostScript*, *PDF*, *DVI*.
Małgorzata Biernacka, Dariusz Biernacki, and Olivier Danvy.
*An Operational Foundation for Delimited Continuations in the CPS Hierarchy.*
August 2005.
iv+43 pp. To appear in the journal *Logical Methods in Computer Science*. This version supersedes BRICS RS-05-11.

**Abstract:** We present an abstract machine and a reduction semantics for the $\lambda$-calculus extended with control operators that give access to delimited continuations in the CPS hierarchy. The abstract machine is derived from an evaluator in continuation-passing style (CPS); the reduction semantics (i.e., a small-step operational semantics with an explicit representation of evalua-

tion contexts) is constructed from the abstract machine; and the control operators are the shift and reset family. At level $n$ of the CPS hierarchy, programs can use the control operators $\text{shift}_i$ and $\text{reset}_i$ for $1 \leq i \leq n$, the evaluator has $n+1$ layers of continuations, the abstract machine has $n+1$ layers of control stacks, and the reduction semantics has $n+1$ layers of evaluation contexts.

We also present new applications of delimited continuations in the CPS hierarchy: finding list prefixes and normalization by evaluation for a hierarchical language of units and products.

[RS-05-23] *PostScript*, *PDF*, *DVI*.
Karl Krukow, Mogens Nielsen, and Vladimiro Sassone.
*A Framework for Concrete Reputation-Systems*.
July 2005.
48 pp. This is an extended version of a paper to be presented at ACM CCS'05.

**Abstract:** In a reputation-based trust-management system, agents maintain information about the past behaviour of other agents. This information is used to guide future trust-based decisions about interaction. However, while trust management is a component in security decision-making, few existing reputation-based trust-management systems aim to provide any formal security-guarantees. We provide a mathematical framework for a class of simple reputation-based systems. In these systems, decisions about interaction are taken based on policies that are exact requirements on agents' past histories. We present a basic declarative language, based on pure-past linear temporal logic, intended for writing simple policies. While the basic language is reasonably expressive, we extend it to encompass more practical policies, including several known from the literature. A naturally occurring problem becomes how to efficiently re-evaluate a policy when new behavioural information is available. Efficient algorithms for the basic language are presented and analyzed, and we outline algorithms for the extended languages as well.

[RS-05-22] *PostScript*, *PDF*, *DVI*.
Małgorzata Biernacka and Olivier Danvy.
*A Syntactic Correspondence between Context-Sensitive Calculi and Abstract Machines*.
July 2005.
iv+39 pp.

**Abstract:** We present a systematic construction of environment-based abstract machines from context-sensitive calculi of explicit substitutions, and we illustrate it with a series of calculi and machines: Krivine's machine with call/cc, the lambda-mu-calculus, delimited continuations, i/o, stack inspection, proper tail-recursion, and lazy evaluation. Most of the machines already exist but have been obtained independently and are only indirectly related to the corresponding calculi. All of the calculi are new and they make it possible to directly reason about the execution of the corresponding machines. In connection with the functional correspondence between evaluation functions and abstract machines initiated by Reynolds, the present syntactic correspondence makes it possible to construct reduction-free normalization functions out of reduction-based ones, which was an open problem in the area of normalization by evaluation.

[RS-05-21] *PostScript*, *PDF*, *DVI*.
Philipp Gerhardy and Ulrich Kohlenbach.
*General Logical Metatheorems for Functional Analysis*.
July 2005.
65 pp.

**Abstract:** In this paper we prove general logical metatheorems which state that for large classes of theorems and proofs in (nonlinear) functional analysis it is possible to extract from the proofs effective bounds which depend only on very sparse local bounds on certain parameters. This means that the bounds are uniform for all parameters meeting these weak local boundedness conditions. The results vastly generalize related theorems due to the second author where the global boundedness of the underlying metric space (resp. a convex subset of a normed space) was assumed. Our results treat general classes of spaces such as metric, hyperbolic, CAT(0), normed, uniformly convex and inner product spaces and classes of functions such as nonexpansive, Hölder-Lipschitz, uniformly continuous, bounded and weakly quasi-nonexpansive ones. We give several applications in the area of metric fixed point theory. In particular, we show that the uniformities observed in a number of recently found effective bounds (by proof theoretic analysis) can be seen as instances of our general logical results.

[RS-05-20] *PostScript*, *PDF*, *DVI*.
Ivan B. Damgård, Serge Fehr, Louis Salvail, and Christian Schaffner.
*Cryptography in the Bounded Quantum Storage Model*.
July 2005.
23 pp.

**Abstract:** We initiate the study of two-party cryptographic primitives with unconditional security, assuming that the adversary's *quantum* memory is of bounded size. We show that oblivious transfer and bit commitment can be implemented in this model using protocols where honest parties need no quantum memory, whereas an adversarial player needs quantum memory of size at least $n/2$ in order to break the protocol, where $n$ is the number of qubits transmitted. This is in sharp contrast to the classical bounded-memory model, where we can only tolerate adversaries with memory of size quadratic in honest players' memory size. Our protocols are efficient, non-interactive and can be implemented using today's technology. On the technical side, a new entropic uncertainty relation involving min-entropy is established.

[RS-05-19] *PostScript*, *PDF*.
Luca Aceto, Willem Jan Fokkink, Anna Ingólfsdóttir, and Bas Luttik.

*Finite Equational Bases in Process Algebra: Results and Open Questions.*
June 2005.
28 pp.

**Abstract:** Van Glabbeek (1990) presented the linear time-branching time spectrum of behavioral equivalences for finitely branching, concrete, sequential processes. He studied these semantics in the setting of the basic process algebra BCCSP, and tried to give finite complete and $\omega$-complete axiomatizations for them. (An axiomatization $E$ is $\omega$-complete when an equation can be derived from $E$ if, and only if, all its closed instantiations can be derived from $E$.) Obtaining such axiomatizations in concurrency theory often turns out to be difficult, even in the setting of simple languages like BCCSP. This has raised a host of open questions that have been the subject of intensive research in recent years. Most of these questions have been settled over BCCSP, either positively by giving a finite complete or $\omega$-complete axiomatization, or negatively by proving that such an axiomatization does not exist. Still some open questions remain. This paper reports on these results, and on the state-of-the-art on axiomatizations for richer process algebras, containing constructs like sequential and parallel composition.

[RS-05-18] *PostScript*, *PDF*, *DVI*.
Peter Bogetoft, Ivan B. Damgård, Thomas Jakobsen, Kurt Nielsen, Jakob Pagter, and Tomas Toft.
*Secure Computing, Economy, and Trust: A Generic Solution for Secure Auctions with Real-World Applications.*
June 2005.
37 pp.

**Abstract:** In this paper we consider the problem of constructing secure auctions based on techniques from modern cryptography. We combine knowledge from economics, cryptography and security engineering and develop and implement secure auctions for practical real-world problems.

In essence this paper is an overview of the research project SCET—Secure Computing, Economy, and Trust— which attempts to build auctions for real applications using secure multiparty computation.

The main contributions of this project are: A generic setup for secure evaluation of integer arithmetic including comparisons; general double auctions expressed by such operations; a real world double auction tailored to the complexity and performance of the basic primitives $+$ and $\leq$; and finally evidence that our approach is practically feasible based on experiments with prototypes.

[RS-05-17] *PostScript*, *PDF*, *DVI*.
Ivan B. Damgård, Thomas B. Pedersen, and Louis Salvail.
*A Quantum Cipher with Near Optimal Key-Recycling.*
May 2005.
29 pp.

**Abstract:** Assuming an insecure quantum channel and an authenticated classical channel, we propose an unconditionally secure scheme for encrypting classical messages under a shared key, where attempts to eavesdrop the ciphertext can be detected. If no eavesdropping is detected, we can securely re-use the entire key for encrypting new messages. If eavesdropping is detected, we must discard a number of key bits corresponding to the length of the message, but can re-use almost all of the rest. We show this is essentially optimal. Thus, provided the adversary does not interfere (too much) with the quantum channel, we can securely send an arbitrary number of message bits, independently of the length of the initial key. Moreover, the key-recycling mechanism only requires one-bit feedback. While ordinary quantum key distribution with a classical one time pad could be used instead to obtain a similar functionality, this would need more rounds of interaction and more communication.

[RS-05-16] *PostScript*, *PDF*, *DVI*.
Dariusz Biernacki, Olivier Danvy, and Kevin Millikin.
*A Dynamic Continuation-Passing Style for Dynamic Delimited Continuations.*
May 2005.
ii+24 pp.

**Abstract:** We present a new abstract machine that accounts for dynamic delimited continuations. We prove the correctness of this new abstract machine with respect to a pre-existing, definitional abstract machine. Unlike this definitional abstract machine, the new abstract machine is in defunctionalized form, which makes it possible to state the corresponding higher-order evaluator. This evaluator is in continuation+state passing style and threads a trail of delimited continuations and a meta-continuation. Since this style accounts for dynamic delimited continuations, we refer to it as 'dynamic continuation-passing style.'

We show that the new machine operates more efficiently than the definitional one and that the notion of computation induced by the corresponding evaluator takes the form of a monad. We also present new examples and a new simulation of dynamic delimited continuations in terms of static ones.

[RS-05-15] *PostScript*, *PDF*, *DVI*.
Małgorzata Biernacka and Olivier Danvy.
*A Concrete Framework for Environment Machines.*
May 2005.
ii+25 pp.

**Abstract:** We materialize the common belief that calculi with explicit substitutions provide an intermediate step between an abstract specification of substitution in the lambda-calculus and its concrete implementations. To this end, we go back to Curien's original calculus of closures (an early calculus with explicit substitutions), we extend it minimally so that it can express one-step reduction strategies, and we methodically derive a series

of environment machines from the specification of two one-step reduction strategies for the lambda-calculus: normal order and applicative order. The derivation extends Danvy and Nielsen's refocusing-based construction of abstract machines with two new steps: one for coalescing two successive transitions into one, and one for unfolding a closure into a term and an environment in the resulting abstract machine. The resulting environment machines include both the idealized and the original versions of Krivine's machine, Felleisen et al.'s CEK machine, and Leroy's Zinc abstract machine.

[RS-05-14] *PostScript*, *PDF*, *DVI*.
Olivier Danvy and Henning Korsholm Rohde.
*On Obtaining the Boyer-Moore String-Matching Algorithm by Partial Evaluation.*
April 2005.
ii+8 pp. Superseded by BRICS RS-05-29.
**Abstract:** We present the first derivation of the search phase of the Boyer-Moore string-matching algorithm by partial evaluation of an inefficient string matcher. The derivation hinges on identifying the 'bad-character-shift' heuristic as a binding-time improvement, bounded static variation. An inefficient string matcher incorporating this binding-time improvement specializes into the search phase of the Horspool algorithm, which is a simplified variant of the Boyer-Moore algorithm. Combining the bad-character-shift binding-time improvement with our previous results yields a new binding-time-improved string matcher that specializes into the search phase of the Boyer-Moore algorithm.

[RS-05-13] *PostScript*, *PDF*, *DVI*.
Dariusz Biernacki, Olivier Danvy, and Chung-chieh Shan.
*On the Dynamic Extent of Delimited Continuations.*
April 2005.
ii+32 pp. Extended version of an article to appear in *Information Processing Letters*. Subsumes BRICS RS-05-2.
**Abstract:** We show that breadth-first traversal exploits the difference between the static delimited-control operator 'shift' (alias 'S') and the dynamic delimited-control operator 'control' (alias 'F'). For the last 15 years, this difference has been repeatedly mentioned in the literature but it has only been illustrated with one-line toy examples. Breadth-first traversal fills this vacuum.

[RS-05-12] *PostScript*, *PDF*, *DVI*.
Małgorzata Biernacka, Olivier Danvy, and Kristian Støvring.
*Program Extraction from Proofs of Weak Head Normalization.*
April 2005.
19 pp. Extended version of an article to appear in the preliminary proceedings of MFPS XXI, Birmingham, UK, May 2005.
**Abstract:** We formalize two proofs of weak head normalization for the simply typed lambda-calculus in first-order minimal logic: one for normal-order reduction, and one for applicative-order reduction in the object language. Subsequently we use Kreisel's

modified realizability to extract evaluation algorithms from the proofs, following Berger; the proofs are based on Tait-style reducibility predicates, and hence the extracted algorithms are instances of (weak head) normalization by evaluation, as already identified by Coquand and Dybjer.

[RS-05-11] *PostScript*, *PDF*, *DVI*.
Małgorzata Biernacka, Dariusz Biernacki, and Olivier Danvy.
*An Operational Foundation for Delimited Continuations in the CPS Hierarchy.*
March 2005.
iii+42 pp. A preliminary version appeared in Thielecke, editor, *4th ACM SIGPLAN Workshop on Continuations*, CW '04 Proceedings, Association for Computing Machinery (ACM) SIGPLAN Technical Reports CSR-04-1, 2004, pages 25–33. This version supersedes BRICS RS-04-29, but is superseded by BRICS RS-05-24.

**Abstract:** We present an abstract machine and a reduction semantics for the $\lambda$-calculus extended with control operators that give access to delimited continuations in the CPS hierarchy. The abstract machine is derived from an evaluator in continuation-passing style (CPS); the reduction semantics (i.e., a small-step operational semantics with an explicit representation of evaluation contexts) is constructed from the abstract machine; and the control operators are the shift and reset family. At level $n$ of the CPS hierarchy, programs can use the control operators shift$_i$ and reset$_i$ for $1 \leq i \leq n$, the evaluator has $n+1$ layers of continuations, the abstract machine has $n+1$ layers of control stacks, and the reduction semantics has $n+1$ layers of evaluation contexts.

We also present new applications of delimited continuations in the CPS hierarchy: finding list prefixes and normalization by evaluation for a hierarchical language of units and products.

[RS-05-10] *PostScript*, *PDF*, *DVI*.
Dariusz Biernacki and Olivier Danvy.
*A Simple Proof of a Folklore Theorem about Delimited Control.*
March 2005.
ii+11 pp. Superseded by BRICS RS-05-25.

**Abstract:** We formalize and prove the folklore theorem that the static delimited-control operators shift and reset can be simulated in terms of the dynamic delimited-control operators control and prompt. The proof is based on a small-step operational semantics that takes the form of an abstract machine.

[RS-05-9] *PostScript*, *PDF*, *DVI*.
Gudmund Skovbjerg Frandsen and Peter Bro Miltersen.
*Reviewing Bounds on the Circuit Size of the Hardest Functions.*
March 2005.
6 pp. To appear in *Information Processing Letters*.

**Abstract:** In this paper we review the known bounds for $L(n)$, the circuit size complexity of the hardest Boolean function on $n$

input bits. The best known bounds appear to be

$$\frac{2^n}{n}(1+\frac{\log n}{n}-O(\frac{1}{n})) \leq L(n) \leq \frac{2^n}{n}(1+3\frac{\log n}{n}+O(\frac{1}{n}))$$

However, the bounds do not seem to be explicitly stated in the literature. We give a simple direct elementary proof of the lower bound valid for the full binary basis, and we give an explicit proof of the upper bound valid for the basis $\{\neg, \wedge, \vee\}$.

[RS-05-8] *PostScript*, *PDF*, *DVI*.
Peter D. Mosses.
*Exploiting Labels in Structural Operational Semantics.*
February 2005.
15 pp. Appears in *Fundamenta Informaticae*, 60:17–31, 2004.

**Abstract:** Structural Operational Semantics (SOS) allows transitions to be labelled. This is fully exploited in SOS descriptions of concurrent systems, but usually not at all in conventional descriptions of sequential programming languages.

This paper shows how the use of labels can provide significantly simpler and more modular descriptions of programming languages. However, the full power of labels is obtained only when the set of labels is made into a category, as in the recently-proposed MSOS variant of SOS.

[RS-05-7] *PostScript*, *PDF*, *DVI*.
Peter D. Mosses.
*Modular Structural Operational Semantics.*
February 2005.
46 pp. Appears in *Journal of Logic and Algebraic Programming*, 60–61:195–228, 2004.

**Abstract:** Modular SOS (MSOS) is a variant of conventional Structural Operational Semantics (SOS). Using MSOS, the transition rules for each construct of a programming language can be given incrementally, once and for all, and do not need reformulation when further constructs are added to the language. MSOS thus provides an exceptionally high degree of modularity in language descriptions, removing a shortcoming of the original SOS framework.

After sketching the background and reviewing the main features of SOS, the paper explains the crucial differences between SOS and MSOS, and illustrates how MSOS descriptions are written. It also discusses standard notions of semantic equivalence based on MSOS. An appendix shows how the illustrative MSOS rules given in the paper would be formulated in conventional SOS.

[RS-05-6] *PostScript*, *PDF*.
Karl Krukow and Andrew Twigg.
*Distributed Approximation of Fixed-Points in Trust Structures.*
February 2005.
41 pp.

**Abstract:** Recently, Carbone, Nielsen and Sassone introduced the trust-structure framework; a semantic model for trust-management in global-scale distributed systems. The framework is based on the notion of trust structures; a set of "trust-levels" ordered by two distinct partial orderings. In the model, a unique global trust-state is defined as the least fixed-point of a collection of local policies assigning trust-levels to the entities of the system. However, the framework is a purely denotational model: it gives precise meaning to the global trust-state of a system, but without specifying a way to compute this abstract mathematical object.

This paper complements q the denotational model of trust structures with operational techniques. It is shown how the least fixed-point can be computed using a simple, totally-asynchronous distributed algorithm. Two efficient protocols for approximating the least fixed-point are provided, enabling sound reasoning about the global trust-state without computing the exact fixed-point. Finally, dynamic algorithms are presented for safe reuse of information between computations, in face of dynamic trust-policy updates.

[RS-05-5] *PostScript*, *PDF*, *DVI*.
Dariusz Biernacki, Olivier Danvy, and Kevin Millikin.
*A Dynamic Continuation-Passing Style for Dynamic Delimited Continuations (Preliminary Version).*
February 2005.
ii+16 pp. Superseded by BRICS RS-05-16.

**Abstract:** We present a new abstract machine that accounts for dynamic delimited continuations. We prove the correctness of this new abstract machine with respect to a definitional abstract machine. Unlike this definitional abstract machine, the new abstract machine is in defunctionalized form, which makes it possible to state the corresponding higher-order evaluator. This evaluator is in continuation+state passing style, and threads a trail of delimited continuations and a meta-continuation. Since this style accounts for dynamic delimited continuations, we refer to it as 'dynamic continuation-passing style.'

We illustrate that the new machine is more efficient than the definitional one, and we show that the notion of computation induced by the corresponding evaluator takes the form of a monad.

[RS-05-4] *PostScript*, *PDF*, *DVI*.
Andrzej Filinski and Henning Korsholm Rohde.
*Denotational Aspects of Untyped Normalization by Evaluation.*
February 2005.
51 pp. Extended version of an article to appear in the FOSSACS 2004 special issue of RAIRO, *Theoretical Informatics and Applications.*

**Abstract:** We show that the standard normalization-by-evaluation construction for the simply-typed $\lambda_{\beta\eta}$-calculus has a natural counterpart for the untyped $\lambda_{\beta}$-calculus, with the central type-indexed logical relation replaced by a "recursively defined" *invariant relation*, in the style of Pitts. In fact, the construction can be seen as generalizing a computational-adequacy argument for an untyped, call-by-name language to normalization instead of evaluation.

In the untyped setting, not all terms have normal forms, so the normalization function is necessarily partial. We establish its correctness in the senses of *soundness* (the output term, if any, is in normal form and $\beta$-equivalent to the input term); *identification* ($\beta$-equivalent terms are mapped to the same result); and *completeness* (the function is defined for all terms that do have normal forms). We also show how the semantic construction enables a simple yet formal correctness proof for the normalization algorithm, expressed as a functional program in an ML-like call-by-value language.

Finally, we generalize the construction to produce an infinitary variant of normal forms, namely *Böhm trees*. We show that the three-part characterization of correctness, as well as the proofs, extend naturally to this generalization.

[RS-05-3] *PostScript*, *PDF*, *DVI*.
Olivier Danvy and Mayer Goldberg.
*There and Back Again*.
January 2005.
iii+16 pp. Extended version of an article to appear in *Fundamenta Informaticae*. This version supersedes BRICS RS-02-12.

**Abstract:** We present a programming pattern where a recursive function defined over a data structure traverses another data structure at return time. The idea is that the recursive calls get us 'there' by traversing the first data structure and the returns get us 'back again' while traversing the second data structure. We name this programming pattern of traversing a data structure at call time and another data structure at return time "There And Back Again" (TABA).

The TABA pattern directly applies to computing symbolic convolutions and to multiplying polynomials. It also blends well with other programming patterns such as dynamic programming and traversing a list at double speed. We illustrate TABA and dynamic programming with Catalan numbers. We illustrate TABA and traversing a list at double speed with palindromes and we obtain a novel solution to this traditional exercise. Finally, through a variety of tree traversals, we show how to apply TABA to other data structures than lists.

A TABA-based function written in direct style makes full use of an ALGOL-like control stack and needs no heap allocation. Conversely, in a TABA-based function written in continuation-passing style and recursively defined over a data structure (traversed at call time), the continuation acts as an iterator over a second data structure (traversed at return time). In general, the TABA pattern saves one from accumulating intermediate data structures at call time.

[RS-05-2] *PostScript*, *PDF*, *DVI*.
Dariusz Biernacki and Olivier Danvy.
*On the Dynamic Extent of Delimited Continuations*.
January 2005.
ii+30 pp.

**Abstract:** We show that breadth-first traversal exploits the difference between the static delimited-control operator shift (alias S) and the dynamic delimited-control operator control (alias F). For the last 15 years, this difference has been repeatedly mentioned in the literature but it has only been illustrated with one-line toy examples. Breadth-first traversal fills this vacuum.

We also point out where static delimited continuations naturally give rise to the notion of control stack whereas dynamic delimited continuations can be made to account for a notion of 'control queue.'.

[RS-05-1] *PostScript*, *PDF*, *DVI*.
Mayer Goldberg.
*On the Recursive Enumerability of Fixed-Point Combinators*.
January 2005.
7 pp. Superseedes BRICS report RS-04-25.

**Abstract:** We show that the set of fixed-point combinators forms a recursively-enumerable subset of a larger set of terms we call non-standard fixed-point combinators. These terms are observationally equivalent to fixed-point combinators in any computable context, but the set of on-standard fixed-point combinators is not recursively enumerable.