

Abstract State Machines Capture Parallel Algorithms: Correction and Extension

ANDREAS BLASS

University of Michigan

and

YURI GUREVICH

Microsoft Research

We consider parallel algorithms working in sequential global time, for example, circuits or parallel random access machines (PRAMs). Parallel abstract state machines (parallel ASMs) are such parallel algorithms, and the parallel ASM thesis asserts that every parallel algorithm is behaviorally equivalent to a parallel ASM. In an earlier article, we axiomatized parallel algorithms, proved the ASM thesis, and proved that every parallel ASM satisfies the axioms. It turned out that we were too timid in formulating the axioms; they did not allow a parallel algorithm to create components on the fly. This restriction did not hinder us from proving that the usual parallel models, like circuits or PRAMs or even alternating Turing machines, satisfy the postulates. But it resulted in an error in our attempt to prove that parallel ASMs always satisfy the postulates. To correct the error, we liberalize our axioms and allow on-the-fly creation of new parallel components. We believe that the improved axioms accurately express what parallel algorithms ought to be. We prove the parallel thesis for the new, corrected notion of parallel algorithms, and we check that parallel ASMs satisfy the new axioms.

Categories and Subject Descriptors: D.1.3 [**Programming Techniques**]: Concurrent Programming—*Parallel programming*; F.1.1 [**Computation by Abstract Devices**]: Models of Computation—*Unbounded-action devices (e.g., cellular automata, circuits, networks of machines)*; F.1.2 [**Computation by Abstract Devices**]: Modes of Computation—*Alternation and nondeterminism, parallelism and concurrency*; I.6.8 [**Simulation and Modeling**]: Types of Simulation—*Parallel*

General Terms: Algorithms, Languages, Theory

Additional Key Words and Phrases: Parallel algorithm, abstract state machine, ASM thesis, postulates for parallel computation, parallel programming

The work of the first author was partially supported by NSF grant DMS-0070723 and by a grant from Microsoft Research.

Authors' addresses: A. Blass, Mathematics Department, University of Michigan, Ann Arbor, MI 48109-1043; email: ablass@umich.edu; Y. Gurevich, Microsoft Research, One Microsoft Way, Redmond, WA 98052; email: gurevich@microsoft.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2008 ACM 1529-3785/2008/06-ART19 \$5.00 DOI 10.1145/1352582.1352587 <http://doi.acm.org/10.1145/1352582.1352587>

ACM Reference Format:

Blass, A. and Gurevich, Y. 2008. Abstract state machines capture parallel algorithms: Correction and extension. *ACM Trans. Comput. Logic*, 9, 3, Article 19 (June 2008), 32 pages. DOI = 10.1145/1352582.1352587 <http://doi.acm.org/10.1145/1352582.1352587>

1. INTRODUCTION

Sequential algorithms, like C programs or sequential abstract state machines (sequential ASMs), work in small steps, that is, steps of bounded complexity. Parallel algorithms, like circuits or parallel random access machines (PRAMs) or parallel ASMs, work in wide shallow steps. The steps are wide in the sense that the algorithm has no fixed bound on the number of components executing in parallel during a single step. The steps are shallow in the sense that the algorithm has a fixed bound, independent of the input or state, on the number of actions executed sequentially during a step.

The sequential ASM thesis asserts that every sequential algorithm is behaviorally equivalent to a sequential ASM. In Gurevich [2000], the second author axiomatized sequential algorithms by means of three convincing postulates, proved the thesis, and checked that every sequential ASM satisfies the postulates. In Blass and Gurevich [2003], we extended that line of investigation to the more challenging case of parallel algorithms. We axiomatized parallel algorithms by means of pretty convincing postulates and we proved the thesis for this case: every parallel algorithm is behaviorally equivalent to a parallel ASM.

Our postulates were, of course, intended to describe what happens in any sort of parallel computation, but it turned out later that they were too restrictive in one respect. They did not allow creation of new proclats (the basic sequential algorithms from which the parallel ones are built up) on the fly during a computation step. The set of proclats could be changed only by updates of the state, as part of the transition at the end of a step.

This restriction on proclat creation did not hinder us from proving that the usual parallel models, like circuits or PRAMs or even alternating Turing machines, satisfy the postulates. But it resulted in an error in our attempt to prove that parallel ASMs always satisfy the postulates. The other parallel computation models we considered did not require on-the-fly proclat creation but parallel ASMs did.

Here we liberalize accordingly the postulates of Blass and Gurevich [2003] and thus expand the notion of parallel algorithms. We show that the main theorem, the ASM thesis for parallel algorithms (Blass and Gurevich [2003], Theorem 10.1), remains true for the new, larger class of parallel algorithms. And we check that parallel ASMs satisfy the new postulates. We use the occasion to correct a couple of other, smaller errors in Blass and Gurevich [2003].

In Section 2, we describe the problems with the examples in Sections 8.4 and 8.5 of Blass and Gurevich [2003], we show how all but one of the problems can easily be corrected, and we sketch the idea needed to correct the one remaining problem. In Section 3 we develop this idea in detail. This involves modifying the postulates from Blass and Gurevich [2003] to allow proclats to create (or

activate) additional procelets on the fly. In Section 4, we show that this modification solves the one remaining problem in Blass and Gurevich [2003], Section 8. Finally, in Section 5, we show how to extend the proof of the ASM thesis for parallel algorithms (Blass and Gurevich [2003], Theorem 10.1), to the wider class of algorithms defined by our modified postulates.

To avoid excessive repetitions, we assume that the reader is acquainted with the content of Blass and Gurevich [2003]. We occasionally give references to particular passages in Blass and Gurevich [2003], and we occasionally give reminders about particular points there, but we try to keep these references and reminders minimal.

Remark 1. An alternative response to the discovery of the flaw in Blass and Gurevich [2003] might be to retain the postulates and try to modify the definition of ASMs so as to capture exactly the algorithms defined by those postulates. Such a modification might be possible, but it seems to be of very limited interest. The postulates should describe the general notion of parallel computation. Since ASMs, as described in Blass and Gurevich [2003] and earlier in Gurevich [1995], are surely parallel algorithms, the postulates should cover them.

The need to liberalize the postulates may lead to the suspicion that further liberalization may be needed in the future and that, in contrast to the present situation, it could require an extension of the notion of ASM. In fact, the study of intrastep interaction of an algorithm with its environment has led to an extension of the ASM framework; see Blass et al. [2006]. But we do not know and couldn't conceive a parallel algorithm that works in wide, shallow steps without intrastep interaction with its environment but does not fit the framework of our postulates. If such algorithms exist, we would be very interested to hear about them.

One can also suspect that the postulates were composed to prove the ASM thesis without reforming the notion of parallel ASM of Gurevich [1995]. All we can do in this respect is to assure the reader that our goal is to understand the nature of algorithms, in this case wide-shallow-step algorithms without intrastep interaction, and, in the case of Blass et al. [2006], intrastep interactive, small-step algorithms. We are ready to go wherever the analysis takes us; let the chips fall where they may. In the case of parallel algorithms, our proof of the ASM thesis required only insignificant changes of the ASM model of Gurevich [1995]. In the case of Blass et al. [2006], the ASM model needed a nontrivial extension.

2. THE PROBLEMS AND HOW TO SOLVE THEM

The flaws that have been found in the work of Blass and Gurevich [2003] concern two of the examples in Section 8 of that article. In this section, we explain what these flaws are, we show how to easily correct all but one of them within the framework of Blass and Gurevich [2003], and we indicate why the one remaining flaw requires a liberalization of the framework. Specifically, we first remove the one error from Section 8.4 of Blass and Gurevich [2003] and then turn our attention to Section 8.5. The latter section contains two errors. One requires

some adjustment of the algorithm proposed in Blass and Gurevich [2003], but the other is more serious, requiring the liberalization of the framework to allow intrastep creation of procllets.

Formally, all of these problems in Section 8 of Blass and Gurevich [2003] had the same underlying cause: the use of comprehension terms in the description of algorithms. Recall that comprehension terms, expressions of the form $\{\{t(x) : x \in r : \varphi(x)\}\}$ denoting multisets, are part of the syntax of parallel ASMs, as introduced in Section 9 of Blass and Gurevich [2003]. They are not, however, terms in the ordinary sense of first-order logic, and that is the sense in which *term* was used in the postulates of Blass and Gurevich [2003], Section 7. The errors in the examples in Sections 8.4 and 8.5 arose from treating comprehension terms as though they were terms in the sense of the postulates.

Inspection of the postulates reveals that the notion of *term* was used there once explicitly and once implicitly. The explicit use was in the last clause of the Background Postulate (Blass and Gurevich [2003], p. 600), where `Procllet` was required to be a variable-free term. The implicit use was in the Procllets Postulate (Blass and Gurevich [2003], pp. 605–606), which required the procllets to execute a sequential algorithm with output. The definition of sequential algorithms with output (Blass and Gurevich [2003], p. 582) incorporated the Bounded Exploration Postulate from Gurevich [2000] (*Bounded Sequentiality* at this point in Blass and Gurevich [2003] was a typo), which required a bounded exploration witness consisting of finitely many terms. These sequential algorithms with output are equivalent to sequential ASMs with `Output` rules, as explained on p. 631 of Blass and Gurevich [2003]. These ASMs use ordinary terms (of the vocabulary described in the Procllets Postulate), not comprehension terms, in their updates, guards, and outputs, because the bounded exploration witness consists of ordinary terms. With these observations in mind, it is easy to see what was wrong in the examples of Sections 8.4 and 8.5 in Blass and Gurevich [2003].

We begin with Section 8.4 of Blass and Gurevich [2003], which dealt with first-order and fixed-point logic. The treatment of first-order logic was correct, but the explanation (on page 621) of a step of the induction leading to a fixed point used a comprehension term in the update of `temp(p)`. The context there (page 621) was that a procllet p is to calculate an inflationary fixed point $\text{IFP}_{P, \bar{x}}(\varphi(P, \bar{x}))(\bar{t})$ by calculating the successive iterates

$$\Phi^0 = \emptyset \quad \text{and} \quad \Phi^{n+1} = \Phi^n \cup \{\bar{a} : \varphi(\Phi^n, \bar{a})\}$$

at successive steps of the algorithm under construction. Other procllets are available to carry out the subsidiary calculations associated with φ and its subformulas. The step from Φ^n to Φ^{n+1} , as described in Blass and Gurevich [2003] involved a comprehension term, in which the procllet p converts its mailbox (the multiset of results from the subsidiary computations using Φ^n) into the desired Φ^{n+1} .

Fortunately, the comprehension term here becomes unnecessary if we modify slightly the format in which p stores (as `temp(p)` in the notation of Blass and Gurevich [2003]) the sets Φ^n and the activity of the subsidiary procllets that compute the value of φ .

Instead of taking $\text{temp}(p)$ to be the set Φ^n itself, let $\text{temp}(p)$ be the characteristic function of this set, regarded as a set of ordered pairs, that is,

$$\{(\bar{a}, \text{true}) : \bar{a} \in \Phi^n\} \cup \{(\bar{a}, \text{false}) : \bar{a} \in M^k - \Phi^n\},$$

where M is the base set of the structure in which the formula is to be evaluated and where k is the arity of P .

Also, when a subsidiary proklet has computed the truth value v of φ at certain arguments \bar{a} , using Φ^n as the interpretation of P , let it push to p the pair $\langle \bar{a}, v' \rangle$ where v' is the disjunction of v and the truth value of $\Phi^n(\bar{a})$.

Note that this v' is the truth value of $\bar{a} \in \Phi^{n+1}$, and so the ordered pair $\langle \bar{a}, v' \rangle$ is one of the pairs that p should have in $\text{temp}(p)$ at the next step of the computation. Thus, the update to be performed by p is simply to give $\text{temp}(p)$ the value of myMail (provided this is different from the previous value of $\text{temp}(p)$). No comprehension term is needed.

Technically, we should mention another modification in the work of the subsidiary procleets for φ and its subformulas. These procleets use the current value of Φ^n , which they pull from p . With the new format for storing Φ^n in $\text{temp}(p)$, the subsidiary procleets will, of course, have to take this format into account in their computations. This modification has no effect on the discussion in Blass and Gurevich [2003], because that discussion didn't include such fine details about the work of the subsidiary procleets.

We turn next to the first difficulty in Section 8.5, namely, in the work of a term-proklet $\langle \hat{u}, \bar{a} \rangle$ where u is a comprehension term $\{\{t(x) : x \in r : \varphi(x)\}\}$ and where \bar{a} is a list of values for the pseudofree variables of u . (In Blass and Gurevich [2003] we wrote t for what we now call u ; the new notation avoids confusion with $t(x)$.) This term-proklet $\langle \hat{u}, \bar{a} \rangle$ has the task of computing the value of u when its pseudofree variables are given the values \bar{a} , and Blass and Gurevich [2003] contained a three-part recipe for how to do this. The first two parts of the recipe are correct: activate the proklet $\langle \hat{r}, \bar{a} \rangle$, which will provide a value for r , say b . Then activate, for each $c \in b$, the procleets $\langle \hat{t}(x), \bar{a} \hat{\sim} c \rangle$ and $\langle \hat{\varphi}(x), \bar{a} \hat{\sim} c \rangle$, which will provide the values of $t(c)$ and $\varphi(c)$. The error is in the final step, where the right $t(c)$'s, namely, those corresponding to c 's for which $\varphi(c) = \text{true}$, are picked out, assembled into a set, and equipped with the correct multiplicities. As it stands, this requires the use of comprehension terms, which are unavailable to procleets.

The key idea for repairing the problem is to arrange for the proklet $\langle \hat{u}, \bar{a} \rangle$ to find the multiset it needs, the value of u , as its mailbox, rather than trying to assemble it. Carrying out this idea will require careful attention to what messages are sent to $\langle \hat{u}, \bar{a} \rangle$ by other procleets. The messages must be just the right $t(c)$'s, each with its right multiplicity. Arranging this will require some additional procleets and static functions, to do some of the work that was previously hidden in a comprehension term, for example, the work of picking out the right $t(c)$'s. We postpone the implementation of this idea until Section 4, because the details depend on the repair of the other error in Section 8.5 of Blass and Gurevich [2003], which we discuss next.

The remaining, serious error is the definition of the "terms" $\text{MDA}(p)$ and $\text{MA}(p)$ on pages 624–625 of Blass and Gurevich [2003]. (There is also an obvious

typographical error in the definition of MDA which, as it stands, never mentions MDA; the intent was to have $\text{MDA}(\langle \hat{A}, \bar{a} \rangle)$ in all six places where $\langle \hat{A}, \bar{a} \rangle$ occurs.) With this definition, $\text{MDA}(p)$ and $\text{MA}(p)$ are not really terms, in the sense required by the postulates for algorithms, because they involve comprehension terms. The intention behind the definitions of MDA and MA was, as explained in Blass and Gurevich [2003], to provide a finite set that contains all proclefs that might be activated by the algorithm. There seems to be no way to achieve this intention without using comprehension terms. That is, terms in the correct, first-order sense do not enable us to name, in advance of executing a particular step of the algorithm, a set guaranteed to contain all proclefs that might be needed during that step. In other words, an appropriate set of proclefs, for a particular step in the computation, can be described only during the execution of this step, not before the step begins. This is the motivation for the following extension of the notion of algorithm.

Do not require the full set of proclefs for any step to be given by a term `Proclet` of the algorithm’s vocabulary. Instead, require some subset, called *primary proclefs*, to be given by such a term, `PriProclet`. Allow primary proclefs to activate (or create) new *secondary proclefs* during the step. Furthermore, allow secondary proclefs to activate further proclefs (which we still call *secondary*, not *tertiary*), etc. All the proclefs, primary as well as secondary, participate in the activities of proclefs described in Blass and Gurevich [2003]—pushing and receiving messages, setting up and pulling displays, and updating the state. But only the primary proclefs are specified as part of the algorithm’s state. The secondary ones are temporary, losing their proclef status at the end of the step. In the terminology of Blass and Gurevich [2003], Section 7, the notion of secondary proclef is given by the `ken`, not the state.

In somewhat more detail, our conventions for the activation of secondary proclefs are as follows. To activate new proclefs, a proclef p , primary or secondary, must mark for activation a finite subset s of the state. The secondary proclefs thereby activated are not, however, the members of s but rather the ordered pairs $\langle q, p \rangle$ where $q \in s$. Thus, a secondary proclef $\langle q, p \rangle$ “knows” its creator p in the sense that it can refer to p by a term in its local state, namely, the term `second(me)`.

This liberalization of the notion of algorithm requires changes in several of the postulates and definitions of Blass and Gurevich [2003]; the next section spells these changes out.

3. POSTULATES

The purpose of this section is to modify the definition of algorithms from Blass and Gurevich [2003] by allowing proclefs to activate other proclefs.

Remark 2. The terminology *activate a proclef* was already used in Blass and Gurevich [2003], but with a different meaning than here.

In Blass and Gurevich [2003], each state of an algorithm determined a set of proclefs. The notion of “active” proclef was not part of the state but rather a convenient and intuitive way for us to informally describe certain aspects of a `ken`, and it could have different meanings in the context of different algorithms.

In particular, an inactive procelet was still a procelet and could, despite the terminology, engage in certain basic activities, for example pulling information to see whether it should become active.

In the present article, each state of an algorithm will determine a set of *primary* procleets. Activation produces additional, *secondary* procleets, which would not be procleets and could not engage in any activities at all if they were not activated. Activation will be a formal notion, an essential ingredient in our postulates, not merely a convenient abbreviation for certain aspects of kens.

The distinction between the two notions of activation may be clarified by noting that the present notion (but not that in Blass and Gurevich [2003]) of activating a procelet could as well be called *proceletizing an element of the state*.

In broad terms, our modification of the notion of parallel algorithm is to allow any procelet p to mark for activation, during the course of a step, some finite subset $\text{Act}(p)$ of the state X . The effect of this activation is that the elements (q, p) for $q \in \text{Act}(p)$ become procleets in their own right.

For this to make sense, a few observations are in order. First, as in Blass and Gurevich [2003], we regard subsets of X as multisets of elements of X in which all the multiplicities are 1. The Background Postulate of Blass and Gurevich [2003] ensures (and will continue to ensure even after we modify it below) that finite multisets of elements of a state X are themselves elements of X . In particular, $\text{Act}(p)$ is an element of X .

Second, what does it mean for p to “mark” a set? It means to assign that set as the value of a certain dynamic, nullary function symbol `myAct` in its local state. (This `myAct` is not part of the global state of the whole algorithm but of the local state of the individual procelet.) This is exactly like the way procleets produced their displays, $\text{Display}(p)$, by setting a value for `myDisplay` in Blass and Gurevich [2003]. We adopt the conventions that the initial value of `myAct`, at the beginning of any step, is the empty set and that, if a procelet assigns to `myAct` a value in X that is not a set, then nothing is thereby activated.

Third, not all procleets can arise from this activation process; there must be some procleets already available at the beginning of a step to get the activation process started. The set of these primary procleets is to be given by the value of a term `PriProcelet` in the state. We could require that there is always just one primary procelet; if more are wanted, they could be activated by the primary one at the beginning of the step. This theoretical simplification, however, seems to bring no real benefit, and would impose a cost: As was shown in Sections 8.1–8.4 of Blass and Gurevich [2003], most known types of parallel algorithms do not require intrastep activation of procleets. If we insisted on starting each step with only a single procelet, we would be forced to include intrastep activation even in algorithms that otherwise have no need for it.

After this broad outline of how we intend to modify the notion of parallel algorithm from Blass and Gurevich [2003], we turn to the details of revising the postulates in Blass and Gurevich [2003], Section 7, to accommodate this picture of activation of procleets.

The Sequential Time Postulate, which was taken unchanged from Gurevich [2000] to Blass and Gurevich [2003], could remain unchanged again, but we take

this opportunity to incorporate a small improvement that was first pointed out in Blass and Gurevich [2006], namely, that the sets $\mathcal{S}(A)$ of states and $\mathcal{I}(A)$ of initial states must be nonempty.

—*Sequential Time Postulate.* An algorithm A is associated with a nonempty set $\mathcal{S}(A)$ of states, a nonempty subset $\mathcal{I}(A) \subseteq \mathcal{S}(A)$ of initial states, and a map $\tau_A : \mathcal{S}(A) \rightarrow \mathcal{S}(A)$ called the *one-step transformation*.

The Abstract State Postulate, also unchanged from Gurevich [2000] to Blass and Gurevich [2003], remains unchanged here.

—*Abstract State Postulate.* All states of A are first-order structures with the same finite vocabulary, which we call the *vocabulary of A* . τ_A does not change the base set of a state. Both $\mathcal{S}(A)$ and $\mathcal{I}(A)$ are closed under isomorphisms. Any isomorphism from a state X to a state Y is also an isomorphism from $\tau_A(X)$ to $\tau_A(Y)$.

In the Background Postulate of Blass and Gurevich [2003], the last item required a variable-free term `Procllet` naming, in each state, a finite set, also called `Procllet`. Modify this by changing `Procllet` to `PriProcllet`. This change, though only notational at the current point in the postulates, reflects the fact that only the set of primary procllets is given with the state; secondary procllets are activated during the computation steps. The notation `Procllet` will be used later to refer to the whole set of procllets, primary and secondary.

—*Background Postulate.* Each state contains the following:

- the elements `true` and `false`, the Boolean operations on them, `undef`, and the equality predicate;
- all ordered pairs of elements of the state, with a binary function symbol for pairing $\langle x, y \rangle$ and unary functions `first` and `second` for extracting the components of a pair;
- all finite multisets of elements of the state, with symbols for the empty multiset \emptyset , singletons $\{\{x\}\}$, binary sum $x \uplus y$, general sum $\biguplus x$, `TheUnique`, and `AsSet`, and
- a variable-free term `PriProcllet` naming a finite set, also called `PriProcllet`.

Remark 3. As in Blass and Gurevich [2003], the Background Postulate ensures the availability of (among other things) the operation `AsSet`, defined as taking any x to the multiset that has the same members as x has but with multiplicity only 1. In Blass and Gurevich [2003], we used this operation only when the argument is a multiset, so that it simply removes the multiplicities. In the present article, there will be an additional use of `AsSet`, one in which the argument need not be a multiset. If x is not a multiset, then it has no members, and so the definition gives $\text{AsSet}(x) = \emptyset$. It follows that, using `AsSet` and equality, we can test for sethood; x is a set if and only if $x = \text{AsSet}(x)$.

In Blass and Gurevich [2003], Definition 7.10, a *ken* K of a state X was defined to consist of X together with two functions `MailboxK` and `DisplayK`, such that each has the set of procllets as its domain and such that the values of

Mailbox_K are multisets. This definition needs significant changes to accommodate activation of new procllets.

The most obvious change is that the ken must include information about which procllets activate which other procllets. Activation takes place entirely within a step. It does not persist across a step boundary from one state to the next; any procllets that should remain procllets for the next step should be added to PriProcllet . Thus, activation information is not part of the state. It is temporary information to be changed and used by the algorithm within a step, just like the mailboxes and displays of Blass and Gurevich [2003]; such information resides in the ken. So a ken K should include a unary function Act_K . The intended meaning of $\text{Act}_K(p)$ is the set marked for activation, as myAct , by the procllet p , but as far as the definition of *ken* is concerned, Act_K is just some function from procllets to sets in X . The intended meaning will be formalized later in the notion of *correct ken*. In these respects, Act_K and myAct behave just like Display_K and myDisplay in Blass and Gurevich [2003].

A more complicated change in the notion of ken arises from the fact that the functions Mailbox_K , Display_K , and Act_K have as their domain the set of procllets. In Blass and Gurevich [2003], that set was given with the state, but now it depends on the ken via the activations described by Act_K . This interdependence between kens and the associated sets of procllets accounts for the greater complexity of the following definition compared to the corresponding Definition 7.10 in Blass and Gurevich [2003].

Definition 4. A *ken* of a state X consists of X together with three functions, Mailbox_K , Display_K , and Act_K , with a common domain $\text{Procllet}_K \subseteq X$ and with values in X , subject to the following requirements.

- The values of Mailbox_K are multisets.
- The values of Act_K are sets.
- $\text{PriProcllet}_X \subseteq \text{Procllet}_K$.
- If $p \in \text{Procllet}_K$ and $q \in \text{Act}_K(p)$, then $\langle q, p \rangle \in \text{Procllet}_K$.
- Procllet_K is the smallest set satisfying the preceding two requirements.

The elements of Procllet_K are called the *procllets* of the ken K .

As usual in recursive definitions, the detailed meaning of the last requirement is that, for all subsets Z of the state X , if $\text{PriProcllet}_X \subseteq Z \subseteq \text{Procllet}_K$ and if, for each $p \in Z$ and each $q \in \text{Act}_K(p)$, we have $\langle q, p \rangle \in Z$, then $Z = \text{Procllet}_K$.

Notice that the closure condition on Procllet_K incorporates our convention that, if q is in the set $\text{Act}(p)$ then it is not q itself but the pair $\langle q, p \rangle$ that becomes a procllet.

Remark 5. This convention allows the procllets $\langle q, p \rangle$ activated by p to “know” which procllet activated them, that is, they can refer to p in their local states (see Definition 7 below) by means of the term $\text{second}(\text{me})$. This sort of knowledge seems intuitively reasonable, and it also serves two technical purposes.

First, it ensures that two proplets will not both activate the same secondary proplet for different purposes.

Second, it provides a way for a proplet p to pass information to the proplets it activates, namely by displaying it. Sending messages would not suffice for this purpose, since the number of messages sent by a proplet during a step will be bounded (because the proplets are sequential algorithms with output) while the number of proplets activated by p need not be bounded. So displaying and pulling are the only ways for p to convey information to all the proplets it activates. For this to work, all these proplets must know p , in order to pull the information.

If comprehension terms were available, then this second purpose would not require our convention of automatically indicating the activator in every secondary proplet. Indeed, if such indications were wanted, then instead of marking a set s for activation, the proplet p could mark $\{\langle q, p \rangle : q \in s : \text{true}\}$, that is, p could attach the activator tags on its own. Indeed, p could similarly convey any bounded amount of additional information to the proplets it creates, by building this information into the proplets themselves. But, since comprehension terms are not available to the proplets, this approach will not work. And in any case, it would not achieve the first purpose indicated above.

Remark 6. One can imagine a more powerful sort of activation, where the tag added to a secondary proplet is not necessarily its activator but some other element of the state chosen by the activator. In addition to `myAct`, there would be another dynamic, nullary function `myTag` (initially `undef`). When a proplet updates `myAct` to a set s and `myTag` to e , the effect is to activate secondary proplets $\langle q, e \rangle$ for $q \in s$.

In this system, it would be possible for several proplets to activate the same secondary proplet, so it would be up to the algorithm to prevent unwanted clashes—for example by always including the activator p as a component of e .

This system provides a powerful means of communication from a proplet p to the proplets q that it activates. It avoids the need for p to display information for these q 's since it can build the information into the proplets themselves.

Indeed, this sort of tagging could, with some awkwardness, replace displaying and pulling as a means of communication. One strategy for doing this is as follows. Let all the proplets x that want to pull from p instead send p a message, say of the form $\langle x, \text{pull} \rangle$. Instead of displaying an entity d , p activates new auxiliary proplets $\langle q, \langle d, p \rangle \rangle$ for all q in its mailbox. That is, it updates `myAct` to `myMail` and updates `myTag` to $\langle d, \text{me} \rangle$. Each new proplet $\langle q, \langle d, p \rangle \rangle$ checks whether its first component q is of the form $\langle x, \text{pull} \rangle$. If so, it sends its second component $\langle d, p \rangle$ to x , which interprets this message as meaning that d is the display of p . A careful presentation of this strategy would have to prevent conflicts between the messages used here, to simulate displays, and any messages that the proplets use for other purposes; we refrain from looking at these details.

In Blass and Gurevich [2003], Definition 7.11, we defined the local state of a proplet, given a ken K for a state X . To accommodate activation, we expand the local states of Blass and Gurevich [2003] to include a dynamic nullary symbol `myAct` whereby a proplet indicates the set it wants to activate. The new

definition, replacing Definition 7.11 of Blass and Gurevich [2003], therefore reads as follows.

Definition 7. Suppose K is a ken of a state X and suppose $p \in \text{Proclet}_K$. An *initial local state* for p in X is the structure X plus

- a static, nullary symbol me , interpreted as p ;
- a static, nullary symbol myMail , interpreted as some multiset in X ;
- a static, unary symbol Display , interpreted as some unary function $X \rightarrow X$;
- a dynamic, nullary symbol myDisplay , interpreted as undef ; and
- a dynamic, nullary symbol myAct , interpreted as \emptyset .

The *initial local state of p given by K* is the initial local state for p in X where

- myMail is interpreted as $\text{Mailbox}_K(p)$ and
- Display is interpreted as Display_K extended¹ by $\text{Display}(x) = \text{undef}$ for $x \notin \text{Proclet}_K$.

Remark 8. The initial local state of p is the only local state of p that we shall work with. In any step of the overall algorithm, a procelet will execute the procelet algorithm only once, in the initial state. The result of this execution can include, in addition to output (messages to other procleets) and updates to the global state (to be executed at the end of the overall algorithm's step), updates to the dynamic symbols myDisplay and myAct . These updates can be regarded as producing a new non-initial local state, but no computation will be done in that state.

We can therefore, when discussing the state in which a procelet computes, speak of its local state, omitting the word *initial*. In fact, in Blass and Gurevich [2003], we didn't even introduce *initial* in this context. We have done so here in order to emphasize that these states use the initial default values for myDisplay and myAct , even in the case of a procelet p and a ken K for which $\text{Display}_K(p)$ and $\text{Act}_K(p)$ have values different from these defaults. The similarity between the names Display and myDisplay does not entail any connection between the value of the former (in some ken) and the initial value of the latter. Only in the case of correct kens (defined later) is the syntactic similarity reflected in a semantical connection, and that connection does not involve the initial value of myDisplay but the final value, after execution of the procelet algorithm (also defined below). The same comments apply to Act and myAct .

As in Blass and Gurevich [2003], we occasionally refer to the states of the entire algorithm as *global* states, to distinguish them from the local states of procleets.

Remark 9. There is an analogy between myAct and myDisplay . Both provide ways for a procelet to make a contribution to the overall ken, specifically to Act_K and Display_K , respectively, at least when we deal (as we soon will) with correct kens. But there are two differences. The lesser of the two is that Act_K is required

¹This extension should also have been in the corresponding Definition 7.11 of Blass and Gurevich [2003].

to take sets as values, whereas the dynamic symbol `myAct` could, in principle, denote any element of the state. We shall (in the definition of *correct ken*) adopt the convention that if a procelet p gives `myAct` a value that is not a set, then the resulting value of $\text{Act}_K(p)$ should be \emptyset .

The second difference is that Display_K is part of the initial local state given by K but Act_K is not. That is, a procelet has access to what other procleets have displayed but not to the activations performed by other procleets. Why?

Allowing procleets to know what other procleets activate would introduce a sneaky means of communication. A procelet p could activate another procelet q , not so that q would participate in the computation but rather so that other procleets, seeing that q is activated, would be able to infer some information that p wants to transfer to them. If we are not careful, such communication could make the computation process circular, that is, there might be no correct ken. For example, p might activate q if and only if it has no incoming message, while r might send a message to p if and only if q is activated. If there are no other procleets that might send a message to p , then no correct ken is possible; the message-sending and activating specified by the procleets' programs are circular and contradictory.

This particular example could be rendered harmless by acknowledging that p is sending information to r and putting an edge from p to r in the information flow digraph. That edge, together with the one arising from the message r might send to p , would constitute a cycle in that digraph, contrary to the Bounded Sequentiality Postulate. So the example would be excluded by this postulate.

But the general problem cannot be removed so easily. If procleets could know, in general, about each other's activations, then the information flow digraph should have edges from every procelet to every procelet—a monstrous contradiction to Bounded Sequentiality. To avoid such a disaster, we would have to perform an analysis of just which procleets really do (in some ken) get information from which other procleets. Such an analysis would be essentially the same as the analysis leading to the definition of *pulls from*, Definition 7.18 in Blass and Gurevich [2003].

There is a simpler way to get activation information to those procleets that might need it: if other procleets should find out what p has activated, let p incorporate $\text{Act}(p)$ into its display for the other procleets to read. In this way, the analysis mentioned above is subsumed by the analysis leading to *pulls from*, there is no need to include sneaky transmission of information in the information flow digraph, and all communication between procleets still fits the push and pull paradigms of Blass and Gurevich [2003].

The Procleets Postulate of Blass and Gurevich [2003] needs one evident addition and some reorganization. The addition is that the initial local state in which a procelet operates should contain the dynamic, nullary symbol `myAct`. The reorganization arises from the following considerations along with a desire to stay conceptually close to the picture in Blass and Gurevich [2003]. The Procleets Postulate of (Blass and Gurevich [2003], Section 7.3) was written so as to refer only to states, not to kens. Thus, for example, it says that `myMail` should denote (in the local state of a procelet p) some multiset, not that this multiset

should be $\text{Mail}_K(p)$ for a specific ken K . Kens appeared in the postulates of Blass and Gurevich [2003] only in the subsequent Section 7.4, which dealt with interaction between proplets. In our present context, we must give up either this ken-independence of the Proplets Postulate or the mention of proplets in the postulate. The reason is, of course, that the notion of proplet now depends on the ken and not merely on the state. We choose the second option. That is, we retain the general structure of the Proplets Postulate, referring only to the state, and we therefore postpone any mention of the proplets themselves in postulates.

Because the postulate is not about proplets but only about their algorithm, we rename it accordingly.

—*Proplet Algorithm Postulate.* The algorithm A determines a single sequential algorithm with output, called the *proplet algorithm*, in the vocabulary of the global algorithm plus the static nullary symbols me and myMail , the static unary symbol Display , and the dynamic nullary symbols myDisplay and myAct . The outputs of the proplet algorithm are ordered pairs (addressee, content) of elements of the state.

During any step of the (overall) algorithm, each element p in the current state X is to be regarded as potentially executing the proplet algorithm for one step in a state consisting of X , p as the denotation of me , some multiset as the denotation of myMail , some unary function as the denotation of Display , undef as the initial value of myDisplay , and \emptyset as the initial value of myAct . *Potentially* here refers to the fact that, once we define the correct ken K , only the elements of Proplet_K will actually run the proplet algorithm. Also, once we define the correct ken K , the initial local state in which a proplet p operates will be the initial local state given by K as in Definition 7.

We emphasize that, in each step of the global algorithm, the proplet algorithm is to be executed for only one step by each proplet. Accordingly, it will be very convenient to use the following abbreviation.

Definition 10. Let K be a ken of a state X , and let $p \in \text{Proplet}_K$. We abbreviate *one step of the proplet algorithm is executed in the initial local state of p given by K as p fires in K* . Here *one step is executed* means that the transition function is applied and outputs are produced once; there is no iteration as in runs of an algorithm.

The remarks following the Proplet Algorithm Postulate in Blass and Gurevich [2003], Section 7.3, apply here with the obvious additions of saying that myAct behaves like myDisplay , being updated during (rather than at the end of) a step, and not retaining its value past the end of a step.

Section 7.4 of Blass and Gurevich [2003], about the interaction between proplets, needs substantial modification, especially because the very notion of proplet now depends on a sort of interaction, namely, activation.

The first modification here, in Definition 7.18 from Blass and Gurevich [2003] of “ q pulls from p ,” is rather minor. The idea behind the definition was that q pulls from p in state X if there are two kens for X , differing only in the values

of $\text{Display}(p)$, such that q behaves differently in its initial local states given by these kens. We do not change this idea, but the detailed meaning of *behaves differently*, namely, producing different updates or sending different mailings, must now be extended to include activating different elements. Formally, the required change in Definition 7.18 of Blass and Gurevich [2003] is that “different updates (in the global state or in myDisplay)” becomes “different updates (in the global state, in myDisplay , or in myAct).”

In addition, we explicitly require the two kens in the definition to agree that p and q are proclets. In fact, this was already implicit. The definition involves $\text{Display}(p)$, which is defined only when p is a procllet. It also involves firing q , that is, executing q in appropriate initial local states; these states involve $\text{Mailbox}(q)$, which is defined only if q is a procllet.

Definition 11. Let p and q be elements of a state X . Then q pulls from p if there are two kens K and K' such that

- both p and q are in $\text{Procllet}_K \cap \text{Procllet}_{K'}$;
- K and K' differ only in the values of $\text{Display}(p)$, and
- when q fires in K and in K' , the results differ either in updates (of the global state or myDisplay or myAct) or in mailings (counting multiplicities).

The information flow digraph of Blass and Gurevich [2003], Definition 7.19, should be modified so that if p activates q then there is an edge from p to q . Intuitively, something has flowed from p to q , perhaps “activeness” or “procllet-hood”; it is not obvious whether this constitutes information. But the following example shows that, whether or not we call it information, it must be included in the digraph.

Example 12. Suppose there is just one primary procllet p , which activates another procllet q if and only if p 's mailbox is empty (and p does nothing else). Suppose further that q 's computation is just to send a message to p . So the information flow digraph has an edge from q to p , representing the possibility of a mailing, and we claim that it should also have an edge from p to q , representing the possibility of activation. By including this second edge, we introduce a cycle into the digraph, so that the Bounded Sequentiality Postulate is violated and the instructions we gave for these two procllets do not constitute an algorithm. Without the second edge, the instructions would satisfy the postulate. Our claim is that the former outcome, “not an algorithm,” is correct. The reason is that these instructions cannot be consistently executed. p must activate q if and only if p 's mailbox is empty, which happens if and only if q is not activated. (Once we define correctness of kens, we can say that the circularity prevents the existence of a correct ken.)

Taking into account the preceding observations, the dependence of the notion of procllet on the ken, and our desire to remain as close to Blass and Gurevich [2003] as these considerations permit, we are led to the following definition of the information flow digraph, replacing Blass and Gurevich [2003], Definition 7.19.

Definition 13. Let X be a (global) state. Its *information flow digraph* has as vertices all those elements $p \in X$ that belong to Proclet_K for at least one ken K . There is an edge from p to q if at least one of the following conditions is satisfied.

- q pulls from p .
- There is a ken K , for which both p and q are procllets, such that, when p fires in K , it sends a message to q .
- q is of the form $\langle r, p \rangle$ and there is a ken K , for which p is a procllet, such that, when p fires in K , it updates myAct to a set containing r .

Remark 14. The three clauses defining the edges of the information flow digraph correspond to information flow by pulling, pushing, and activating, respectively. In all three clauses, as well as in the definition of vertices, we have included everything that *might* be involved for *some* reasonable ken.

Remark 15. We explain briefly why our formulation of the third clause in Definition 13, the one about activation, is preferable to two plausible-looking alternatives. The first alternative is to replace the part about p updating myAct to a set containing r with the simpler requirement $r \in \text{Act}_K(p)$. The second is to require not only p but also q to be a procllet of K . (Our definition requires q to be a procllet of some ken, in order to be a vertex of our digraph, but it need not be a procllet of the same ken K mentioned in the third clause.) Readers for whom our definition is obviously better than these alternatives are invited to skip the rest of this remark.

The defect in both alternatives arises from the fact that, in arbitrary kens K (as opposed to the correct kens, which will be defined later), no connection is required between the function Act_K and the results of the procllets' computations of myAct in their local states given by K .

On the one hand, this means that we could have $r \in \text{Act}_K(p)$ even when r is not at all the sort of thing that p might activate under the procllet algorithm, indeed, even if the procllet algorithm is such that myAct can never be updated. As a result, the first proposed alternative could put into our digraph a great many activation edges that have nothing to do with any activation that the algorithm would ever perform. Indeed, we would have edges from each vertex p to all vertices of the form $\langle r, p \rangle$, for arbitrary r in the state. In this situation, the Bounded Sequentiality Axiom would be excessively difficult to satisfy.

On the other hand, if a procllet p , firing in K , activates $\langle r, p \rangle$ by updating myAct to a set that contains r , there is no guarantee that $r \in \text{Act}_K(p)$, and therefore there is no guarantee that $\langle r, p \rangle$ is a procllet of K . As a result, the second proposed alternative would miss many possible activations.

Both alternative formulations would be admissible if we were working only with correct kens, but we are not, and for good reason. The existence and uniqueness of correct kens (Theorem 23 below) depends on the Bounded Sequentiality Postulate, which in turn is formulated in terms of the information flow digraph. For more information about the need to work with all kens rather than only correct ones, see Sections 7.4 and 12 of Blass and Gurevich [2003].

With the revised definition of the information flow digraph, we can retain the Bounded Sequentiality Postulate from Blass and Gurevich [2003].

—*Bounded Sequentiality Postulate.* There is a uniform bound B , depending only on the algorithm and not on the state, for the lengths of all directed walks in the information flow digraphs of global states.

Definition 16. Let X be a state and p a vertex of its information flow digraph. The *level* of p in this digraph is the length of the longest walk in the digraph that ends at p . By *length* here we mean (as in Blass and Gurevich [2003]) the number of vertices in the walk, not the number of edges, so levels begin with 1, not 0.

Clearly, the level of p always exists and is no greater than the uniform bound B given by the Bounded Sequentiality Postulate. It also follows immediately from the definition that, if the information flow digraph has an edge from p to q , then the level of p is strictly smaller than that of q .

Our next task is to prove the analog, in our new situation, of Theorem 7.22 of Blass and Gurevich [2003], which asserts the existence of a unique correct ken for each state. We begin with a definition of correctness, just like that in Blass and Gurevich [2003] except that it takes activation into account, both explicitly in a clause relating `Act` to `myAct` and implicitly by using the new notion of ken.

Definition 17. Let X be a global state and K a ken for X . Then K is a *correct* ken for X if the following three conditions are satisfied.

- For each $p \in \text{Proclet}_K$, the members of $\text{Mailbox}_K(p)$ are the messages sent to p by the procllets $q \in \text{Proclet}_K$ when each such q fires in K . (The multiplicity of a message m in $\text{Mailbox}_K(p)$ is the sum, over all $q \in \text{Proclet}_K$, of the multiplicity with which q sends m to p , that is, the multiplicity of $\langle p, m \rangle$ in the output multiset of q 's execution of the procllet algorithm.)
- For each $p \in \text{Proclet}_K$, the value of $\text{Display}_K(p)$ is the value that `myDisplay` obtains when p fires in K .
- For each $p \in \text{Proclet}_K$, the value of $\text{Act}_K(p)$ is the value that `myAct` obtains when p fires in K , provided this value is a set. Otherwise, $\text{Act}_K(p) = \emptyset$.

For technical purposes, it is useful to have a name for the following consequence of correctness.

Definition 18. Let X be a global state and K a ken for X . Then K is a *plausible* ken for X if, whenever $p \in \text{Proclet}_K$ and $q \in \text{Act}_K(p)$, then $\langle q, p \rangle$ is at a higher level than p in the information flow digraph for X .

In connection with this definition, notice first that from $p \in \text{Proclet}_K$ and $q \in \text{Act}_K(p)$ it follows that $\langle q, p \rangle$ is a procllet of K and is therefore a vertex of the information flow digraph.

Notice also that, if K satisfies the third requirement in the definition of *correct*, then it is plausible. Indeed, if we assume the third requirement and also assume $p \in \text{Proclet}_K$ and $q \in \text{Act}_K(p)$, then p firing in K produces a value of `myAct` that contains q . That provides an edge in the information flow

digraph from p to $\langle q, p \rangle$ and thus ensures that $\langle q, p \rangle$ is at a higher level than p .

Before proving that every state has a unique correct ken, we record some preliminary information that will be needed in that proof.

Definition 19. Let K be a ken for a state X , and let l be a positive integer. $H(K, l)$ is defined to be the ken that is the same as K except that $\text{Display}_{H(K, l)}(p) = \text{undef}$ for those $p \in \text{Proclet}_K$ that have level $\geq l$ in the information flow digraph for X . We refer to $H(K, l)$ as the result of *hiding the displays* of K from level l up.

Notice that $\text{Act}_{H(K, l)} = \text{Act}_K$ and therefore $\text{Proclet}_{H(K, l)} = \text{Proclet}_K$. In particular, in passing from K to $H(K, l)$, we need not adjust the domains of the functions.

LEMMA 20. *Let K be a ken for a state X , p a proklet of K , and l its level in the information flow digraph of X . Consider the one-step executions of the proklet algorithm by p in two initial local states, the one given by K and the other by $H(K, l)$. These two executions produce the same updates (to the global state and to myDisplay and myAct) and the same output multiset of messages.*

PROOF. Inspection of the definitions of initial local states and hiding reveals that the two initial local states mentioned in the lemma differ only in the values of $\text{Display}(q)$ when $q \in \text{Proclet}_K = \text{Proclet}_{H(K, l)}$ and q has level $\geq l$. Of these q 's, only finitely many are relevant to the updates and outputs produced when p executes the proklet algorithm, because this algorithm, being a sequential algorithm with output, satisfies the Bounded Exploration Postulate. So we can connect K to $H(K, l)$ by a finite sequence of kens, in which

- at the first step we pass from K to a ken in which $\text{Display}(q)$ has been changed to undef for all the irrelevant q 's of level $\geq l$, so the new ken differs only finitely from $H(K, l)$, and
- at each subsequent step, we change the value of $\text{Display}(q)$ for only one $q \in \text{Proclet}_K$.

We already know that at the first step in this sequence the updates and outputs of p are unchanged. It remains to check that the same is true at all subsequent steps in our sequence of kens. So consider, for the rest of this proof, a particular pair of consecutive kens in the rest of our sequence; suppose, toward a contradiction, that they do not agree as to p 's computation; and let q be the unique element where their Display functions differ.

Because all the kens in the sequence have the same Act function, they all have the same prockets as K . In particular, our q is a proklet for both of the consecutive kens under consideration. This and the assumed disagreement as to p 's computation imply that p pulls from q and so there is an edge from q to p in the information flow digraph. This is absurd, as q has level $\geq l$ and p has level l . \square

LEMMA 21. *Let X be a state, K and K' two plausible kens for it, and l a positive integer. Assume that, for all $q \in \text{Proclet}_K \cap \text{Proclet}_{K'}$,*

- if q has level $\leq l$ then $\text{Mailbox}_K(q) = \text{Mailbox}_{K'}(q)$,
- if q has level $< l$ then $\text{Display}_K(q) = \text{Display}_{K'}(q)$, and
- if q has level $< l$ then $\text{Act}_K(q) = \text{Act}_{K'}(q)$.

Let $p \in \text{Proclet}_K$ have level l . Then p is also in $\text{Proclet}_{K'}$, and the executions of p in its initial local states given by K and by K' produce the same updates and the same outputs.

PROOF. We proceed by induction on l and observe that our assumptions about K and K' imply that the same assumptions also hold if l is replaced by any smaller l' . So, by induction hypothesis, any procelet of K with level $< l$ is also a procelet of K' . Since p is a procelet of K , it is either in PriProclet_X or activated by some other $q \in \text{Proclet}_K$ (i.e., $p = \langle r, q \rangle$ for some $r \in \text{Act}_K(q)$). In the former case, we immediately have that p is also a procelet of K' , because PriProclet is given by the state, not the ken. In the latter case, the assumption of plausibility implies that q has level $< l$. So $q \in \text{Proclet}_{K'}$ and, by hypothesis, $\text{Act}_K(q) = \text{Act}_{K'}(q)$. In particular, q activates p in K' , and so $p \in \text{Proclet}_{K'}$, as desired.

To prove that p produces the same updates and outputs whether it fires in K or in K' , we may, thanks to Lemma 20, work with the kens $H(K, l)$ and $H(K', l)$ obtained by hiding displays from level l up. But we claim that the initial local states of p given by these kens coincide. Clearly, proving the claim would suffice to complete the proof of the lemma.

To prove the claim, notice first that the only effect of a ken on the initial local state of p is to provide the interpretations of `myMail` and `Display`. Since p is at level l , the first assumption of the lemma ensures that K and K' and therefore also $H(K, l)$ and $H(K', l)$ agree as to `Mailbox`(p), which provides the value of `myMail` for p 's initial local state. As for `Display`, the second assumption of the lemma ensures that all of K , K' , $H(K, l)$, and $H(K', l)$ agree as to the displays of those elements of level $< l$ that are procleets with respect to both K and K' . But the induction hypothesis implies, as in the first paragraph of this proof, that an element of level $< l$ is a procelet for all or none of K , K' , $H(K, l)$, and $H(K', l)$. So, for q of level $< l$, `Display`(q) will have the same value for all of K , K' , $H(K, l)$, and $H(K', l)$, either by the second assumption of the lemma or because of the convention, in the definition of initial local states, that `Display` of a non-procelet is always `undef`. The same convention combined with the definition of hiding ensures that `Display`(q) has the same value, namely `undef`, for $H(K, l)$ and $H(K', l)$ whenever q has level $\geq l$. (Notice that such a q might be a procelet for just one of K and K' . In that case, `Display` _{$H(K, l)$} (q) and `Display` _{$H(K', l)$} (q) would be `undef` for different reasons—once because of hiding and once because of the convention concerning `Display` of non-procleets. The need to deal with this situation blocks the easier argument of Blass and Gurevich [2003], Lemma 7.25 and motivated the notion of hiding.) This completes the proof that the initial local states of p given by $H(K, l)$ and $H(K', l)$ are identical, as required. \square

Remark 22. The hypothesis of the lemma may seem a bit awkward, because it concerns `Mailbox` up to and including level l but `Display` and `Act` only strictly below level l . We claim, however, that this situation is semantically natural even

if syntactically awkward. The reason is that the parts of the ken covered by this hypothesis at stage l are exactly the parts relevant to the computations done by elements p at level l . Indeed, the information relevant to the computation of such a p consists of (1) whether p is a proclat and therefore should be doing a computation at all, (2) the mailbox of p , and (3) the displays that p reads. Here (1) involves `Proclat` at level l (namely, at p), which is determined by `Act` at strictly lower levels (because of plausibility); (2) involves `Mailbox` at level l ; and (3) involves `Display` at lower levels.

Another way to view the situation is that the hypothesis covers exactly the information that is determined, in a correct ken, by the activity of proclats at strictly lower levels than l . As a special case, for $l = 1$, we have just the information that is determined by the state without any need for computation by proclats, namely, the proclats at level 1 (determined by `PriProclat` as plausibility prevents any activation of proclats of level 1) and their mailboxes (empty as nothing can send messages to them).

The general structure of this hypothesis, referring to `Mailbox` (and implicitly to `Proclat`) for one level higher than `Display` and `Act`, will recur in subsequent arguments.

THEOREM 23. *For every state X , there is a unique correct ken.*

PROOF. We first prove uniqueness. Suppose both K and K' are correct kens for the state X . We prove the following by induction on l .

- (1) `ProclatK` and `ProclatK'` have the same elements of levels $\leq l$. So for levels $\leq l$ we can speak of proclats without specifying which of the two kens we mean.
- (2) `MailboxK` and `MailboxK'` agree on all proclats of level $\leq l$.
- (3) `DisplayK` and `DisplayK'` agree on all proclats of level $< l$.
- (4) `ActK` and `ActK'` agree on all proclats of level $< l$.

For the base case, $l = 1$, items (3) and (4) are vacuous. Item (1) follows from the fact that a proclat at level 1 cannot be activated by another proclat in a correct (or just plausible) ken. So at level 1 the proclats for any plausible ken are just those in `PriProclatX`. Item (2) at level 1 follows from the fact that messages always go from lower to higher levels (because they produce edges) in the information flow digraph. In particular, no message ever goes to a proclat of level 1. Since K and K' are correct, it follows that `MailboxK` = `MailboxK'` = \emptyset .

For the induction step, assume the assertions hold for a certain l . As noted in item (1), it makes sense to speak of proclats at level l without specifying K or K' . For each such proclat p , we can apply Lemma 21 to conclude that its updates and outputs are the same whether K or K' provides its initial local state. By correctness, it follows that `DisplayK(p)` = `DisplayK'(p)` and `ActK(p)` = `ActK'(p)`. Thus, we have items (3) and (4) for $l + 1$ in place of l .

Furthermore, the mail sent by p is the same for K as for K' . Since this holds for all proclats p at level l and also for all proclats at lower levels (by the same argument), and since proclats at level $l + 1$ can receive mail only

from proclats at levels $\leq l$ (because message-sending produces edges in the information flow digraph), we conclude that every proclat at level $l + 1$ receives the same messages, with the same multiplicity, whether the proclats are firing K or in K' . Another application of correctness gives us item (2) for proclats at level $l + 1$. Since we already had this item for proclats of lower level, we now have item (2) with $l + 1$ in place of l .

Finally, since proclats can be activated only by proclats of lower level, and since PriProclat depends only on the state, not the ken, the proclats at level $l + 1$ are determined by $\text{Act}(q)$ for proclats of levels $< l + 1$. Thus, item (1) for $l + 1$ follows from item (4) for $l + 1$.

This completes the induction and thus, if we take l to be the number of levels, the proof of uniqueness. To prove existence, we construct the desired K by induction on levels l . After stage l , we will have constructed a ken $K(l)$, intended to have the following properties. Its proclats are the proclats of the correct ken up to and including level l , and its Mailbox function agrees with that of the correct ken on these proclats. Furthermore, its Display and Act functions agree with those of the correct ken on proclats at levels $< l$. On proclats of level $\geq l$, its Display and Act functions have the default values undef and \emptyset , respectively.

For $l = 1$, let $\text{Mailbox}_{K(1)}$, $\text{Display}_{K(1)}$, and $\text{Act}_{K(1)}$ be the constant functions with domain PriProclat_X and with values \emptyset , undef , and \emptyset , respectively. This is clearly a plausible ken, with $\text{Proclat}_{K(1)} = \text{PriProclat}_X$.

For the induction step, suppose $K(l)$ is already defined. Let each proclat $p \in \text{Proclat}_{K(l)}$ of level $\leq l$ fire in $K(l)$. Temporarily define $\text{Display}_{K(l+1)}(p)$ to be the resulting value of myDisplay for each such p and to be undef for all other elements p of the state. Similarly, temporarily define $\text{Act}_{K(l+1)}(p)$ to be the resulting value of myAct for each such p where this value is a set, and to be \emptyset for all other elements p of the state. (We have defined these functions on too large a domain and will correct for this later; this was the only reason for saying “temporarily.”) Use this temporary $\text{Act}_{K(l+1)}$ to define $\text{Proclat}_{K(l+1)}$, as in the definition of ken. Then restrict $\text{Display}_{K(l+1)}$ and $\text{Act}_{K(l+1)}$ to $\text{Proclat}_{K(l+1)}$ as required in the definition of ken. Note that, by restricting $\text{Act}_{K(l+1)}$ we have not changed what $\text{Proclat}_{K(l+1)}$ should be, since the characterization of Proclat in terms of Act (in the last three clauses in the definition of ken) uses Act only applied to elements of Proclat .

To complete the induction step, we must define $\text{Mailbox}_{K(l+1)}(p)$ for all elements $p \in \text{Proclat}_{K(l+1)}$. Define it to be the multiset of messages sent to p by the elements q of $\text{Proclat}_{K(l)}$ firing as in the preceding paragraph, multiplicities being summed over q . This completes the definition of $K(l)$.

In the inductive step, if $q \in \text{Act}_{K(l+1)}(p)$, with p and therefore also q in $\text{Proclat}_{K(l+1)}$, then, by definition of $\text{Act}_{K(l+1)}$, the ken $K(l)$ witnesses that the information flow digraph has an edge from p to $\langle q, p \rangle$. Thus, $K(l+1)$ is plausible. Since $K(1)$ is vacuously plausible, all $\text{Act}_{K(1)}(p)$ being empty, we conclude that $K(l)$ is plausible for all l .

We shall show that the sequence of kens $K(l)$ gradually stabilizes in the following sense. For each l , the kens $K(l)$ and $K(l + 1)$ agree as to Proclat and Mailbox up to and including level l of the information flow digraph, and

they agree as to Display and Act at all levels strictly below l . The proof is by induction on l .

For the basis of the induction, $l = 1$, notice that the assertions of agreement concerning Display and Act are vacuous. Concerning Proclet, the definition says that it contains only elements of PriProclet when the ken is $K(1)$, but when the ken is $K(2)$ it can contain also secondary $\langle r, p \rangle$, where $r \in \text{Act}_{K(2)}(p)$. As $K(2)$ is plausible, such secondary proclefs cannot be at level 1. Thus, $K(1)$ and $K(2)$ have the same proclefs at level 1. These proclefs have, by definition, empty mailboxes under $K(1)$. Under $K(2)$, their mailboxes contain the messages sent to them by proclefs q firing in $K(1)$. Any such message would cause an edge from q to p in the information flow digraph, which is impossible since p is at level 1. This completes the verification of the claimed stabilization between $K(1)$ and $K(2)$ and thus the basis of our induction.

For the induction step, suppose we have the desired stabilization between $K(l)$ and $K(l + 1)$. To prove the stabilization between $K(l + 1)$ and $K(l + 2)$, compare the definitions of these two kens. The former fires the proclefs of $K(l)$ at levels $\leq l$ in their initial local states given by $K(l)$; the latter fires the proclefs of $K(l + 1)$ at levels $\leq l + 1$ in their initial local states given by $K(l + 1)$. By induction hypothesis, these two sets of firings involve the same proclefs p at levels $\leq l$ (though there may be different proclefs at higher levels). Furthermore, for each such procler p , its two computations yield the same updates and outputs, by Lemma 21 since the kens are plausible. This means, in particular, that, in the definitions of $K(l + 1)$ and $K(l + 2)$, the temporary versions of Display and Act agree up to and including level l .

We claim next that that Proclet is the same, for $K(l + 1)$ and $K(l + 2)$, up to and including level $l + 1$. Suppose, toward a contradiction, that some q of level $\leq l + 1$ is a procler in one of the kens $K(l + 1)$ and $K(l + 2)$ but not in the other. Consider such a q of minimum possible level in the information flow digraph. Clearly, q cannot be in PriProclet_X , for then it would be a procler of both kens. So it is of the form $q = \langle r, p \rangle$ where, for one of the kens p is a procler and $r \in \text{Act}(p)$, while for the other ken either p is not a procler or $r \notin \text{Act}(p)$. As the kens are plausible, p is of lower level than q , so, by our choice of q to minimize the level, p is a procler of both kens. So we must have $r \in \text{Act}(p)$ for one ken and not for the other. We already saw that the Act functions of these two kens agree up to and including level l , so p must have level $\geq l + 1$. That contradicts the fact that p has lower level than q . This contradiction completes the proof that the kens $K(l + 1)$ and $K(l + 2)$ have the same proclefs up to and including level $l + 1$.

Finally, for these common proclefs at levels $\leq l + 1$, the two kens $K(l + 1)$ and $K(l + 2)$ have the same mailboxes, since these mailboxes come from the outputs of the computations by proclefs of levels $\leq l$, in initial local states given by $K(l)$ and $K(l + 1)$, and we have seen that these outputs are the same in both cases.

This completes the inductive proof of the claimed agreement between $K(l)$ and $K(l + 1)$ up to level l . Apply this result with l greater than the maximum length B of walks allowed by the Bounded Sequentiality Postulate. For such an l , we have $K(l) = K(l + 1)$; the whole kens are identical. Rereading the

definition of $K(l + 1)$ in the light of this equality, we find that it says precisely that $K(l)$ is correct. \square

Now that we have the existence and uniqueness of the correct ken for any state, we can formulate the final postulate.

—*Update Postulate.* The update set of the algorithm in a (global) state is the set of all the updates of global dynamic functions produced by all the procllets of the correct ken, firing in the correct ken.

This is exactly like the corresponding postulate in Blass and Gurevich [2003] except that the notion of procllet now depends explicitly on the ken.

Remark 24. The fact that, in each step of an algorithm, each procllet fires just once is formally contained in the Update Postulate. This postulate refers to a single firing, and so do the definitions on which it depends, like the definition of correct kens. And this postulate describes the whole influence of the procllets on the overall computation, since only the updated state persists to the next step.

The intuitive notion of bounded sequentiality requires bounds not only on the number of procllets that act in sequence, as formalized in the Bounded Sequentiality Postulate, but also on the sequentiality in the actions of a single procllet. The latter bound is ensured, in our postulates, by the requirements that procllets execute a sequential algorithm and that they fire only once per step.

Remark 25. The Update Postulate implies that the updates of global dynamic functions produced by the various procllets do not clash. This is because the update set of an algorithm, as defined in connection with the Sequential Time Postulate in Blass and Gurevich [2003], will never contain conflicting updates.

The next definition is like that in Blass and Gurevich [2003] but using our modified postulates.

Definition 26. A *parallel algorithm* is an algorithm satisfying the Sequential Time, Abstract State, Background, Procllet Algorithm, Bounded Sequentiality, and Update Postulates.

4. ABSTRACT STATE MACHINES AS ALGORITHMS

This section is devoted to showing that ASMs can be viewed as parallel algorithms in the sense defined above. That is, we prove what was claimed in Section 8.5 of Blass and Gurevich [2003]. We do not repeat all the parts of Section 8.5 that are correct but concentrate on correcting the errors.

It will be useful first to make a few comments about the nature of the examples in Section 8 of Blass and Gurevich [2003]. Those examples involved various approaches to parallel computation—PRAMs, circuits, alternating Turing machines, first-order logic, fixed-point logic, and ASMs—and indicated how the algorithms of these models fit our postulates. The most important work here was to analyze these algorithms down to the level of procllets in order to say what

the proclats and the proclat algorithm should be. Another aspect was adjoining, if necessary, the multisets, ordered pairs, etc., required by the Background Postulate, but this second aspect was fairly routine. Except for `Proclet`, everything required by the Background Postulate is determined once the “atomic” elements of the state are known. Thus, to make these other models of parallel computation fit our postulates, the procedure is roughly as follows. First analyze the algorithms to see what entities are directly involved. The small sequential processes that make up the parallel algorithm are among these entities, as proclats, and the way they work is the proclat algorithm. Then, if Boolean values, multisets, and ordered pairs are not already present, adjoin them, along with the basic operations on them. (See Blass and Gurevich [2003], Section 7, particularly Remarks 7.3 and 7.5, for comments on the naturality of adjoining these things.)

In the particular case of ASMs, the definition in Blass and Gurevich [2003], Section 9, required their vocabularies to contain everything listed in the Background Postulate. The ASMs constructed in the proof of the main result, Theorem 10.1, of Blass and Gurevich [2003] will indeed contain all these things, because they are behaviorally equivalent to algorithms, as defined by the postulates; behavioral equivalence demands that the states and therefore the vocabulary are the same. But one can also consider other ASMs, for example, the parallel ASMs of Gurevich [1995] (but without interactive features like external functions and importing reserve elements), and they should also fit our postulates when equipped with suitable proclats and a background containing multisets and ordered pairs.

For simplicity, we consider in detail only one version of ASMs, namely, that defined in Section 9 of Blass and Gurevich [2003], with one clarification. Where we said that the vocabulary of an ASM should contain “all the items required by the Background Postulate”, the last item in that postulate, “a variable-free term `Proclet ...`” (which would now become `PriProclet`) is to be omitted. If an ASM’s vocabulary happens to contain a nullary function symbol `PriProclet`, then that symbol should be renamed to avoid a conflict with the `PriProclet` involved in our postulates. Note that a symbol `PriProclet` in some arbitrary ASM need not have anything to do with the actual proclats, the sequential subprocesses of which the algorithm is composed; the name `PriProclet` could be purely accidental. We prefer to reserve this name for the actual primary proclats.

Although we concentrate on one version of ASMs, the same ideas could be used to handle other versions, for example, without multisets or without ordered pairs in the background. (These things, if missing from the ASM states, would, of course, be added when we defined the states to be used in the postulates.) We believe that the version we consider exhibits all the difficulties of other natural versions, so that, given the following treatment of it, the reader will be able to treat the others also. The vocabulary of the algorithm we describe will consist of the vocabulary of the given ASM plus two new, static, unary function symbols defined as follows:

$$\text{Mult}(x, y) = \text{multiplicity of element } x \text{ in multiset } y$$

and

$$\text{pred}(n) = \{0, 1, \dots, n - 1\}$$

for natural numbers n . (Mult and pred abbreviate *multiplicity* and *predecessors*, respectively. What exactly the natural numbers are, in a state of our algorithm, is irrelevant to the functioning of the algorithm; we postpone discussion of the issue to Remark 27 after the presentation of the algorithm.)

The exposition in Blass and Gurevich [2003], Section 8.5, began with a rough but intuitively understandable description of the algorithm, containing three explicitly acknowledged difficulties, and it continued with a discussion of how to circumvent these difficulties. In an attempt to retain intuitive understandability for our present, somewhat more complicated (but correct) algorithm, we again want to separate the main idea from the circumvention of the old difficulties. In this way, we can concentrate attention on the new aspects of the construction. Furthermore, of the three difficulties mentioned and circumvented in Blass and Gurevich [2003], two can be treated here in the same way as there, and one disappears entirely. We explain this first, in order to get all of these difficulties out of the way. (In Blass and Gurevich [2003], we described the algorithm first and eliminated the difficulties afterward, but in the present context it seems clearer to handle the difficulties first.)

The first difficulty was that the rough description assumed that the ASM never produced conflicting updates. It was solved by showing, in Blass and Gurevich [2003], Section 9.2, how to convert any ASM into one that never produces such clashes. The same solution applies in the present context.

The second difficulty was that there are infinitely many proclats, though only finitely many become active in any step. This difficulty disappears in the present context. Our postulates require only the set of primary proclats to be finite, and the algorithm we describe will involve only a single primary proclat. Any element of the state is potentially a secondary proclat, and by activating some of these we can obtain all the proclats used by the construction in Blass and Gurevich [2003]. Thus, we no longer need the terms $\text{MDA}(p)$ and $\text{MA}(p)$ that we defined—incorrectly because we used comprehension terms—on pages 624–625 of Blass and Gurevich [2003].

The third difficulty concerned cycles in the information flow digraph, where one proclat activated another and gave it some information, and then the second proclat returned some information to the first. This difficulty was solved by replacing each proclat by two or three others, each performing a part of the original proclat's task. We shall refer to these two or three proclats as *incarnations* of the original one. Thus, the first incarnation of one proclat might activate another proclat, but the reply from the second proclat would then go to the second or third incarnation of the first. (More precisely, the first incarnation of the first proclat might activate the first incarnation of the second proclat, and eventually the last incarnation of the second proclat would reply to the second or third incarnation of the first proclat.) This solution continues to work in the present setting. *Activate* has, of course, a new meaning, the meaning given by our postulates, rather than merely displaying some information that is read by the proclat to be activated. But the idea remains the same. In fact, we shall

describe our algorithm in the same format as in Blass and Gurevich [2003], numbering a procelet's tasks in a way that indicates which tasks are to be done by which incarnations.

In addition to getting these difficulties out of the way, we must make another preliminary comment, on the nature of the procleets used in our algorithm. In the rough description (before addressing the three difficulties) in Blass and Gurevich [2003], pp. 622–623, the procleets were ordered pairs $\langle \hat{Z}, \bar{a} \rangle$, where Z is an occurrence of a term or rule in the given ASM, \hat{Z} is a “name” for it, and \bar{a} is a tuple of values for its pseudofree variables. Recall that the names were assigned rather arbitrarily in Blass and Gurevich [2003]; they just need to be elements of the state that can be explicitly named, and this is easy to arrange since there are only finitely many of them. Recall also that a variable x is *pseudofree* in an occurrence Z if Z lies in the scope of a comprehension term or a parallel rule that binds x ; since an ASM program has no free variables, all the free variables of a term or rule are among its pseudofree variables. To list the values of these variables in a tuple, we implicitly assume a particular ordering of the variables. It will be convenient to assume that the pseudofree variables of any term or rule Z are listed in decreasing order of their scopes; these scopes are linearly ordered because they all contain Z . (This convention was already tacitly used in Blass and Gurevich [2003].)

Our procleets will be more complicated for two reasons. First, as indicated above in the solution of the third difficulty, each of the procleets in the rough description will actually have two or three incarnations, different procleets that divide the work in such a way as to avoid cycles. Some care will be needed in the choice of just which elements of the state serve as the second and third incarnations of procleets; we postpone the details until after we have presented enough of the algorithm to motivate the details. For now, the reader can use the following simple picture of incarnations. Because the necessary number of incarnations of a procelet $\langle \hat{Z}, \bar{a} \rangle$ is entirely determined by the nature of the term or rule Z , we can simply include, with \hat{Z} , a marker distinguishing the various incarnations of the procelet. Thus, instead of $\langle \hat{Z}, \bar{a} \rangle$, we would have $\langle \widehat{Z}, k, \bar{a} \rangle$ for the k th incarnation of $\langle \hat{Z}, \bar{a} \rangle$.

The second complication in our procleets arises from the convention that each secondary procelet p is an ordered pair whose second component is the procelet that activated p . Thus, where one might intuitively think of a chain of activations, say

$$x \in \text{PriProcelet}, \quad y \in \text{Act}(x), \quad z \in \text{Act}(y), \quad w \in \text{Act}(z),$$

the actual chain would look like

$$x \in \text{PriProcelet}, \quad y \in \text{Act}(x), \quad z \in \text{Act}(\langle y, x \rangle), \quad w \in \text{Act}(\langle z, \langle y, x \rangle \rangle),$$

the last procelet activated here being $\langle w, \langle z, \langle y, x \rangle \rangle$. In our situation, this means that our procleets will not have the simple form $\langle \widehat{Z}, k, \bar{a} \rangle$ but will be ordered pairs whose first components have this form and whose second components are themselves procleets. The procelet previously called $\langle \widehat{Z}, k, \bar{a} \rangle$ will thus encode the list of all the occurrences of terms and rules within which Z lies.

Both of these complications in the proclats—exhibiting the incarnation numbers and the ancestors of the proclats—will contribute little to the essential ideas of the algorithm that we shall describe, but they threaten to obscure the ideas by cluttering the notation. Accordingly, we adopt the convention of writing simply $\langle \hat{Z}, \bar{a}, \dots \rangle$; the intention is that the “...” reminds us of all the extra information coded in the proclat, but we refrain from exhibiting this information and thus remain close to the $\langle \hat{Z}, \bar{a} \rangle$ notation used in Blass and Gurevich [2003].

In addition to the proclats $\langle \hat{Z}, \bar{a}, \dots \rangle$ that will play essentially the same roles as $\langle \hat{Z}, \bar{a} \rangle$ in the rough description in Blass and Gurevich [2003], our algorithm will involve some additional proclats. Most of these arise from the need to avoid comprehension terms in the proclat algorithm. In describing the activity of proclats $\langle \hat{u}, \bar{a} \rangle$ where u is a comprehension term, we incorrectly used comprehension terms in Blass and Gurevich [2003]. Now, the “work” done by these comprehension terms will be described honestly, using additional proclats.

There will also be some additional proclats in our description of the activity of proclats $\langle \hat{R}, \bar{a}, \dots \rangle$ when R is a parallel rule `do forall $x \in r$, $R_0(x)$ enddo`. These are needed in order to get the right format for the subsidiary proclats $\langle \widehat{R_0(x)}, \bar{a} \frown c, \dots \rangle$.

The work of all these additional proclats will be described in the context of the work of their activators (or activators of activators). The “main” proclats $\langle \hat{Z}, \bar{a}, \dots \rangle$ serve the same purpose as $\langle \hat{Z}, \bar{a} \rangle$ did in Blass and Gurevich [2003]. When Z is a term, the purpose is to compute its value v , using \bar{a} for the values of (pseudo-)free variables, and to push v to the parent (i.e., to the activator). More precisely, it pushes the pair $\langle v, \langle \hat{Z}, \bar{a}, \dots \rangle \rangle$, so the parent knows which child computed this v . When Z is a rule, the purpose is to ensure the execution of the updates that the rule produces, again using \bar{a} to supply the values of free variables. To “ensure the execution” here means either to execute the updates or to activate enough other proclats that will ensure the execution.

In the following description of our algorithm, it is to be understood that, when a proclat (of the rough description) has several incarnations (in the precise description), then each of these incarnations except the last is to activate the next one. This activation is to be added to the activations explicitly mentioned in the following description.

In the light of the preceding discussion, we can use, in our present algorithm, much of what was done in the rough description in Blass and Gurevich [2003]. Specifically, the activity of the proclats corresponding to variables, to terms of the form $f(t_1, \dots, t_n)$, to update rules, and to conditional rules can be described exactly as in Blass and Gurevich [2003], pp. 622–623, with just two modifications:

- change every $\langle \hat{Z}, \bar{a} \rangle$ to $\langle \hat{Z}, \bar{a}, \dots \rangle$;
- delete the parenthetical comment that activation is done “by displaying an appropriate signal,” since activation is now done by updating `myAct`.

Notice that, in each of these cases, a proclat activates a bounded set of other proclats, so the desired `myAct` can be explicitly given by the proclat algorithm.

There remain the proclats corresponding to comprehension terms and to

parallel rules—the two ASM constructs that introduce unbounded parallelism. For these proclats, we cannot proceed exactly as in Blass and Gurevich [2003]. Indeed, the instructions in Blass and Gurevich [2003] for the second incarnation of such a proclat (i.e., the instructions labeled (2) in the rough description) involve activating an unbounded set of secondary proclats. To do so, a proclat would update `myAct` to mark a set for activation, but the necessary set could be described only by a comprehension term. The following instructions for these two sorts of proclats avoid this difficulty by marking for activation only a set directly available to the activating proclat, that is, a set that can be named in the proclat’s initial local state. The price for this is that the comprehension term to be avoided involved some parallel work, which must now be done by additional proclats. Here are the details.

Let p be the proclat $\langle \hat{u}, \bar{a}, \dots \rangle$ where u is a comprehension term $\{\{t(x) : x \in r : \varphi(x)\}\}$. As in Blass and Gurevich [2003], the work of p will be in three parts (i.e., p will have three incarnations).

- (1) activate $\langle \hat{r}, \bar{a}, \dots \rangle$;
- (2) after receiving the value b computed by this secondary proclat, display b and mark `AsSet(b)` for activation;
- (3) push `(myMail, me)` to your parent.

Obviously, for this to be correct, a lot has to happen between instructions (2) and (3). Specifically, the proclats activated in (2) must somehow ensure that the mailbox of p is exactly the multiset that p is supposed to compute, the value of $\{\{t(x) : x \in r : \varphi(x)\}\}$ when the free variables have values given by \bar{a} . That is achieved as follows.

The proclats activated as a result of (2) are $\langle c, p, \dots \rangle$ for all $c \in b$. (The “...” here refers only to an indication of the fact that the activator is the second incarnation of p .) Each such $\langle c, p, \dots \rangle$ does the following.

- (1) Read b from the display of (the second incarnation of) p and mark for activation

$$\{\{\widehat{t(x)}, \bar{a} \frown c\}, \{\widehat{\varphi(x)}, \bar{a} \frown c\}\} \uplus \text{pred}(\text{Mult}(c, b)).$$

Recall that $\text{pred}(\text{Mult}(c, b)) = \{0, 1, \dots, m - 1\}$ where m is the multiplicity of c as an element of b .

- (2) When the proclats $\langle \widehat{t(x)}, \bar{a} \frown c \rangle$ and $\langle \widehat{\varphi(x)}, \bar{a} \frown c \rangle$ return values, if $\varphi(c) = \text{true}$ then display $\{t(c)\}$; otherwise display \emptyset .

The additional proclats $\langle k, \dots \rangle$ (for $0 \leq k < \text{Mult}(c, b)$) activated here read the display of (the second incarnation of) $\langle c, p, \dots \rangle$ and, if it is nonempty, extract its element (using `TheUnique`) and mail that to (the third incarnation of) p . As a result of all this work, the proclats activated by $\langle c, p, \dots \rangle$ will contribute to the mailbox of p either nothing, if $\varphi(c) = \text{false}$, or exactly $\text{Mult}(c, b)$ copies of $t(c)$, if $\varphi(c) = \text{true}$. Combining the results from all elements c of `AsSet(b)`, we obtain as the mailbox of p exactly the required multiset, the value of $\{\{t(x) : x \in r : \varphi(x)\}\}$ with free variables interpreted according to \bar{a} .

Finally, we consider the somewhat easier case of a parallel rule. Let p be

the proclat $\langle \hat{R}, \bar{a}, \dots \rangle$, where R is the rule `do forall $x \in r, R_0(x)$ enddo`. Its instructions are

- (1) activate $\langle \hat{r}, \bar{a}, \dots \rangle$; and
- (2) after receiving the value b computed by this secondary proclat mark `AsSet(b)` for activation.

Each proclat $\langle c, p, \dots \rangle$ activated in (2) merely activates $\langle \widehat{R_0(x)}, \bar{a} \frown c, \dots \rangle$. These proclats will then ensure the execution of $R_0(c)$ for all $c \in b$, as required.

To complete the description of our algorithm, we still have two tasks to accomplish. We must provide the term `PriProclat`, and we must keep our promise to provide details about which elements of the state serve as the second and third incarnations of our proclats.

The first of these tasks is remarkably easy. We need only a single primary proclat, the one corresponding to the whole ASM program Π , considered as a rule; all other proclats that we need are activated during the execution of the algorithm. We therefore define `PriProclat` to be $\{\langle \hat{\Pi}, \langle \rangle \rangle\}$. (Here $\langle \rangle$ is the empty tuple, since Π has no pseudofree variables.)

Finally, we turn to incarnations. When a proclat p activates its next incarnation p' , it does so by including an appropriate element q in its `myAct`, so that $p' = \langle q, p \rangle$. But what is an appropriate q ? The proclat $\langle q, p \rangle$ had better be distinct from all the other proclats activated by p . We need not worry about coincidences with other secondary proclats, activated by proclats other than p , for these will have their activators, not p , as their second components. We also need not worry about a coincidence with the primary proclat $\langle \hat{\Pi}, \langle \rangle \rangle$, since its second component $\langle \rangle$ is distinct from p . In most cases, our description of the algorithm tells what the other proclats activated by p are, and it is easy to find a suitable q . If, as suggested in Blass and Gurevich [2003], we use certain multisets as the codes \hat{Z} , then `true` will work as the required q in all but two cases. To see this, just use the fact that `true` is not a multiset or an ordered pair or a natural number. (See Remark 27 below about natural numbers.) The two cases where this choice might fail are those where p is the second incarnation of $\langle \hat{Z}, \bar{a}, \dots \rangle$ and Z is a comprehension term or a parallel rule. In those cases, the other proclats activated by p include $\langle c, p \rangle$ for all members c of b (in the notation used for describing the algorithm above). So we must ensure that $q \notin b$. Fortunately, there is a standard (in set theory) trick for getting an object that is not a member of a given (multi)set b , namely, to take the object b itself. (There are also other options, for example, $\{\{b\}\}$.) Thus, we can obtain the next incarnation, in the two cases where $q = \text{true}$ might not work, by taking $q = b$ instead. One final comment is needed here, namely, that, in these cases, p should display b , so that the proclats it creates can tell, by reading the display, whether they are the next incarnation of p or one of the other proclats, $\langle c, p \rangle$ for $c \in b$, that p activated.

Remark 27. Our algorithm depended on the availability of natural numbers and the functions `Mult` and `pred`. There are at least two plausible ways to represent natural numbers by elements of the states of our algorithms. One is the coding, due to von Neumann, that has become standard in set theory. It sets

(inductively) $n = \{0, 1, \dots, n - 1\}$. With this coding pred is simply the identity function. An alternative coding, quite natural when multisets are available, is to represent n by a multiset consisting of n copies of some standard entity, for example true or \emptyset .

But in fact, we could do without any particular representation of the natural numbers. The algorithm never used Mult and pred individually but only in the combination $\text{pred} \circ \text{Mult}$, and it never mattered that the elements of $\text{pred} \circ \text{Mult}(c, b)$ were numbers, only that they were $\text{Mult}(c, b)$ distinct objects, all distinct from true . (Distinctness from true was used only in our discussion of using $q = \text{true}$ for producing second and third incarnations.) Thus, we could simply assume the existence of some such function to serve in place of $\text{pred} \circ \text{Mult}$.

5. ALGORITHMS ARE ABSTRACT STATE MACHINES

Our goal in this section is to show that the ASM thesis, that all algorithms are behaviorally equivalent to ASMs, holds for parallel algorithms in the sense of Definition 26. That is, the thesis is not damaged by our extension of the notion of parallel algorithm from Blass and Gurevich [2003] to allow intrastep creation of proplets. We do not change the definitions of ASMs (see Blass and Gurevich [2003], Section 9.1), and of behavioral equivalence (see Blass and Gurevich [2003], Definition 2.3). Recall that behavioral equivalence is a very strong equivalence relation, requiring the same states, the same initial states, and the same one-step transition function. In particular, when we construct an ASM equivalent to a given algorithm A , it is obvious what the states and initial states of the ASM must be; the only issue is constructing an ASM program that produces the same transition function as A .

THEOREM 28. *Every parallel algorithm is behaviorally equivalent to an ASM.*

PROOF. The proof is very similar to the proof of the corresponding, weaker result, Theorem 10.1, in Blass and Gurevich [2003], so we only explain the changes that are required by our present, broader notion of algorithm.

The first part of the proof in Blass and Gurevich [2003], expressing the proplet algorithm as a sequential ASM with output, Π , requires only a notational change: Since the initial local state now interprets the dynamic, nullary symbol myAct , this symbol is added to the vocabulary of Π .

The next part of the proof in Blass and Gurevich [2003], starting at the bottom of page 631, gives a rough description of how the ASM's computation will proceed. As pointed out there, the most natural approach to computing the correct ken , namely, imitating the level-by-level recursion used in the proof of Theorem 7.22, needs to be modified because information about the levels is not available to the proplets. So in Blass and Gurevich [2003] the rough description of the ASM's work is called a "rough description of the modification" of the natural approach.

We must now make a further modification because the notion of proplet is not fixed by the state (as it was in Blass and Gurevich [2003]) but changes from

phase to phase as a result of activations. The work of the ASM still proceeds in B phases, with all procllets executing the procllet algorithm in certain initial local states at each phase, but the set of procllets here depends on the phase. In the first phase, the procllets are the elements of PriProclet . At any later phase, say phase k , the procllets are the elements of PriProclet and the elements activated by procllets at phase $k - 1$. The latter are, of course, the elements of the form $\langle q, p \rangle$ where p was a procllet at phase $k - 1$ and it updated its myAct to a set containing q .

At each phase, the procllets of that phase execute the procllet algorithm in an initial local state where myMail and Display are interpreted as the results of the preceding phase (with $\text{myMail} = \emptyset$ and $\text{Display}(q) = \text{undef}$ in phase 1), exactly as in Blass and Gurevich [2003]. The argument in Blass and Gurevich [2003], p. 632, showing that, in phase k , all procllets of level $\leq k$ are computing in the initial local states given by the correct ken carries over to the present context. It follows that these procllets do the correct pushing, displaying, and activating at this phase and that, as a result, the next phase has the correct set of procllets up to and including level $k + 1$. (The reference to Lemma 7.25 in Blass and Gurevich [2003] is now replaced with a reference to Lemma 21.)

Just as in Blass and Gurevich [2003], the rough description must be supplemented with a decision to suppress all updates of the global state until phase B , during which the ASM is using the correct set of procllets with the correct initial local states at all levels.

In the formal presentation of the ASM, starting on page 633 of Blass and Gurevich [2003], the following extensions are needed to handle activation of procllets. First, in addition to the terms (or more precisely term schemas) $\text{Outmail}(p, M, D)$ and $\text{Dspl}(p, M, D)$ used there, we also have $\text{Asp}(p, M, D)$. Here Asp stands for *activate secondary procllets*; the intended meaning is that procllet p , with M as its mailbox and D as the display function, would execute updates of myAct , and $\text{Asp}(p, M, D)$ is the multiset of all the elements x such that p would update myAct to x , just as in Blass and Gurevich [2003]. $\text{Dspl}(p, M, D)$ is the multiset of those x such that p updates myDisplay to x . (The ASM program Π can be arranged so that it produces at most one update of myAct and myDisplay , but it is convenient to give the definitions in a general form that does not presuppose this uniqueness.)

The definitions of Outmail_R and Dspl_R by induction on rules R , as given in Blass and Gurevich [2003], must be supplemented with clauses to define Asp_R . To formulate these clauses, we use for terms t the notation t' , defined just as in Blass and Gurevich [2003] with the extra clause that myAct in t is to be replaced with \emptyset in t' . Now the clauses defining Asp_R are exactly analogous to the clauses for Dspl_R in Blass and Gurevich [2003]:

- If R is an update rule of the form $\text{myAct} := t$, then $\text{Asp}_R(p, M, D)$ is $\{\{t'\}\}$;
- If R is any other update rule, then $\text{Asp}_R(p, M, D)$ is \emptyset ;
- If R is $\text{Push } t_0 \text{ to } t_1$, then $\text{Asp}_R(p, M, D)$ is \emptyset ;
- If R is $\text{do in parallel } R_0, \dots, R_k \text{ enddo}$; then $\text{Asp}_R(p, M, D)$ is the sum $\text{Asp}_{R_0}(p, M, D) \uplus \dots \uplus \text{Asp}_{R_k}(p, M, D)$;

—If R is if φ then R_0 else R_1 endif, then $\text{Asp}_R(p, M, D)$ is

$$\{\{z : z \in \text{Asp}_{R_0}(p, M, D) : \varphi'\} \uplus \{z : z \in \text{Asp}_{R_1}(p, M, D) : \neg\varphi'\}.$$

Continuing in analogy with how we handled Dsp1 , we define $\text{Asp}(p, M, D)$ to be $\text{TheUnique}(\text{AsSet}(\text{Asp}_\Pi(p, M, D)))$. One more definition is needed, to incorporate both the automatic tagging of secondary proclefs with their activators and the convention that, if a proclef marks for activation something other than a set, then it thereby activates nothing. Accordingly, we define $\text{TP}(p, M, D)$ to be

$$\{\langle q, p \rangle : q \in \text{Asp}(p, M, D) : \text{AsSet}(\text{Asp}(p, M, D)) = \text{Asp}(p, M, D)\};$$

the notation TP stands for *tagged proclefs*.

Next, we modify the definitions in Blass and Gurevich [2003], p. 634, of $\text{Mailbox}_k(p)$ and $\text{Display}_k(p)$ to take into account the possible variation of the set of proclefs from one phase to another. The (unique) occurrence of Proclef in these definitions is to be replaced with Proclef_k , which in turn is defined by induction on k simultaneously with $\text{Mailbox}_k(p)$ and $\text{Display}_k(p)$, as follows:

— Proclef_0 is PriProclef ;

— Proclef_{k+1} is $\text{PriProclef} \uplus \{\{\text{TP}(p, \text{Mailbox}_k(p), \text{Display}_k) : p \in \text{Proclef}_k : \text{true}\}\}$.

It was shown in Blass and Gurevich [2003] that $\text{Mailbox}_k(p)$ and $\text{Display}_k(p)$ give the mailbox and display functions after k phases of the computation in our description of how the desired ASM works. The argument there extends to show that $\text{TP}_k(p)$ gives the set of proclefs activated by p in phase k and therefore that Proclef_k is the set of proclefs that execute during phase $k + 1$. Arguing as on page 635 of Blass and Gurevich [2003], we find that the updates of the given algorithm A are matched by the ASM program

do forall $p \in \text{Proclef}_{B-1}$ $\Pi^*(p)$ enddo

(the same as in Blass and Gurevich [2003] except for the subscript $B - 1$ on Proclef), where $\Pi^*(p)$ is obtained from Π by the same substitutions as in Blass and Gurevich [2003], except that Skip replaces not only updates of myDisplay but also updates of myAct . \square

ACKNOWLEDGMENTS

We thank Dean Rosenzweig for pointing out that we had incorrectly used comprehension terms in Blass and Gurevich [2003], Section 8, and for subsequent helpful discussions.

REFERENCES

- BLASS, A. AND GUREVICH, Y. 2003. Abstract state machines capture parallel algorithms. *ACM Trans. Computat. Log.* 4, 578–651.
- BLASS, A. AND GUREVICH, Y. 2006. Ordinary interactive small-step algorithms, I. *ACM Trans. Computat. Log.* 7, 363–419.
- BLASS, A., GUREVICH, Y., ROSENZWEIG, D., AND ROSSMAN, B. 2006. Interactive small-step algorithms I: Axiomatization, and II: Abstract state machines and the characterization theorem. Microsoft

Research Tech. Rep. MSR-TR-2006-170 and MSR-TR-2006-171. Microsoft Research, Redmond, WA. *Logic. Math. Comput. Sci.* To appear.

GUREVICH, Y. 1995. Evolving algebra 1993: Lipari guide. In *Specification and Validation Methods*, E. Börger, Ed. Oxford University Press, Oxford, 9–36.

GUREVICH, Y. 2000. Sequential abstract state machines capture sequential algorithms. *ACM Trans. Computat. Log.* 1, 77–111.

Received August 2006; accepted January 2007