University of Pennsylvania
ScholarlyCommons

Technical Reports (CIS)

Department of Computer & Information Science

April 1994

Efficient Compilation of High-Level Data Parallel Algorithms

Dan Suciu
University of Pennsylvania

Val Tannen
University of Pennsylvania, val@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

Dan Suciu and Val Tannen, "Efficient Compilation of High-Level Data Parallel Algorithms", . April 1994.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-94-17.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/514
For more information, please contact repository@pobox.upenn.edu.

Efficient Compilation of High-Level Data Parallel Algorithms

Abstract

We present a high-level parallel calculus for nested sequences, *NSC* offered as a possible theoretical "core" of an entire class of collection-oriented parallel languages. *NSC* is based on while-loops as opposed to general recursion. A formal machine independent definition of the parallel time complexity and the work complexity of programs in *NSC* is given. Our main results are: (1) We give a translation method for a particular form of recursion, called *map-recursion* into *NSC*, that preserves the time complexity and adds an arbitrarily small overhead to the work complexity and (2) We give a compilation method for *NSC* into a very simple vector parallel machine which preserves the time complexity and again adds an arbitrarily small overhead to the work complexity.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-94-17.

Efficient Compilation of High-Level Data Parallel Algorithms

MS-CIS-94-17
LOGIC & COMPUTATION 78

Dan Suciu
Val Tannen

University of Pennsylvania
School of Engineering and Applied Science
Computer and Information Science Department
Philadelphia, PA 19104-6389

April 1994

Efficient Compilation of High-Level Data Parallel Algorithms

Dan Suci

Val Tannen*

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389, USA
Email: (suci, val)@saul.cis.upenn.edu

Abstract

We present a high-level parallel calculus for nested sequences, $\mathcal{NSC}$, offered as a possible theoretical “core” of an entire class of collection-oriented parallel languages. $\mathcal{NSC}$ is based on *while*-loops as opposed to general recursion. A formal, machine independent definition of the parallel time complexity and the work complexity of programs in $\mathcal{NSC}$ is given. Our main results are: (1) We give a translation method for a particular form of recursion, called *map*-recursion, into $\mathcal{NSC}$, that preserves the time complexity and adds an arbitrarily small overhead to the work complexity, and (2) We give a compilation method for $\mathcal{NSC}$ into a very simple vector parallel machine, which preserves the time complexity and again adds an arbitrarily small overhead to the work complexity.

1 Introduction

There are many advantages to programming in a high-level language. However, while sequential algorithms are most of the time designed and evaluated in reasonably high-level terms, the situation with parallel algorithms is - by necessity, so far - more complicated. The issue is intimately connected with the existing efforts to bridge the gap between the theoretical design of parallel algorithms and practical programming on massively parallel computers.

In the case of *data parallelism*, the work of Blelloch [Ble90, Ble93] and Blelloch and Sabot [BS90] has made substantial progress on this issue. For example, if we manage to represent an algorithm in a high-level language such as NESL with a certain work and time (a.k.a. element or step) complexity and if the representation satisfies certain restrictions then we are guaranteed an implementation of the same algorithm with the same asymptotic time and work complexity in terms of a low-level parallel vector model, which in turn admits efficient implementations on various architectures, for example the CM2. The present paper is proposing a different treatment of similar goals.

We start with a somewhat abstract high-level language which represents and manipulates mostly nested sequences (lists) and so we called it $\mathcal{NSC}$, for *nested sequence calculus* (section 3). We regard $\mathcal{NSC}$ as a possible theoretical “core” of an entire class of collection-oriented parallel languages. In keeping with the tenets of data parallelism [HS86], $\mathcal{NSC}$ ’s only parallel operation is *map* (apply-to-all). We give a precise high-level definition of parallel complexity (in the work and time framework [Jaj92]) for $\mathcal{NSC}$ programs.

Blelloch [Ble90, Ble93] gives convincing evidence that nested *map*’s on nested sequences (what he calls *nested parallelism*) can enhance the expressiveness of a data parallel language. But these high-level features are quite removed from concrete parallel architectures or even the parallel vector model and need to be compiled away. *Unnesting* the nested parallelism is at the center of the compilation technique of [Ble90, BS90, Ble93]. However, in a language with general recursion, this technique is guaranteed to preserve the asymptotic parallel complexity only for programs that satisfy a certain semantic condition called *containment*.

$\mathcal{NSC}$ is based on *while*-loops rather than general recursion. This will surely impose some limitations, although not that many: our *first main result* consists of showing that a large and practically relevant class of programs, called *map*-recursive, can be translated into $\mathcal{NSC}$ while asymptotically preserving the time complexity and adding an arbitrarily small overhead to the work complexity (theorem 4.2). It even turns out that some recursive programs which are not contained in the sense of [Ble90] are in fact *map*-recursive. The major benefit however is that we can compile $\mathcal{NSC}$ without the need for an unbounded stack of vectors, as general recursion would require. Avoiding the stack is a good idea because SIMD architectures associate a relatively small memory with each processor. A program that generates many entries in its vector stack will run out of memory even if the vectors are very short and hence much of the *total* amount of memory of the machine remains unused. We believe that our compilation technique can lead to better memory management. Of course, this needs to be tested in practice.

Following Blelloch, we define a simple parallel vector model in order to describe abstractly the class of target architectures for our compilation method (section 2). Our BVRAM (Bounded Vector Random Access Machine) differs from the VRAM [Ble90] primarily in that it has a finite number of vector registers. This emphasizes the absence of a run-time vector stack. Of course the number of registers needed

*The authors were partially supported by NSF Grant CCR-90-57570

depends on the source program being compiled. Another important difference is that we need less powerful communication primitives. The BVRAM has no general permutation instruction, and its communication primitives can be implemented on a butterfly network with $n \log n$ nodes in $O(\log n)$ steps. The BVRAM can be efficiently implemented on SIMD architectures such as CM2 and MasPar MP-1, and it has the potential of efficient implementation on MIMD machines as well, such as CM5, Paragon XP/S, KSR1 etc.

Our *second main result* is a technique that compiles any $\mathcal{N}SC$ program into a BVRAM program again while asymptotically preserving the time complexity and adding an arbitrarily small overhead to the work complexity (theorem 7.1). Along the way we also give a simulation that allows us to understand $\mathcal{N}SC$ complexity in terms of the complexity of computations on a certain flavor of PRAM (proposition 3.2), we show how to implement the BVRAM instructions on a butterfly network (proposition 2.1), we connect $\mathcal{N}SC$ with some standard parallel complexity classes (proposition 6.2), we show how to represent in $\mathcal{N}SC$ Valiant's $O(\log n \log \log n)$ time sorting algorithm [Val75, Jaj92] (section 5), and, as part of the compilation process, we define an intermediate abstract language - the *sequence algebra* - which has the same power as BVRAM's but may prove more flexible in connecting to the designs of the future (section 7).

$\mathcal{N}SC$ borrows heavily from our experience with languages for collection types [BTS91, BBW92] and it is worthwhile mentioning that many of its operations make as much sense for sets and bags (multisets) as for lists (sequences). It matters to us, though it may not be so relevant to the goals of this paper, that $\mathcal{N}SC$ is based on a clear, statically checkable type system, that we understand the meaning of $\mathcal{N}SC$ programs independently of their parallel execution, and that we know how to reason about them - for example how to validate source to source optimizations. We have in mind applications to databases and this naturally brings up important complexity issues. In a previous paper we have shown a tight connection between a related data parallel language for sets and the class NC [SBT94]. This in turn has led us to the more practical questions addressed here.

2 The Target: Bounded Vector Random Access Machines

To compile the higher level programming language described in section 3 only a very simple vector parallel model is needed. The Bounded Vector Random Access Machine, BVRAM, is a restriction of the VRAM introduced in [Ble90], in that it only admits a fixed number of registers, and has only particular communication primitives, not a general permutation. The BVRAM can be efficiently implemented on a wide range of parallel architectures, because: (1) only a simple, rather particular form of communication is needed to implement every instruction of the BVRAM, and (2) memory management at each processor is simplified by having only a bounded number of vector registers, as opposed to an unbounded number in the VRAM model.

A BVRAM, M , consists of a fixed number of **vector registers** $V_1, \dots, V_r$. Each V_i can hold a sequence (a vector) of natural numbers of arbitrary, but finite length. To keep the model simple, we don't include *scalar* registers: a number is represented by a sequence of length 1. A *program* for M is a sequence of labeled instructions, from the following instruction set. For some of the instructions below, it is convenient

to view a pair of registers V_i, V_j in which the length of the first equals the sum of the numbers in the second as a *nested sequence*. E.g., intuitively we view $[x_0, x_1, z_0, z_1, z_2], [2, 0, 3]$ as standing for the nested sequence $[[x_0, x_1], [], [z_0, z_1, z_2]]$.

- Move instruction: $V_i \leftarrow V_j$.
- Arithmetic operations, of the form $V_i \leftarrow V_j \text{ op } V_k$. Here op is an arithmetic operation from a set Σ . V_j and V_k must be arrays of the same length, and the operation op is applied simultaneously on all elements of V_j and V_k from the same positions, and the result is stored in V_i . In general we leave Σ unspecified, but mention here that for theorems 4.2 and 7.1 Σ has to contain $+, -, *, /$, *right-shift*, $\log_2$, while for proposition 6.2 we require that all operations in Σ be in NC. *Minus*, written $m - n$, is defined as $m - n$ when $m \geq n$ and 0 otherwise.
- Sequence oriented operations: $V_i \leftarrow []$ loads the empty sequence in V_i . $V_i \leftarrow [n]$, where $n \in \mathbb{N}$ loads the singleton sequence $[n]$ into V_i . $V_i \leftarrow V_j @ V_k$ appends V_j and V_k and stores the result in V_i . $V_i \leftarrow [\text{length}(V_j)]$ computes the length of V_j . $V_i \leftarrow \text{enumerate}(V_j)$ loads the sequence $[0, 1, \dots, n-1]$ into V_i , where n is the length of V_j .
- Bounded monotone routing $V_i \leftarrow \text{bm_route}(V_j, V_k, V_l)$; here V_k and V_l must have the same length. The effect is that each element in V_l is replicated a number of times equal to the corresponding number in V_k . In addition, it is required that the result matches in length the sequence V_j (i.e. initially V_j, V_k represent a nested sequence). E.g. if $V_j = [x_0, x_1, z_0, z_1, z_2]$, $V_k = [2, 0, 3]$ and $V_l = [a, b, c]$, then the instruction $V_i \leftarrow \text{bm_route}(V_j, V_k, V_l)$ stores $[a, a, c, c, c]$ into V_i .
- Segmented bounded monotone routing $V_i \leftarrow \text{sbm_route}(V_j, V_k, V_l, V_m)$. Here, V_j, V_k and V_l, V_m must be nested sequences, and $\text{length}(V_k) = \text{length}(V_m)$. Then, the subsequences of V_l are replicated according to the numbers in V_k and the result is stored in V_i . E.g., suppose $V_j = [x_0, x_1, z_0, z_1, z_2]$, $V_k = [2, 0, 3]$, $V_l = [a_0, a_1, b_0, b_1, b_2, c_0, c_1, c_2]$ and $V_m = [2, 3, 3]$. Then, after $V_i \leftarrow \text{sbm_route}(V_j, V_k, V_l, V_m)$, V_i will hold the value $[a_0, a_1, a_0, a_1, c_0, c_1, c_2, c_0, c_1, c_2, c_0, c_1, c_2]$. In the particular case when V_k, V_m have length one, this computes the *cartesian product* of V_j and V_l . Note that the length of the output is $\leq \text{length}(V_j) * \text{length}(V_l)$ and that bm_route can be expressed with two sbm_route instructions.
- Selection $V_i \leftarrow \sigma(V_j)$. The effect is that the nonzero values of V_j are packed and moved into V_i . E.g. if $V_j = [3, 0, 1, 0, 0, 4]$, then $[3, 1, 4]$ is stored in V_i .
- The unconditional jump `goto l` and the conditional jump `if empty?( $V_i$ ) then goto l`, where l is a label of some instruction. The conditional jump is taken iff V_i currently holds the empty sequence.
- *halt*, stops the program.

We associate with each BVRAM program P two numbers: r_i, r_o , the number of *input* and *output* registers. P expects r_i inputs in the registers $V_1, \dots, V_{r_i}$, and returns r_o outputs, in $V_1, \dots, V_{r_o}$. For some input, the result of P

might be undefined, if P enters an infinite loop, or if an error occurs. For a terminating execution of P , we define the **parallel time complexity** T to be the total number of instructions executed by P , i.e. each instruction is considered to have parallel time complexity 1. Similarly, we define the **work complexity** W as the sum of the work complexities of all instructions executed by P , where the work complexity of some instruction is defined to be the sum of the lengths of its input and output registers.

As opposed to VRAMs [Ble90] there is no general permutation instruction on a BVRAM (but one can be computed with an increase in the time or work complexity). This may lead to more efficient implementations on fixed-connection networks, as exemplified by the following proposition.

Proposition 2.1 *Any BVRAM instruction of work complexity W can be implemented in time $O(\log n)$ on a butterfly network with $n \log n$ nodes, where $n = O(W)$, using only oblivious routing algorithms.*

Proof. (Sketch) The arithmetic operations involve no communication at all, thus can be implemented in $O(1)$ steps. The append operation $V_i \leftarrow V_j @ V_k$ only requires a *monotone routing* of the values in V_k . This can be done in $O(\log n)$ steps, using the greedy routing algorithm, see [Lei92], pp. 534. *bm_route* is implemented by a monotone routing, and takes $O(\log n)$ steps with the greedy algorithm. For *sbm_route*, suppose first that $\text{length}(V_j) = \text{length}(V_i) = 1$, i.e. *sbm_route* computes the cartesian product of V_i and V_k . Also, suppose that the length of V_i and V_k are powers of 2, namely 2^p and 2^q respectively. Take $n = 2^{p+q}$; then we have 2^p packets residing in the first 2^p rows of a butterfly with 2^{p+q} rows, and we have to route the packet with address $00 \dots 0u_{p-1} \dots u_1u_0$ to all addresses of the form $v_{q-1} \dots v_1v_0u_{p-1} \dots u_1u_0$. This is done in q stages, starting with the higher dimension, using the greedy algorithm. In the general case of *sbm_route*, we have to replicate a number of smaller sequences. First, round upwards to the closest power of 2 the length of each such subsequence, and spread the sequences such that each sequence of length m starts at an address divisible by m . Next, perform in parallel all replications, as described above. $\square$

When the number n of available processors is less than the number W of elements in an array, then we group $\frac{W}{n}$ adjacent elements of the array in the same processor. The above proposition can be extended to this case: some instruction of complexity W can be implemented on a butterfly network in $O(\frac{W}{n} \log n)$ steps.

3 The Source: The Nested Sequence Calculus ($\mathcal{NSC}$)

We use *types* to explain the structure of $\mathcal{NSC}$ and classify its features. The types are given by the grammar $t ::= \text{unit} \mid \mathbb{N} \mid t \times t \mid t + t \mid [t]$. *unit* has exactly one value: the empty tuple $()$. $\mathbb{N}$ is the type of nonnegative integers. The values of the **product type** $s \times t$ are pairs (x, y) , with $x \in s, y \in t$. $[t]$ is the **finite sequences type** over t : it contains all sequences $[x_0, \dots, x_{n-1}]$, with $n \geq 0$ and $x_0, \dots, x_{n-1} \in t$. $s + t$ is the **disjoint union type** of s and t ; its values are of the form $\text{in}_1(x)$ with $x \in s$ and $\text{in}_2(y)$ with $y \in t$. We define the **boolean type** $\mathbb{B} \stackrel{\text{def}}{=} \text{unit} + \text{unit}$, and identify its values $\text{in}_1(())$ and $\text{in}_2(())$ with *true* and *false* respectively. Extending the list of built-in types with reals, strings, etc., can be done while preserving all results.

The primitives of $\mathcal{NSC}$ are chosen to be operations naturally associated to its types. Its expressions belong to one of two distinct syntactic categories: **terms**, denoted by M, N, P, U, V , etc., which have some type t , and **functions**, denoted by F, G , etc., have associated two types, the domain s and codomain t . By abuse of the language we say in this case that the “type” of some function F is $s \rightarrow t$. However $s \rightarrow t$ is not a type per se, which makes constructs like $s \rightarrow (t_1 \rightarrow t_2)$ or $(s_1 \rightarrow s_2) \rightarrow t$ impossible. Both terms and functions may contain free variables. See appendix A for a full and formal description of the language):

- Variables x , error Ω , constants n (where $n \in \mathbb{N}$), arithmetic operations $M \text{ op } N$, where $\text{op} \in \Sigma$ (recall from section 2 that $\Sigma = \{+, -, *, /, \dots\}$), and equality $M = N$.
- Constructs associated with the product type: $()$, π_1, π_2 , (M, N) . Here $()$ denotes the empty tuple, (M, N) is a pair, while π_1, π_2 are the projections, with the meaning $\pi_1(x, y) \stackrel{\text{def}}{=} x, \pi_2(x, y) \stackrel{\text{def}}{=} y$.
- Constructs associated with the sum type: $\text{in}_1(M)$, $\text{in}_2(N)$, and *case* M of $\text{in}_1(x) \Rightarrow N \text{ in}_2(x) \Rightarrow P$. The latter is defined to be equal to $N[U/x]$ when $M = \text{in}_1(U)$, and respectively to $P[V/y]$, when $M = \text{in}_2(V)$.
- Constructs associated with functions: $\lambda x : s. M$ and $F(M)$. The former is called a *lambda abstraction*, and is a *function* (as opposed to a *term*), of type $s \rightarrow t$, provided that M is a term of type t . The second construct, $F(M)$, is a term called *function application* having type t , provided that F is some function of type $s \rightarrow t$, and M is a term of type s . Although the type s is part of the syntax of $\lambda x : s. M$, we shall drop it when it is clear from the context. Note that $\lambda x : s. F$, where F is a function, is not a legal construct in $\mathcal{NSC}$, nor is $\lambda x : s \rightarrow t. M$, i.e. no higher order functions are allowed.
- Iteration: *while* (P, F) is some function of type $t \rightarrow t$, provided that P and F are functions of type $t \rightarrow \mathbb{B}$ and $t \rightarrow t$ respectively.
- Constructs associated with collections (these constructs work on sequences but also make sense for other kinds of collections, like sets and bags [BBW92]): $[]$, $[M]$, $M @ N$, *flatten* (M) , *length* (M) , *get* (M) , and *map* (F) . Here $[]$ denotes the empty sequence, $[M]$ is the singleton sequence, and $@$ is the append operator. Next, *flatten* $([x_0, \dots, x_{n-1}]) \stackrel{\text{def}}{=} x_0 @ x_1 @ \dots @ x_{n-1}$, and *length* (M) returns the length of some sequence. *get* is defined by *get* $([x]) \stackrel{\text{def}}{=} x$, *get* $([]) \stackrel{\text{def}}{=} \text{get}([x_0, x_1, \dots]) \stackrel{\text{def}}{=} \Omega$. Finally *map* (F) is a function of type $[s] \rightarrow [t]$, provided that F is a function of type $s \rightarrow t$. Its meaning is: *map* $(F)([x_0, \dots, x_{n-1}]) \stackrel{\text{def}}{=} [F(x_0), \dots, F(x_{n-1})]$.
- Constructs associated only to sequences, and not to other kinds of collections: *zip* (M, N) , *enumerate* (M) , and *split* (M, N) . The meanings are: *zip* $([x_0, \dots, x_{n-1}], [y_0, \dots, y_{n-1}]) \stackrel{\text{def}}{=} [(x_0, y_0), \dots, (x_{n-1}, y_{n-1})]$ (*zip* is undefined if its two arguments have different lengths), *enumerate* $([x_0, \dots, x_{n-1}]) \stackrel{\text{def}}{=} [0, \dots, n-1]$. Finally

$\text{split}(M, N)$ splits M according to the numbers contained in N ; e.g. $\text{split}([a, b, c, d, e, f], [3, 0, 1, 0, 2]) \stackrel{\text{def}}{=} [[a, b, c], [], [d], [], [e, f]]$. It is defined only if the sum of elements in N equals the length of M .

Note that any function in $\mathcal{NSC}$ can, in a fixed amount of time, only increase the size of its input by some polynomial. Had we introduced as a primitive in the language something like $\iota(n) \stackrel{\text{def}}{=} [0, \dots, n-1]$, which generates an arbitrarily long list out of a number, this property would fail. From this small set of primitives, we can *derive* a rich set of functions. Some examples:

Database projections. $\Pi_i : [t_1 \times t_2] \rightarrow [t_i]$, $\Pi_i \stackrel{\text{def}}{=} \text{map}(\pi_i)$.

Conditionals. if $x = y$ then M else N is expressed by $\text{case } (x = y) \text{ of } in_1(u) \Rightarrow M \text{ } in_2(v) \Rightarrow N$, where u, v are variables of type *unit*, not occurring in M, N .

Broadcasting. $\rho_2(x, [y_0, \dots, y_{n-1}])$, which is defined to be $[(x, y_0), \dots, (x, y_{n-1})]$, can be expressed as $\rho_2(x, y) \stackrel{\text{def}}{=} \text{map}(\lambda(v). (x, v))(y)$ and has the type $\rho_2 : s \times [t] \rightarrow [s \times t]$. When x itself is a sequence, $\rho_2(x, y)$ essentially computes the cartesian product of x and y . (The name ρ_2 is motivated by other considerations [BBW92].)

Bounded monotone routing. $\text{bm_route}((u, d), x)$, expressed as $\Pi_1(\text{flatten}(\text{map}(\rho_2)(\text{zip}(x, \text{split}(u, d)))))$, has type $\text{bm_route} : ([s] \times [\mathbb{N}]) \times [t] \rightarrow [t]$. bm_route is essentially the same operation as described in section 2. E.g. $\text{bm_route}([u_0, u_1, u_2, v_0, v_1], [3, 0, 2], [a, b, c]) = [a, a, a, c, c]$. The bound u prohibits us from constructing a very long sequence in constant parallel time. An *unbounded* monotone routing $m_route : [\mathbb{N}] \times [t] \rightarrow [t]$ can be defined in $\mathcal{NSC}$ (with *while*), but requires more than a constant number of parallel steps. This is indeed necessary, since $m_route([n], [a])$, produces the sequence $[a, a, \dots, a]$ of length n , whose size is not polynomially bounded by the input. Finally, note that in the context of nested sequences, our bounded monotone routing is not truly “monotone”. Indeed, $\text{bm_route}([([()], [2]), [a, b, c]]) = [[a, b, c], [a, b, c]]$, and the relative order of a, b and c has not been preserved. This forces us to introduce sbm_route in the BVRAM model.

Selections. $\sigma_1 : [s + t] \rightarrow [s]$, $\sigma_2 : [s + t] \rightarrow [t]$. $\sigma_1(x)$ selects from some sequence x only those elements which have the form $in_1(u)$, while $\sigma_2(x)$ selects only the elements of the form $in_2(v)$. E.g. if $x = [in_1(a), in_2(b), in_2(c), in_2(d), in_1(e), in_2(f)]$, then $\sigma_1(x) = [a, e]$, $\sigma_2(x) = [b, c, d, f]$. σ_1 is defined by $\sigma_1(x) \stackrel{\text{def}}{=} \text{flatten}(\lambda(u). \text{case } u \text{ of } in_1(u') \Rightarrow [u'] \text{ } in_2(u'') \Rightarrow []) (x)$, and σ_2 is defined similarly.

Operations on lists. *first* and *tail* can be defined by:

$$\begin{aligned} \text{first}(x) &\stackrel{\text{def}}{=} \text{get}(\text{get}(\text{bm_route}([([], [1, 0]), \\ &\quad \text{split}(x, [1, \text{length}(x) - 1])))) \\ \text{tail}(x) &\stackrel{\text{def}}{=} \text{get}(\text{bm_route}([([], [0, 1]), \\ &\quad \text{split}(x, [1, \text{length}(x) - 1])))) \end{aligned}$$

If x is empty, *split* will produce an error. Similarly we can define *last* and *remove_last*, which return the last element, and delete the last element from a sequence, respectively. In general, we can access any element of some sequence of length n in $O(1)$ parallel time, and with $O(n)$ work complexity (we formally define below the time and work complexity). Using *map*, we can produce an arbitrary permutation in $O(1)$ parallel time, but with an increase of the work complexity to

$O(n^2)$. Using radix sort in base n^ϵ , for some arbitrary $\epsilon > 0$, we can even compute an arbitrary permutation in $O(1)$ parallel time with $O(n^{1+\epsilon})$ work complexity. Alternatively, we can use an optimal sorting algorithm (see e.g. [Jaj92]), which reduces the work complexity to $O(n)$ by increasing the time complexity (e.g. the sorting algorithm described in section 5 has $T = O(\log n \log \log n)$). Thus, the cost of performing an arbitrary permutation is visible in the higher level language.

The compilation theorem 7.1 is robust enough to hold if $\mathcal{NSC}$ is extended with additional primitives, like a general permutation *permute* or scan operations, provided that corresponding instructions are added to the BVRAM model. E.g. theorem 7.1 can be extended to prove that $\mathcal{NSC} + \text{permute}$ can be efficiently compiled into BVRAM+*permute*. But in its present form theorem 7.1 is stronger, because it proves a general permutation is not necessary in a BVRAM in order to compile efficiently a high-level language like $\mathcal{NSC}$. This is of importance in view of the high cost of implementing a general permutation on existing massively parallel architectures [KLGLS90].

As promised, we will give a high-level definition of parallel time complexity T and work complexity W for $\mathcal{NSC}$ programs, in a machine independent way. The idea is for the parallel complexity of a program to be inferred from its structure in the same way in which the sequential complexity is inferred from the structure of a program in a sequential language. In our case, all primitive operations (including @ and *flatten*) take one parallel step, while in a $\text{map}(F)([x_0, \dots, x_{n-1}])$, the n executions of F are done in parallel. The iteration construct however may count for several steps hence our definition cannot be done solely by induction on programs. This is handled by providing a formal operational semantics and then counting the depth of derivations in it. The work complexity is tied to the size of the data that is being manipulated.

Formally, we start by defining **S-objects** by the grammar: $C ::= () \mid n \mid (C, C) \mid in_1(C) \mid in_2(C) \mid [C, \dots, C]$ where $n \in \mathbb{N}$. We only consider typed S-objects. We adopt a *unit size complexity measure*, and define the *size* of some S-object by $\text{size}() = \text{size}(n) = 1$, $\text{size}((C, D)) = 1 + \text{size}(C) + \text{size}(D)$, $\text{size}(in_1(C)) = \text{size}(in_2(C)) = 1 + \text{size}(C)$, $\text{size}([C_0, \dots, C_{n-1}]) = 1 + \sum_{i=0, n-1} \text{size}(C_i)$. We use *true* and *false* as abbreviations for $in_1()$ and $in_2()$.

Next, we define the **evaluation** of some term (also called the operational semantics) in a natural semantics style, as in [Kah87]. This consists of rules which simultaneously define a binary relation $M \Downarrow C$ meaning that the term M evaluates to the S-object C and a ternary relation $F(C) \Downarrow C'$ meaning that the function F applied to the S-object C evaluates to C' . E.g. if $F = \lambda x. \text{flatten}(x) \text{@} [100]$ and $C = [[3, 5], [2]]$, then $F(C) \Downarrow [3, 5, 2, 100]$. Some representative rules are:

$$\begin{aligned} &\frac{M \Downarrow m \quad N \Downarrow n}{M + N \Downarrow m + n} \\ &\frac{M \Downarrow (C, D)}{\pi_1(M) \Downarrow C} \quad \frac{M \Downarrow (C, D)}{\pi_2(M) \Downarrow D} \\ &\frac{M \Downarrow C \quad N \Downarrow D}{(M, N) \Downarrow (C, D)} \quad \frac{M \Downarrow C \quad F(C) \Downarrow D}{F(M) \Downarrow D} \\ &\frac{M \Downarrow [C_0, \dots, C_{m-1}] \quad N \Downarrow [D_0, \dots, D_{n-1}]}{M @ N \Downarrow [C_0, \dots, C_{m-1}, D_0, \dots, D_{n-1}]} \\ &\frac{F(C_0) \Downarrow D_0 \quad \dots \quad F(C_{n-1}) \Downarrow D_{n-1}}{\text{map}(F)([C_0, \dots, C_{n-1}]) \Downarrow [D_0, \dots, D_{n-1}]} \\ &\frac{P(C) \Downarrow \text{false}}{\text{while}(P, F)(C) \Downarrow C} \end{aligned}$$

$$\frac{P(C) \Downarrow \text{true} \quad F(C) \Downarrow C' \quad \text{while}(P, F)(C') \Downarrow D}{\text{while}(P, F)(C) \Downarrow D}$$

The complete set of rules is given in appendix B where we explain a technical complication caused by the presence of bound variables (lambda abstraction) in the language, namely the need to use *environments* as in [Cur88].

Thus, to *evaluate* some closed term M , one has to construct a *proof tree*, whose nodes are labeled with rules of the operational semantics, such that its root is labeled with some rule with conclusion $M \Downarrow C$. Based on this operational semantics, we now define the time and work complexity of $\mathcal{NSC}$ in a machine independent way.

Definition 3.1 Consider some $\mathcal{NSC}$ term M . The time and work complexity $T(M), W(M)$ of $M \Downarrow C$ are defined by induction on the proof of $M \Downarrow C$. The induction is done simultaneously with the definition of the time and work complexity $T(F, C)$ and $W(F, C)$ of some evaluation $F(C) \Downarrow D$, where F is a $\mathcal{NSC}$ function, and C, D are S -objects. Except for the rules for *map* and *while*, for every rule of the form:

$$\frac{M_1 \Downarrow C_1, \dots, M_n \Downarrow C_n}{M \Downarrow C}$$

we define:

$$T(M) \stackrel{\text{def}}{=} 1 + \sum_{i=1}^n T(M_i) \quad W(M) \stackrel{\text{def}}{=} \text{SIZE} + \sum_{i=1}^n W(M_i)$$

where *SIZE* is the total size of all S -objects mentioned in the rule (in the premises and the conclusion, including the environments). For the *map*-rule, the definition of W remains the same, while the definition of T becomes:

$$T(M) \stackrel{\text{def}}{=} 1 + \max_{i=1, \dots, n} (T(M_i))$$

(this corresponds to the fact that the function is applied in parallel on all objects in the sequence). For the *while* rule we do not include in *SIZE* the size of the output D (otherwise, the final output D of *while* would be counted as many times as many iterations are performed by *while*). More precisely, if the last rule of $\text{while}(P, F)(C) \Downarrow D$ was:

$$\frac{P(C) \Downarrow \text{true} \quad F(C) \Downarrow C' \quad \text{while}(P, F)(C') \Downarrow D}{\text{while}(P, F)(C) \Downarrow D}$$

then:

$$T(\text{while}(P, F), C) \stackrel{\text{def}}{=} 1 + T(P, C) + T(F, C) + T(\text{while}(P, F), C')$$

$$W(\text{while}(P, F), C) \stackrel{\text{def}}{=} \text{size}(C) + \text{size}(C') + W(P, C) + W(F, C) + W(\text{while}(P, F), C')$$

(i.e. $\text{size}(D)$ is not included explicitly in $W(\text{while}(P, F), C)$)

The language $\mathcal{NSC}$ together with its notions of time and work complexity is a model of parallel computation in its own right but parallel algorithms are most commonly given in terms of one of the several known flavors of PRAM. To facilitate comparisons, we offer the following efficient simulation ($\mathcal{NSC}$'s version of Brent's scheduling principle, as it were):

Proposition 3.2 Any $\mathcal{NSC}$ function of time complexity T and work complexity W can be simulated on a CREW PRAM with scan primitives using p processors with asymptotic complexity $O(T + W/p)$.

Proof. (Sketch) Given some function f in $\mathcal{NSC}$, first flatten f for an extended version of a BVRAM, with unbounded many vector registers and indirect addressing (essentially the VRAM of [Ble90], but with the communication primitives described in section 2). The resulting extended-BVRAM program has the same time and work complexity as f : see remark 7.3. Next use the simulation of an extended BVRAM on a CREW with scan primitives, in the spirit of [Ble90]. We need a CREW instead of a EREW in order to simulate *bm_route* and *sbm_route*. $\square$

4 Expressing map-recursive functions in $\mathcal{NSC}$

Although it is described in a concise, mathematical style (notice that we called it a “calculus” rather than a “language”) $\mathcal{NSC}$ can be easily extended to a more user-friendly language, by allowing a certain amount of block structure: definitions of global/local variables and of nonrecursive functions. There is a straightforward translation of such an extension back into $\mathcal{NSC}$, which we omit from this extended abstract. Accomodating recursive functions though, is a more delicate problem, which we address here.

Consider the following limited form of recursion:

Definition 4.1 A function definition is *map-recursive* if it has the form

$$\text{fun } f(x) = c(x, \text{map}(f)(d(x)))$$

First, it is easy for a compiler to check whether a recursive definition is of this form (in contrast, containment [Ble90] is an undecidable property). Second, this form is general enough to express many existent parallel algorithms: tail recursive definitions, and what is usually meant by divide-and-conquer recursion (for instance the worked example in section 5) are *map-recursive*. Here are some recursion schemata and a sketch of how to convert them into *map-recursive* form (and in the process “parallelize” them) :

$$\begin{aligned} \text{fun } g(x) &= \text{if } p(x) \text{ then } s(x) \text{ else } c(g(d_1(x)), g(d_2(x))) \\ \text{fun } h(x) &= \text{if } p(x) \text{ then } s(x) \text{ else } c(h(d(x))) \\ \text{fun } k(x) &= \text{if } p(x) \text{ then } s(x) \text{ else} \\ &\quad \text{if } p'(x) \text{ then } c(k(d_1(x)), k(d_2(x))) \\ &\quad \text{else } c'(k(d'_1(x)), k(d'_2(x)), k(d'_3(x))) \end{aligned}$$

For g , we construct a list of length 2, and recursively map g on it (Quicksort has this form). For h , the list will have length 1 (tail recursion is a particularization of this form). k is more interesting, since it divides its input into either two or three subproblems. Note that it is not contained [Ble90], so the compilation techniques described here work on some cases on which those of [Ble90] don't. In converting k , the list will have length 1, 3 or 4, where the first element is a tag, and k is slightly modified to return the identity on the tag (a sum of types is used here).

The first of our two main results states that *map-recursion* can be translated (in a source-to-source manner) into a $\mathcal{NSC}$ expression, while preserving its time complexity and “almost” preserving its work complexity.

Theorem 4.2 Consider some function f defined in $\mathcal{NSC}$ extended with *map-recursion*, with time and step complexity T, W . Then, for any $\varepsilon > 0$, one can construct a function f' in $\mathcal{NSC}$ which is equivalent to f and which has time and work complexity $T' = O(T)$ and $W' = O(W^{1+\varepsilon})$ respectively. Moreover, if the divide and conquer tree of f is balanced, then $W' = O(W)$.

Proof. For illustration, we consider only the function g from above. Suppose the types are: $g : s \rightarrow t$, $d_1, d_2 : s \rightarrow s$, and $c : t \times t \rightarrow t$. Not surprisingly, g can be expressed in $\mathcal{NRA}$, without recursion, in two steps, called *divide phase* and *combine phase* in [MH88]:

Divide Phase Start with the singleton sequence $y = [x]$ of type $[s]$, and apply repeatedly the function $\text{flatten} \circ \text{map}(\lambda x. \text{if } p(x) \text{ then } [x] \text{ else } [d_1(x), d_2(x)])$ having the type $[s] \rightarrow [s]$, until all its elements satisfy the predicate p . (We need to tag the elements resulting from $[x]$, to avoid applying p repeatedly on them; we omit the details.) Call y the resulting sequence.

Combine Phase Start by *map*-ing the function s on y , and then apply repeatedly c to adjacent elements of y : some additional bookkeeping is necessary to make sure c is applied to the correct pairs (e.g., it suffices to store the depth in the divide and conquer tree for each element in y , and only combine adjacent elements if they have the same depth). Stop when there is only one element in the resulting list.

Obviously, the translated g will have time complexity $O(T)$. The work complexity is also preserved, in the case in which the divide and conquer tree for the computation of $g(x)$ is perfectly balanced. When the tree is unbalanced, the leaves which are reached sooner have to coexist in the same sequence with those nodes which need more divide steps, thus adding to the total work complexity. Let ν be the number of different levels in the divide and conquer tree which contain leaves. E.g. in an almost perfectly balanced tree, $\nu = 1$ or $\nu = 2$, while in a total “unbalanced” tree, ν can be equal to the total number of leaves, but still $\nu \leq W(g, x)$. We can compute ν in time and work complexity $O(T), O(W)$, by simulating only the *divide* phase, without retaining the results. Let $\varepsilon > 0$. We improve the *divide* phase, such that the time and work complexities of the translation of g into $\mathcal{NSC}$ become $O(T)$ and $O(\nu^\varepsilon W)$ respectively. Namely, we start with $\frac{1}{\varepsilon} + 1$ variables z_i , $i = 0, \dots, \frac{1}{\varepsilon}$, initialized to $[]$, and with y initialized to the singleton $[x]$. We apply repeatedly the divide phase on y ; whenever some leaves are reached, we move them into z_0 . We only allow z_0 to be touched ν^ε times, after which we move its entire content into z_1 , and empty z_0 . We repeat this process, but only allow z_1 to be touched ν^ε times, at which point, we empty z_1 , by moving everything into z_2 . In general, we allow z_i to accumulate only ν^ε times, after which we empty it, by moving everything into z_{i+1} . Obviously, a number of $\nu^{i\varepsilon}$ levels of leaves must be discovered, before making one move into z_i ; thus, z_1 will be filled exactly once, with the leaves from all ν levels. To compute the total additional work complexity, observe that each leaf travels exactly once through $z_0, z_1, \dots, z_{\frac{1}{\varepsilon}}$, and in each z_i is “touched” exactly ν^ε times. Thus, the total work complexity is bounded by $(\frac{1}{\varepsilon} + 1)\nu^\varepsilon W = O(\nu^\varepsilon W)$. Of course, rather complicated bookkeeping is necessary to keep all elements in z_i sorted. The *combine phase* is done similarly, but in reverse. $\square$

The technique of theorem 4.2 seems to extend to more general recursion schemas than the limited recursion. The main kind of recursion to which this technique does not apply is one in which some recursive call to f uses an argument which is computed with a recursive call itself, in the style of the Ackerman function: $A(x, y) = A(x - 1, A(x, y - 1))$. We argue that very few practical algorithms make indeed use of such recursion schemas.

5 An $O(\log n \log \log n)$ Mergesort Algorithm Expressed in $\mathcal{NSC}$

As evidence for the *practical* expressiveness of $\mathcal{NSC}$ we describe in it Valiant’s fast mergesort algorithm [Val75, Jaj92], see the program in figures 1, 2, 3. As we have explained at the beginning of section 4 we are free to use block structure (we choose a syntax close to ML [MTH90]). More importantly, in view of theorem 4.2 we are free to use *map*-recursive definitions, or other recursive schemas which are convertible to *map*-recursion. The main function *mergesort* in figure 1 has the same recursion schema as the function g of section 4 and hence can be converted to a *map*-recursive form. Its parallel time complexity is $O(\log n \log \log n)$.

The fast, $O(\log \log m)$ time *merge* function exhibits a more complicated kind of *map*-recursion. To merge two sequences $A = [a_0, \dots, a_{m-1}]$, $B = [b_0, \dots, b_{n-1}]$, we divide A into $\lceil \sqrt{m} \rceil$ subsequences of length $\leq \sqrt{m}$; let $AA = [A_0, \dots, A_{\sqrt{m}-1}]$ be the resulting nested sequence. Next, we find for each subsequence A_i the corresponding subsequence B_i in B , with which A_i has to be merged, and apply recursively *merge* on all pairs (A_i, B_i) ; let $BB = [B_0, \dots, B_{\sqrt{m}-1}]$. Thus, the general structure of *merge* is:

```
fun merge(A, B) =
  if length(A) ≤ 2 then direct_merge(A, B)
  else let ... compute AA, BB as explained
        in flatten(map(merge)(zip(AA, BB))) end
```

which can be obviously translated into a *map*-recursion.

Figures 2 and 3 contain some auxiliary functions used in *merge*. The function *index(C, I)* expects a sorted sequence of indexes $I = [i_0, \dots, i_{k-1}]$ and, for $C = [C_0, \dots, C_{n-1}]$, returns the sequence $[C_{i_0}, \dots, C_{i_{k-1}}]$: it has constant time complexity and work complexity $= O(n + k)$. The function *index_split(C, I)* splits C according to the indexes in I , again provided that I is sorted, with similar time and work complexity. We use the construct *filter(P) : [t] → [t]*, which for some predicate $P : t \rightarrow \mathbb{B}$ returns the sequence of all elements satisfying P . It is expressible in $\mathcal{NSC}$ by:

```
filter(P)(x) = flatten(map(λu. if P(u) then [u] else [])(x))
```

The functions *first*, *tail*, *last*, *remove_last* and *bm_route* are defined in section 3.

Using the techniques described in [Jaj92], the *merge* function can be transformed to become optimal, i.e. to reduce its work complexity from $O((m + n) \log \log m)$ to $O(m + n)$. This also gives us an optimal (i.e. with $O(n \log n)$ work complexity), $O(\log n \log \log n)$ -time sorting function. The divide-and-conquer trees for both the sorting and the merging function are balanced, hence the translation of theorem 4.2 gives us an optimal $O(\log n \log \log n)$ -time sorting function in $\mathcal{NSC}$.

```

fun mergesort(A) =
  if length(A) ≤ 1 then A
  else let val n = length(A)
        val AA = split(A, [n - n/2, n/2])
        in merge(mergesort(first(AA)),
                  mergesort(last(AA)))
        end

fun merge(A, B) =
  if length(A) ≤ 2 then direct_merge(A, B)
  else
    let val m = length(A)
        val n = length(B)
        val A' = sqrt_positions(A)
        val B' = sqrt_positions(B)
        (* A', B' have lengths  $\sqrt{m}$  and  $\sqrt{n}$  respectively *)
        val R' = direct_rank(A', B')
        val BB1 = sqrt_split(B)
        (* split B into  $\sqrt{n}$  blocks *)
        val a_B = zip(A', index(BB1, R'))
        (* group each a' with its block *)
        val RR' = map(rank_one)(a_B)
        (* rank each a' in its block *)
        val R = map( $\lambda(x, y). (x \dot{-} 1) * \sqrt{n} + y$ )
                (zip(R', RR'))
        val AA = sqrt_split A
        val BB = index_split(B, R)
    in flatten(map(merge)(zip(AA, BB)))
    end

```

Figure 1: Valiant's $O(\log n \log \log n)$ sorting algorithm.

```

fun rank_one(a, B) = length(filter( $\lambda b. b \leq a$ )(B))

fun direct_rank(A, B) = map( $\lambda a. rank\_one(a, B)$ )(A)

fun sqrt_positions(C) =
  let val n = length(C)
      val I = filter( $\lambda i. i \bmod \sqrt{n} = 0$ )(enumerate(C))
  in index(C, I)
  end

fun sqrt_split(C) =
  index_split(C, sqrt_positions(enumerate(C)))

fun direct_merge(A, B) =
  let val R = direct_rank(A, B)
      val BB = index_split(B, R)
  in first(BB)@
     flatten(map( $\lambda(a, B). [a]@B$ )(zip(A, tail(BB))))
  end

```

Figure 2: Auxiliary functions used in *merge*.

```

fun index(C, I) =
  let val n = length(C)
      val k = length(I)
      val zero_to_k = enumerate(I)@[k]
      val delta_I = map( $\dot{-}$ )(zip(I@[n], [0]@I))
      val P = bm_route((C, delta_I), zero_to_k)
      val delta_P = map( $\dot{-}$ )(zip(P, remove_last([0]@P)))
  in bm_route((I, delta_P), C)
  end

fun index_split(C, I) =
  let val n = length(C)
  in split(C, map( $\dot{-}$ )(zip(I@[n], [0]@I)))
  end

```

Figure 3: The functions *index* and *index_split*.

6 Theoretical Expressive Power

In this section we give evidence that $\mathcal{NSC}$ is not too restrictive, as a tool for designing parallel algorithms. Namely, let $\text{CRCW-TIME-PROC}(T(n), P(n))$ be the set of functions computable on a CRCW PRAM in time $T(n)$ using $P(n)$ processors, and $\mathcal{NSC}\text{-TIME-WORK}(T(n), W(n))$ the set of functions expressible in $\mathcal{NSC}$ with time and work complexity $T(n), W(n)$.

Proposition 6.1 *For $T(n), W(n)$, that are suitable (in the sense of [SV84]), we have:*

$$\text{CRCW-TIME-PROC}(O(T(n)), O(W(n))) \subseteq \mathcal{NSC}\text{-TIME-WORK}(O(T(n)), W(n)^{O(1)})$$

More, we get equality, if in the definition of $\mathcal{NSC}$ we restrict the arithmetic operations to the set $\Sigma = \{+, \dot{-}\}$, and if we replace the unit size complexity ($\text{size}(n) \stackrel{\text{def}}{=} 1$ - see section 3) with the logarithmic size complexity ($\text{size}(n) \stackrel{\text{def}}{=} \log n$), in the definition of the work complexity of $\mathcal{NSC}$.

The proof uses a theorem in [SV84], credited to Ruzzo and Tompa, relating CRCW PRAM's to Alternating Turing Machines, and is omitted from this extended abstract. Using the above proposition and proposition 3.2 we can establish that NC coincides with the functions in $\mathcal{NSC}$ with polylogarithmic time and polynomial work complexity. Recall that $\mathcal{NSC}$ is parameterized by a set Σ of arithmetic operations.

Proposition 6.2 *Suppose all arithmetic operations in Σ are in NC. Then:*

$$NC = \mathcal{NSC}\text{-TIME-WORK}(\log^{O(1)} n, n^{O(1)})$$

7 Efficient Compilation of $\mathcal{NSC}$ to BVRAM

Theorem 7.1 (Compilation Theorem) *For every function f in $\mathcal{NSC}$ with time and work complexity T, W , there is a BVRAM, M , such that: $\forall \epsilon > 0$, there is some program P for M , equivalent to f , having time complexity $T' = O(T)$ and $W' = O(W^{1+\epsilon})$.*

Note that, in contrast to theorem 4.2, the number of registers only depends on f and not on ϵ . A *while*-construct can be rewritten as a tail recursive function, hence is *contained*, according to the definition in [Ble90], and therefore the compilation technique described there (for a VRAM, with unbounded many vector registers) preserves its step and work complexity. However, we cannot apply that compilation technique here. Indeed, when viewed as tail recursive function, the work complexity of *while* may increase significantly, because the final result after iterating n steps is touched n additional times, as the tail recursive function returns from its calls. In the definition of the work complexity for *while*, these n additional touches are not counted (see definition 3.1). So the tail recursive translation has a higher work complexity than the original *while* construct. We need a stronger compilation technique in order to stay within the lower work complexity. Moreover, we also only have a bounded number of vector registers.

The proof goes through the following steps:

- *Variable Elimination*. We translate $\mathcal{NSC}$ into a rather similar, but variable free language called Nested Relational Algebra, $\mathcal{NSA}$. The new language only contains functions $fs \rightarrow t$, i.e. no terms. Some term M in $\mathcal{NSC}$, of type t and with free variables $x_1 : s_1, \dots, x_n : s_n$, will be translated into a function $f_M : s_1 \times \dots \times s_n \rightarrow t$ in $\mathcal{NSA}$. The primitive functions and the constructs in $\mathcal{NSA}$ correspond roughly to those in $\mathcal{NSC}$, with only one additional primitive: the function $\rho_2 : s \times [t] \rightarrow [s \times t]$ (see section 3 for its definition). The step and work complexity of functions expressed in $\mathcal{NSC}$ and $\mathcal{NSA}$ are the same. We omit the description of $\mathcal{NSA}$ from this extended abstract; it can be found in appendix C.
- *Flattening*. We define a language for flat sequences, called Sequence Algebra $\mathcal{SA}$, and translate $\mathcal{NSA}$ into $\mathcal{SA}$. Namely, for any $\epsilon > 0$, we show how to translate a function f of $\mathcal{NSA}$ with time and work complexity T, W into an equivalent function in $\mathcal{SA}$ (thus using only flat types), with time and work complexity $O(T)$ and $O(W^{1+\epsilon})$. Of course, any function in $\mathcal{SA}$ can be expressed in $\mathcal{NSA}$ with the same time and work complexity.
- We show that $\mathcal{SA}$ and BVRAM are equivalent, in the sense that any function in $\mathcal{SA}$ can be simulated by a BVRAM with the same time and work complexity, and conversely. One direction of this equivalence helps us completing the compilation, while the other direction allows us to perform optimizations at the level of the language $\mathcal{SA}$, instead of BVRAM.

7.1 The Sequence Algebra, $\mathcal{SA}$

The Sequence Algebra, $\mathcal{SA}$, only has *flat types*. More precisely, we define first *scalar types* by the grammar: $s ::= \text{unit} \mid \mathbb{N} \mid s \times s \mid s + s$, and next define the *flat types* by the grammar: $t ::= \text{unit} \mid [s] \mid t \times t \mid t + t$.

$\mathcal{SA}$ was designed by choosing some set of functions expressible in $\mathcal{NSA}$ (or, equivalent, $\mathcal{NSC}$) over flat types, which seemed to be enough to allow the language $\mathcal{NSA}$ to be translated (flattened) into $\mathcal{SA}$. In addition, $\mathcal{SA}$ is defined in an inductive way, which enables us to prove, by induction, properties about the functions expressible in $\mathcal{SA}$, e.g. lemma 7.2. $\mathcal{SA}$ stands in the same relationship to $\mathcal{NSA}$

as the relational algebra stands to the nested relational algebra [AB88].

Similar to $\mathcal{NSA}$, $\mathcal{SA}$ is a variable-free language, containing some primitive functions, and a set of rules for combining them in order to get more complex functions. We briefly describe $\mathcal{SA}$ below. A complete description of the language can be found in appendix D.

- Error, viewed as a function $\Omega : \text{unit} \rightarrow t$.
- map' s of scalar functions, $\text{map}(\varphi) : [s] \rightarrow [s']$, where $\varphi : s \rightarrow s'$ is a *scalar function*, i.e., informally, a function defined in $\mathcal{NSA}$ (or, equivalently $\mathcal{NSC}$) having only scalar types as input, output, and intermediate types, and without *while*.
- Operations on sequences: the empty sequence $[]$, append $@$, length of a sequence, defined as $\text{length}(x) = [n]$, where n is the length of x , *zip*, *bm_route*, *sbm_route*, selections σ_1, σ_2 (see section 3), and the emptiness test *empty?*, of type $[s] \rightarrow \mathbb{B}$.
- Functions over flat types: the identity $\text{id} : t \rightarrow t$, composition of functions $g \circ f$, projections $\pi_i : t_1 \times t_2 \rightarrow t_i$, pairing of functions (f, g) , injections $\text{in}_i : t_i \rightarrow t_1 + t_2$, and sum of functions $f_1 + f_2 : t_1 + t_2 \rightarrow t$, where $f_i : t_i \rightarrow t$ (an if construct can be derived from this). The latter is defined by: $(f_1 + f_2)(\text{in}_1(x)) \stackrel{\text{def}}{=} f_1(x)$ and $(f_1 + f_2)(\text{in}_2(x)) \stackrel{\text{def}}{=} f_2(x)$.
- Iteration: $\text{while}(p, f)$ is a function of type $t \rightarrow t$, whenever $f : t \rightarrow t$ and $p : t \rightarrow \mathbb{B}$ (recall that $\mathbb{B} = \text{unit} + \text{unit}$ and, thus, is a type of $\mathcal{SA}$).

As for $\mathcal{NSC}$ we define the the time and work complexity for some evaluation $f(C) \Downarrow$, where f is a function in $\mathcal{SA}$ and C is its input (a flat S-object). Note that in the absence of a general *map* there is no nested parallelism in $\mathcal{SA}$.

Although $\mathcal{SA}$ does not contain nested types, like $[\mathbb{N} \times [\mathbb{N} \times \mathbb{N}]]$, it is strong enough to allow such types to be encoded into flat types. The key technical tool for that is to encode some nonflat type $[t]$, where t is a flat type, by some flat type $\text{SEQ}(t)$. For this we use *segment descriptors*, as in [Ble90]. Formally, we transform some *flat type* t into another flat type $\text{SEQ}(t)$, defined by induction on t : (1) $\text{SEQ}(\text{unit}) \stackrel{\text{def}}{=} [\mathbb{N}]$ (2) $\text{SEQ}([s]) \stackrel{\text{def}}{=} [\mathbb{N} \times [s]]$, (3) $\text{SEQ}(t \times t') \stackrel{\text{def}}{=} \text{SEQ}(t) \times \text{SEQ}(t')$, (4) $\text{SEQ}(t + t') \stackrel{\text{def}}{=} [\mathbb{B}] \times \text{SEQ}(t) \times \text{SEQ}(t')$. The idea is that $\text{SEQ}(t)$, although a flat type, can encode sequences of elements from t , i.e. values of type $[t]$. The main technical fact enabling us to prove efficient compilation is the following map lemma.

Lemma 7.2 (The Map Lemma). *Let $f : t \rightarrow t'$ be some function in $\mathcal{SA}$, and let T, W be the time and work complexity of $\text{map}(f)$ (recall that $\text{map}(f)$ is in $\mathcal{NSC}$, but not in $\mathcal{SA}$). Then, for every $\epsilon > 0$, there exists some function $\text{SEQ}(f) : \text{SEQ}(t) \rightarrow \text{SEQ}(t')$ in $\mathcal{SA}$, of time complexity $O(T)$ and work complexity $O(W^{1+\epsilon})$, which simulates $\text{map}(f) : [t] \rightarrow [t']$. More, the structure of $\text{SEQ}(f)$ is independent of ϵ , which implies that “number of vector registers” used by $\text{SEQ}(f)$ is independent of ϵ .*

Proof. (Sketch) This is done by induction on the structure of f . When f is *map* of a scalar function, $\text{SEQ}(f)$ is essentially the same *map*. When f is some operation on

a sequence, we only mention that $SEQ(empty?)$ is essentially a selection, $SEQ(\sigma_1)$ essentially σ_1 , $SEQ(bm_route)$ is a sbm_route , while $SEQ(sbm_route)$ is another sbm_route . The only difficult case is when f is $while(p, g)$. We describe very informally how to compute $SEQ(while(p, g))(x)$, with $x = [x_0, \dots, x_{n-1}]$, of a BVRAM. We could use the same idea as in theorem 4.2, but then the number of registers would depend on ε . Suppose x is in register V_0 . We will use only two additional registers, V_1 and V_2 , which are initially empty. Let t_i be the number of iterations of $while(p, g)(x_i)$, and assume without loss of generality that $t_0 < t_1 < \dots < t_{n-1}$ (we conceptually group all x_i 's having the same t_i), which implies $t_i \geq i$. Let $\delta = n^\varepsilon$, $w_i = W(while(p, g), x_i)$ and $r = \frac{1}{\varepsilon} - 1$. For the moment, assume that in the sequence $x_i, g(x_i), g^{(2)}(x_i), \dots$, the last value (on position t_i) has the smallest size, denoted by s_i , so $s_i t_i \leq w_i$. The simulation proceeds in r stages. The first stage starts by repeatedly applying $SEQ(g)$ on x : whenever some x_i 's reach the end of the iteration, move them into V_1 , until the first $\frac{n}{\delta^r} (\leq n^\varepsilon)$ values are extracted from V_0 , namely $x_i, i = 1, \frac{n}{\delta}$. The additional work complexity due to repeatedly touching the values in V_1 is $O(n^\varepsilon W)$. At this point, we move the entire V_1 into V_2 . For each of the remaining stages $k = 1, r-1$, apply repeatedly $SEQ(g)$ on x , and move, when they terminate, the elements $x_i, i = \frac{n}{\delta^{r-k+1}}, \frac{n}{\delta^{r-k}}$ from V_0 to V_1 : at the end of stage k , we move the entire V_1 into V_2 . The additional work complexity due to repeatedly touching some element x_i in V_1 at this stage is $\leq s_i \frac{n}{\delta^{r-k}}$. But since $i \geq \frac{n}{\delta^{r-k+1}}$, we have that $t_i \geq i \geq \frac{n}{\delta^{r-k+1}} \frac{1}{\delta}$, hence the additional work complexity for x_i is $\leq s_i t_i \delta \leq w_i n^\varepsilon$, which, when added up, accounts for only $O(n^\varepsilon W)$ for stage k , which adds up to at most $O(\frac{1}{\varepsilon} n^\varepsilon W) = O(n^\varepsilon W)$ for all r stages. During all r stages, V_2 is touched only r times, for an additional $O(W)$ work complexity. At the end of the last stage, all x_i 's ($i = 1, n$) end up in V_2 , so V_2 contains the result of $SEQ(while(p, g))(x)$.

Finally we have to show how to define $SEQ(while(p, g))(x)$ in the general case, when the sequence $x_i, g(x_i), g^{(2)}(x_i), \dots, g^{(t_i)}(x_i)$ has a minimum size on some position m_i which is not necessarily the last one. In that case we first compute m_i , for each i : this can be done with complexities $O(T)$ and $O(W)$, by simply applying $SEQ(g)$ repeatedly, and eliminating those elements which reach the end of their iteration. Next we split the whole iteration $SEQ(while(p, g))(x)$ in two parts, essentially by synchronizing the n parallel iterations at the moment when they reach their minimum size, namely: (1) perform the n parallel iterations, as described above, but stop the iteration over x_i at step m_i , (2) continue the n parallel iterations, from step m_i to t_i , using the same technique, but in reverse (because now the minimum sizes are at the beginning). $\square$

Remark 7.3 Had we had arbitrarily many registers instead of a bounded number, we could have designed $SEQ(f)$ with time and work complexity $O(T)$ and $O(W)$ (instead of $O(T)$ and $O(W^{1+\varepsilon})$), which is used in the proof of proposition 3.2. Indeed, for $f = while(p, g)$, assume again that, $\forall i = 1, n$, the smallest size, denoted s_i , in the sequence $x_i, g(x_i), g^{(2)}(x_i), \dots, g^{(t_i)}(x_i)$ is on the last position. Then $SEQ(while(p, g))$ is simulated by placing, upon completion, each element x_i in some different register V_i . At the end we have to combine the registers $V_1, \dots, V_n$, which we do in the following order: combine V_n with V_{n-1} , the result with V_{n-2} , $\dots$, the

result with V_1 . The additional work complexity for the combine phase due to x_i is $s_i i$, which is bounded by w_i , because of our assumption about s_i . We can extend the simulation to the case when the smallest sizes s_i are reached at arbitrary moments, using the same technique as above.

Finally, we flatten the language $\mathcal{NSA}$ into $\mathcal{SA}$. We start by flattening the types. For every type s of $\mathcal{NSA}$, we define $COMPILE(s)$ to be a flat type, which encodes s . Namely:

$$\begin{aligned} COMPILE(unit) &\stackrel{\text{def}}{=} unit \\ COMPILE(\mathbb{N}) &\stackrel{\text{def}}{=} [\mathbb{N}] \\ COMPILE(s \times s') &\stackrel{\text{def}}{=} COMPILE(s) \times COMPILE(s') \\ COMPILE(s + s') &\stackrel{\text{def}}{=} COMPILE(s) + COMPILE(s') \\ COMPILE([s]) &\stackrel{\text{def}}{=} SEQ(COMPILE(s)) \end{aligned}$$

Also, we define the functions $encode_s : s \rightarrow COMPILE(s)$ and $decode_s : COMPILE(s) \rightarrow s$ in $\mathcal{NSA}$, with time complexity $O(1)$ and work complexity linear in the size of the input, with the property $decode_s(encode_s(x)) = x$, for every $x \in s$. The definition of the functions $encode$ and $decode$ are rather standard, and are omitted from this extended abstract.

Finally, we can prove:

Proposition 7.4 Let $f : s \rightarrow s'$ be some function in $\mathcal{NSA}$ with time and work complexity T, W . Then, for every $\varepsilon > 0$, there is some function $COMPILE(f) : COMPILE(s) \rightarrow COMPILE(s')$ in $\mathcal{SA}$ which “simulates f ”, i.e. for every x , $COMPILE(f)(encode(x)) = encode(f(x))$, with time and work complexity $O(T)$, $O(W^{1+\varepsilon})$. Moreover, f' requires “the same number of BVRAM registers” for every ε .

Proof. (Sketch) By induction on the structure of f . All cases are straightforward, except for the case when $f = map(g)$, where we use the Map lemma. $\square$

7.2 Equivalence of $\mathcal{SA}$ and BVRAM

The types in $\mathcal{SA}$ are slightly richer than those of the BVRAM: $\mathcal{SA}$ allows for types like $[unit + \mathbb{N} + \mathbb{N} \times \mathbb{N}] + [\mathbb{N} \times \mathbb{N}] \times [\mathbb{N}] + unit$, while the types on the BVRAM are only of the form $[\mathbb{N}] \times \dots \times [\mathbb{N}]$. However, encoding of $\mathcal{SA}$ types into BVRAM types is straightforward.

Proposition 7.5 $\mathcal{SA}$ and BVAM are equivalent, i.e. any function f in $\mathcal{SA}$ with time and work complexity T, W can be simulated on a BVRAM with the same time and work complexity, and conversely.

Proof. Simulating some function of $\mathcal{SA}$ by a BVRAM program is easily done by induction on the structure of that function. The converse is slightly more involved. Indeed, let r be the number of registers of a BVRAM M , and h some function in $\mathcal{SA}$ of type $[\mathbb{N}] \times ([\mathbb{N}])^r \rightarrow [\mathbb{N}] \times ([\mathbb{N}])^r$ performing one step of the program of M (where the program counter is encoded by a singleton sequence, on the first position). By iterating h we indeed achieve the desired time complexity, but not the work complexity, since at each step, the function h touches all r registers. To avoid this, we define a sequence of r functions $f_i, i = 1, r$. The inputs and outputs for f_i are: the values of the i “smallest” registers, at some particular moment, the indexes of these i registers, the size S of the

next largest register, and the program counter. f_i iterates the one-step function as long as it only affects the i registers it sees, and as long as all the i sizes stay less than S . If any of these conditions is violated, f_i stops. To do its job, f_i calls f_{i-1} , which iterates steps on M by only looking at the smallest $i-1$ registers: when f_{i-1} finishes, f_i tries to do one more step by taking into account the i 's smallest register as well, which f_{i-1} ignores. If it cannot, then it returns (to f_{i+1}). Else, it performs the operation, and calls f_{i-1} again, possibly with a different set of $i-1$ registers, from the set of i registers it sees. $\square$

Although only one direction of proposition is actually needed for the compilation theorem 7.1, the converse is significant from the point of view of optimizations: it implies that any optimizations done for the BVRAM can also be performed at the level of the $\mathcal{SA}$ language.

8 Conclusions

We intend to use $\mathcal{NSC}$ as a core for a “real” parallel language for querying nested collections, by adding proven features such as those encountered in functional languages like ML. Guaranteed complexity bounds such as those emerging from this paper can serve as useful guidelines for language design, especially in the database area. Of course, the techniques we have used in the translation of *map*-recursion and in the unnesting of nested parallelism need to be validated by practical implementations. Equally important is to continue to investigate the practical expressiveness of $\mathcal{NSC}$ by attempting to represent various known efficient parallel algorithms. Another direction of investigation is to develop optimization techniques for this language by using ideas that have been proved useful in databases.

References

- [AB88] Serge Abiteboul and Catriel Beeri. On the power of languages for the manipulation of complex objects. In *Proceedings of International Workshop on Theory and Applications of Nested Relations and Complex Objects*, Darmstadt, 1988. Also available as INRIA Technical Report 846.
- [BBW92] Val Breazu-Tannen, Peter Buneman, and Limsoon Wong. Naturally embedded query languages. In J. Biskup and R. Hull, editors, *LNCS 646: Proceedings of 4th International Conference on Database Theory, Berlin, Germany, October, 1992*, pages 140–154. Springer-Verlag, October 1992. Available as UPenn Technical Report MS-CIS-92-47.
- [Ble90] Guy E. Blelloch. *Vector Models for Data-Parallel Computing*. MIT Press, Cambridge, Massachusetts, 1990.
- [Ble93] Guy Blelloch. NESL: A nested data-parallel language. Technical Report CMU-CS-93-129, Carnegie Mellon University, Pittsburgh, PA 15213, 1993.
- [BS90] Guy Blelloch and Gary Sabot. Compiling collection-oriented languages onto massively parallel computers. *Journal of Parallel and Distributed Computing*, 8:119–134, 1990.
- [BTS91] V. Breazu-Tannen and R. Subrahmanyam. Logical and computational aspects of programming with Sets/Bags/Lists. In *LNCS 510: Proceedings of 18th International Colloquium on Automata, Languages, and Programming, Madrid, Spain, July 1991*, pages 60–75. Springer Verlag, 1991.
- [Cur88] P. L. Curien. The $\lambda\rho$ -calculus: An abstract framework for environment machines. Technical Report URA 725, Laboratoire d’Informatique, Departement de Mathematiques et d’Informatique, Ecole Normale Supérieure, 45 Rue d’Ulm, 75230 Paris Cedex 05, France, 1988.
- [HS86] Daniel Hillis and Guy Steele. Data parallel algorithms. *Communications of the ACM*, 29(12):1170–1183, 1986.
- [Jaj92] Joseph Jaja. *An Introduction to Parallel Algorithms*. Addison-Wesley, 1992.
- [Kah87] Gilles Kahn. Natural semantics. In *Proceedings of Symposium on Theoretical Aspects of Computer Science*, pages 22–39. Springer-Verlag, 1987.
- [KLGLS90] Kathleen Knobe, Joan D. Lukas, and Jr. Guy L. Steele. Data optimization: Allocation of arrays to reduce communication on SIMD machines. *Journal of Parallel and Distributed Computing*, 8:102–118, 1990.
- [Lei92] Frank Thomson Leighton. *Introduction to Parallel Algorithms and Architectures : Arrays, Trees, Hypercubes*. Morgan Kaufmann Publishers, 1992.
- [MH88] Zhijing Mou and Paul Hudak. An algebraic model for divide-and-conquer and its parallelism. *Journal of Supercomputing*, 2:257–278, 1988.
- [MTH90] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. MIT Press, 1990.
- [SBT94] Dan Suciu and Val Breazu-Tannen. A query language for NC. In *Proceedings of 13th ACM Symposium on Principles of Database Systems*, Minneapolis, Minnesota, May 1994. To appear. See also UPenn Technical Report MS-CIS-94-05.
- [SV84] Larry Stockmeyer and Uzi Vishkin. Simulation of parallel random access machines by circuits. *SIAM Journal of Computing*, 13:409–422, May 1984.
- [Val75] L. G. Valiant. Parallelism in comparison problems. *SIAM Journal of Computing*, 4(3):348–355, 1975.

A The Nested Sequence Calculus $\mathcal{NSC}$

We define a **type context** Γ to be a set of the form $\Gamma = \{x_1 : s_1, \dots, x_n : s_n\}$, where x_i are variables and s_i are types. We write $\Gamma \triangleright M : t$, or $\Gamma \triangleright F : s \rightarrow t$, when we want to say that, under the type assumptions of Γ , the term M has type t , or the function F has type $s \rightarrow t$. Below are the rules defining the language. Recall that $\mathbf{B} \stackrel{\text{def}}{=} \text{unit} + \text{unit}$.

Variables, Errors, Constants, Arithmetic

$$\frac{}{x : t, \Gamma \triangleright x : t} \quad \frac{}{\Gamma \triangleright \Omega^t : t} \quad \frac{}{\Gamma \triangleright n : \mathbb{N}} \quad (n \in \mathbb{N})$$

$$\frac{\Gamma \triangleright M : \mathbb{N} \quad \Gamma \triangleright N : \mathbb{N}}{\Gamma \triangleright M \text{ op } N : \mathbb{N}} \quad (op \in \Sigma) \quad \frac{\Gamma \triangleright M : \mathbb{N} \quad \Gamma \triangleright N : \mathbb{N}}{\Gamma \triangleright M = N : \text{bool}}$$

Type products

$$\frac{}{\Gamma \triangleright () : \text{unit}} \quad \frac{\Gamma \triangleright M : s, \Gamma \triangleright N : t}{\Gamma \triangleright (M, N) : s \times t} \quad \frac{\Gamma \triangleright M : s \times t}{\Gamma \triangleright \pi_1(M) : s} \quad \frac{\Gamma \triangleright M : s \times t}{\Gamma \triangleright \pi_2(M) : t}$$

Type sums

$$\frac{\Gamma \triangleright M : s}{\Gamma \triangleright \text{in}_1(M) : s + t} \quad \frac{\Gamma \triangleright M : t}{\Gamma \triangleright \text{in}_2(M) : s + t} \quad \frac{\Gamma \triangleright M : s + t \quad x : s, \Gamma \triangleright N : u \quad y : t, \Gamma \triangleright P : u}{\Gamma \triangleright \text{case } M \text{ of } \text{in}_1(x) \Rightarrow N \mid \text{in}_2(y) \Rightarrow P : u}$$

Functions

$$\frac{x : s, \Gamma \triangleright M : t}{\Gamma \triangleright \lambda x : s. M : s \rightarrow t} \quad \frac{\Gamma \triangleright F : s \rightarrow t, M : s}{\Gamma \triangleright F(M) : t}$$

Iteration

$$\frac{\Gamma \triangleright P : t \rightarrow \text{bool} \quad \Gamma \triangleright F : t \rightarrow t}{\Gamma \triangleright \text{while}(P, F) : t \rightarrow t}$$

Collections

$$\frac{}{\Gamma \triangleright [] : [t]} \quad \frac{\Gamma \triangleright M : t}{\Gamma \triangleright [M] : [t]} \quad \frac{\Gamma \triangleright M : [t] \quad \Gamma \triangleright N : [t]}{\Gamma \triangleright M @ N : [t]} \quad \frac{\Gamma \triangleright M : [[t]]}{\Gamma \triangleright \text{flatten}(M) : [t]}$$

$$\frac{\Gamma \triangleright M : [t]}{\Gamma \triangleright \text{length}(M) : \mathbb{N}} \quad \frac{\Gamma \triangleright M : [t]}{\Gamma \triangleright \text{get}(M) : t} \quad \frac{\Gamma \triangleright F : s \rightarrow t}{\Gamma \triangleright \text{map}(F) : [s] \rightarrow [t]}$$

Sequences

$$\frac{\Gamma \triangleright M : [s] \quad \Gamma \triangleright N : [t]}{\Gamma \triangleright \text{zip}(M, N) : [s \times t]} \quad \frac{\Gamma \triangleright M : [t]}{\Gamma \triangleright \text{enumerate}(M) : [\mathbb{N}]} \quad \frac{\Gamma \triangleright M : [t] \quad \Gamma \triangleright N : [\mathbb{N}]}{\Gamma \triangleright \text{split}(M, N) : [[t]]}$$

Weakening

$$\frac{\Gamma \triangleright M : t}{x : s, \Gamma \triangleright M : t}$$

B Operational Semantics

We define an **environment** to be a finite set of the form $\rho = \{x_1 = C_1, \dots, x_n = C_n\}$, where $x_1, \dots, x_n$ are variables, and $C_1, \dots, C_n$ are S-objects. We say that ρ is associated to some type context Γ iff ρ and Γ mention exactly the same variables and if the type of C_i is the type of the variable x_i in Γ .

The following rules define the ternary relation $\rho \bullet M \Downarrow C$ and the 4-ary relation $\rho \bullet F(C) \Downarrow C'$, where ρ is associated to some type context Γ such that $\Gamma \triangleright M : t$, or $\Gamma \triangleright f : s \rightarrow t$ respectively.

Variables, Errors, Constants, Arithmetic

$$\frac{}{x = C, \rho \bullet x \Downarrow C} \quad \frac{}{\rho \bullet n \Downarrow n}$$

$$\frac{\rho \bullet M \Downarrow m \quad \rho \bullet N \Downarrow n}{\rho \bullet M + N \Downarrow m + n} \quad (\text{similar for all } op \in \Sigma \text{ and } =)$$

Type products

$$\frac{}{\rho \bullet () \Downarrow ()} \quad \frac{\rho \bullet M \Downarrow C \quad \rho \bullet N \Downarrow D}{\rho \bullet (M, N) \Downarrow (C, D)} \quad \frac{\rho \bullet M \Downarrow (C, D)}{\pi_1(M) \Downarrow C} \quad \frac{\rho \bullet M \Downarrow (C, D)}{\pi_2(M) \Downarrow D}$$

Type sums

$$\frac{\rho \bullet M \Downarrow C}{\rho \bullet \text{in}_1(M) \Downarrow \text{in}_1(C)} \quad \frac{\rho \bullet M \Downarrow C}{\rho \bullet \text{in}_2(M) \Downarrow \text{in}_2(C)} \quad \frac{\rho \bullet M \Downarrow \text{in}_1(C) \quad x = C, \rho \bullet N \Downarrow D}{\rho \bullet \text{case } M \text{ of } \text{in}_1(x) \Rightarrow N \mid \text{in}_2(x) \Rightarrow P \Downarrow D}$$

Functions

$$\frac{\rho \bullet N \Downarrow C \quad \rho \bullet F(C) \Downarrow D}{\rho \bullet F(N) \Downarrow D} \quad \frac{x = C, \rho \bullet M \Downarrow D}{\rho \bullet (\lambda x. M)(C) \Downarrow D}$$

Iteration

$$\frac{\rho \bullet P(C) \Downarrow \text{false}}{\rho \bullet \text{while}(P, F)(C) \Downarrow C} \quad \frac{\rho \bullet P(C) \Downarrow \text{true} \quad \rho \bullet F(C) \Downarrow C' \quad \rho \bullet \text{while}(P, F)(C') \Downarrow D}{\rho \bullet \text{while}(P, F)(C) \Downarrow D}$$

Collections

$$\frac{}{\rho \bullet [] \Downarrow []} \quad \frac{\rho \bullet M \Downarrow C}{\rho \bullet [M] \Downarrow [C]} \quad \frac{\rho \bullet M \Downarrow [C_0, \dots, C_{m-1}] \quad \rho \bullet N \Downarrow [D_0, \dots, D_{n-1}]}{\rho \bullet M @ N \Downarrow [C_0, \dots, C_{m-1}, D_0, \dots, D_{n-1}]}$$

$$\frac{\rho \bullet M \Downarrow [[C_{00}, C_{01}, \dots], [C_{10}, C_{11}, \dots], \dots]}{\rho \bullet \text{flatten}(M) \Downarrow [C_{00}, C_{01}, \dots, C_{10}, C_{11}, \dots]} \quad \frac{\rho \bullet M \Downarrow [C_0, \dots, C_{n-1}]}{\rho \bullet \text{length}(M) \Downarrow n}$$

$$\frac{\rho \bullet M \Downarrow [C]}{\rho \bullet \text{get}(M) \Downarrow C} \quad \frac{\rho \bullet F(C_0) \Downarrow D_0 \quad \dots \quad \rho \bullet F(C_{n-1}) \Downarrow D_{n-1}}{\rho \bullet \text{map}(F)([C_0, \dots, C_{n-1}]) \Downarrow [D_0, \dots, D_{n-1}]}$$

Sequences

$$\frac{\rho \bullet M \Downarrow [C_0, \dots, C_{n-1}] \quad \rho \bullet N \Downarrow [D_0, \dots, D_{n-1}]}{\rho \bullet \text{zip}(M, N) \Downarrow [(C_0, D_0), \dots, (C_{n-1}, D_{n-1})]} \quad \frac{\rho \bullet M \Downarrow [C_0, \dots, C_{n-1}]}{\rho \bullet \text{enumerate}(M) \Downarrow [0, \dots, n-1]}$$

$$\frac{\rho \bullet M \Downarrow [C_0, \dots, C_{n_0+\dots+n_{m-1}}] \quad \rho \bullet N \Downarrow [n_0, \dots, n_{m-1}]}{\rho \bullet \text{split}(M, N) \Downarrow [[C_0, \dots, C_{n_0-1}], [C_{n_0}, \dots, C_{n_0+n_1-1}], \dots, [C_{n_0+\dots+n_{m-2}}, \dots, C_{n_0+\dots+n_{m-1}}]]}$$

Weakening

$$\frac{\rho \bullet M \Downarrow C}{x = C', \rho \bullet M \Downarrow C} \quad \frac{\rho \bullet F(C) \Downarrow D}{x = C', \rho \bullet F(C) \Downarrow D}$$

C The Nested Sequence Algebra $\mathcal{NSA}$

Errors, Constants, Arithmetic

$$\frac{}{\Omega^t : \text{unit} \rightarrow t} \quad \frac{n \in \mathbb{N}}{n : \text{unit} \rightarrow \mathbb{N}} \quad \frac{op \in \Sigma}{op : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}} \quad \frac{}{=: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{B}}$$

Function identity and composition

$$\frac{}{id_t : t \rightarrow t} \quad \frac{f : r \rightarrow s \quad g : s \rightarrow t}{g \circ f : r \rightarrow t}$$

Type products

$$\frac{}{!_t : t \rightarrow \text{unit}} \quad \frac{f_1 : s \rightarrow t_1 \quad f_2 : s \rightarrow t_2}{(f_1, f_2) : s \rightarrow t_1 \times t_2} \quad \frac{}{\pi_1 : t_1 \times t_2 \rightarrow t_1} \quad \frac{}{\pi_2 : t_1 \times t_2 \rightarrow t_2}$$

Type sums

$$\frac{}{in_1 : t_1 \rightarrow t_1 + t_2} \quad \frac{}{in_2 : t_2 \rightarrow t_1 + t_2} \quad \frac{f_1 : s_1 \rightarrow t \quad f_2 : s_2 \rightarrow t}{f_1 + f_2 : s_1 + s_2 \rightarrow t} \quad \frac{}{\delta : (t_1 + t_2) \times t \rightarrow t_1 \times t + t_2 \times t}$$

Iteration

$$\frac{p : t \rightarrow \mathbb{B} \quad f : t \rightarrow t}{\text{while}(p, f) : t \rightarrow t}$$

Collections

$$\frac{}{[] : \text{unit} \rightarrow [t]} \quad \frac{}{\text{singleton} : t \rightarrow [t]} \quad \frac{}{@@ : [t] \times [t] \rightarrow [t]} \quad \frac{}{\text{flatten} : [[t]] \rightarrow [t]} \\ \frac{}{\text{length} : [t] \rightarrow \mathbb{N}} \quad \frac{}{\text{get} : [t] \rightarrow t} \quad \frac{f : s \rightarrow t}{\text{map}(f) : [s] \rightarrow [t]}$$

Sequences

$$\frac{}{\text{zip} : [s] \times [t] \rightarrow [s \times t]} \quad \frac{}{\text{enumerate} : [t] \rightarrow \mathbb{N}} \quad \frac{}{\text{split} : [t] \times [\mathbb{N}] \rightarrow [[t]]}$$

Broadcast This replaces the “free variables” present in $\mathcal{NSC}$.

$$\frac{}{\rho_2 : s \times [t] \rightarrow [s \times t]}$$

The evaluation relation $f(C) \Downarrow C'$, for f some function in $\mathcal{NSA}$ of type $s \rightarrow t$ and C, C' S-objects of type s and t respectively, is defined in a way similar to the definition for $\mathcal{NSC}$, but simpler because functions in $\mathcal{NSA}$ do not have free variables, hence there is no need for an environment. The time and work complexity $T(f, C)$ and $W(f, C)$ are defined accordingly.

Proposition C.1 *Any closed function $f \in \mathcal{NSC}$ with time and work complexity T, W is expressible in $\mathcal{NSA}$ by some function f' with time and work complexity $O(T), O(W)$, and vice versa. Thus, $\mathcal{NSC}$ and $\mathcal{NSA}$ have the same expressive power.*

D The Sequence Algebra $\mathcal{SA}$

Scalar types are: $s ::= \text{unit} \mid \mathbb{N} \mid s \times s \mid s + s$. Scalar functions $\varphi : s \rightarrow s'$ are given by:

Constants, Arithmetic

$$\frac{n \in \mathbb{N}}{n : \text{unit} \rightarrow \mathbb{N}} \quad \frac{op \in \Sigma}{op : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}} \quad \frac{}{=: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{B}}$$

Function identity and composition

$$\frac{}{id_s : s \rightarrow s} \quad \frac{\varphi : s \rightarrow s' \quad \psi : s' \rightarrow s''}{\psi \circ \varphi : s \rightarrow s''}$$

Scalar type products

$$\frac{}{!_s : s \rightarrow \text{unit}} \quad \frac{}{\pi_1 : s_1 \times s_2 \rightarrow s_1} \quad \frac{}{\pi_2 : s_1 \times s_2 \rightarrow s_2} \quad \frac{\varphi_1 : s \rightarrow s_1 \quad \varphi_2 : s \rightarrow s_2}{(\varphi_1, \varphi_2) : s \rightarrow s_1 \times s_2}$$

Scalar type sums

$$\frac{}{in_1 : s_1 \rightarrow s_1 + s_2} \quad \frac{}{in_2 : s_2 \rightarrow s_1 + s_2} \quad \frac{\varphi_1 : s_1 \rightarrow s \quad \varphi_2 : s_2 \rightarrow s}{\varphi_1 + \varphi_2 : s_1 + s_2 \rightarrow s} \quad \frac{}{\delta : (s_1 + s_2) \times s \rightarrow s_1 \times s + s_2 \times t}$$

(Cont'd next page)

Flat types are: $t ::= \text{unit} \mid [s] \mid t \times t \mid t + t$. Functions in $\mathcal{SA}$ $f : t \rightarrow t'$ are given by:

Errors and Scalar operations

$$\frac{}{\Omega^f : \text{unit} \rightarrow t} \quad \frac{\varphi : s \rightarrow s' \text{ a scalar function}}{\text{map}(\varphi) : [s] \rightarrow [s']}$$

Function identity and composition

$$\frac{}{id_t : t \rightarrow t} \quad \frac{f : t \rightarrow t' \quad g : t' \rightarrow t''}{g \circ f : t \rightarrow t''}$$

Flat type products

$$\frac{}{!_t : t \rightarrow \text{unit}} \quad \frac{}{\pi_1 : t_1 \times t_2 \rightarrow t_1} \quad \frac{}{\pi_2 : t_1 \times t_2 \rightarrow t_2} \quad \frac{f_1 : t \rightarrow t_1 \quad f_2 : t \rightarrow t_2}{(f_1, f_2) : t \rightarrow (t_1, t_2)}$$

Flat type sums

$$\frac{}{in_1 : t_1 \rightarrow t_1 + t_2} \quad \frac{}{in_2 : t_2 \rightarrow t_1 + t_2} \quad \frac{f_1 : t_1 \rightarrow t \quad f_2 : t_2 \rightarrow t}{f_1 + f_2 : t_1 + t_2 \rightarrow t} \quad \frac{}{\delta : (t_1 + t_2) \times t \rightarrow t_1 \times t + t_2 \times t}$$

Iterations

$$\frac{p : t \rightarrow \mathbb{B} \quad f : t \rightarrow t}{\text{while}(p, f) : t \rightarrow t}$$

Collections

$$\frac{}{[] : \text{unit} \rightarrow [s]} \quad \frac{}{\text{singleton} : \text{unit} \rightarrow [\text{unit}]} \quad \frac{}{@ : [s] \times [s] \rightarrow [s]} \quad \frac{}{\text{length} : [s] \rightarrow [\mathbb{N}]} \\ \frac{}{\text{empty?} : [s] \rightarrow \mathbb{B}} \quad \frac{}{\sigma_1 : [s_1 + s_2] \rightarrow [s_1]} \quad \frac{}{\sigma_2 : [s_1 + s_2] \rightarrow [s_2]}$$

Sequences

$$\frac{}{\text{zip} : [s] \times [s'] \rightarrow [s \times s']} \quad \frac{}{\text{enumerate} : [s] \rightarrow [\mathbb{N}]} \quad \frac{}{\text{bm_route} : ([s] \times [\mathbb{N}]) \times [s'] \rightarrow [s']} \\ \frac{}{\text{sbm_route} : ([s] \times [\mathbb{N}]) \times ([s'] \times [\mathbb{N}]) \rightarrow [s']}$$

Example D.1 Informally we show how to compute $\text{combine} : [\mathbb{B}] \times [s] \times [s] \rightarrow [s]$, where $\text{combine}(f, x, y)$ combines the lists x and y , according to the flags given by f . The resulting list will have the same length as f , and will contain some x_i on those positions where f is true, and some y_j where f is false. E.g. when $f = [\text{true}, \text{false}, \text{false}, \text{true}, \text{false}, \text{true}, \text{true}]$ and $x = [x_0, x_1, x_2, x_3]$, $y = [y_0, y_1, y_2]$, then $\text{combine}(f, x, y)$ must be $[x_0, y_0, y_1, x_1, y_2, x_2, x_3]$. To compute combine in $\mathcal{SA}$, start by enumerating f , to get $[0, 1, 2, 3, 4, 5, 6]$, and by transforming the booleans into 0 and 1, to get $[1, 0, 0, 1, 0, 1, 1]$. Now apply bm_route to select from the first list those elements having a 1 in the second list, and obtain $[0, 3, 5, 6]$. Similarly, we obtain $[1, 2, 4]$. These two lists tell us on which position each element of x and y must end up. Next, we subtract each number in this list from its right neighbor (by considering $7 = \text{length}(f)$ to be the right neighbor of the last element), with the exception of the first position, where we also add the number itself. I.e., we get: $[0 + (3 - 0), 5 - 3, 6 - 5, 7 - 6] = [3, 2, 1, 1]$ and $[1 + (2 - 1), 4 - 2, 7 - 4] = [2, 2, 3]$. Now we bm_route x and y , using these two lists as replication sequences, and get $[x_0, x_0, x_0, x_1, x_1, x_2, x_3]$ and $[y_0, y_0, y_1, y_1, y_2, y_2, y_2]$ respectively (both have the length of f). Finally, we zip them together with f , and map some scalar function which selects x_i or y_i according to the flag.