MAX-PLANCK-INSTITUT FÜR INFORMATIK

Fast Parallel Space Allocation, Estimation and Integer Sorting (revised)

Holger Bast and Torben Hagerup

MPI–I–93–123 June 1993

INFORMATIK

Im Stadtwald
W 6600 Saarbrücken
Germany

Fast Parallel Space Allocation, Estimation and Integer Sorting (revised)

Holger Bast and Torben Hagerup

MPI–I–93–123 June 1993

Fast Parallel Space Allocation, Estimation and Integer Sorting (revised) *

HOLGER BAST

TORBEN HAGERUP

Max-Planck-Institut für Informatik

D–66123 Saarbrücken, Germany

Abstract

The following problems are shown to be solvable in $O(\log^* n)$ time with optimal speedup with high probability on a randomized CRCW PRAM using $O(n)$ space:

- Space allocation: Given n nonnegative integers $x_1, \ldots, x_n$, allocate n nonoverlapping blocks of consecutive memory cells of sizes $x_1, \ldots, x_n$ from a base segment of $O(\sum_{j=1}^{n} x_j)$ consecutive memory cells;
- Estimation: Given n integers in the range $1..n$, compute "good" estimates of the number of occurrences of each value in the range $1..n$;
- Semisorting: Given n integers $x_1, \ldots, x_n$ in the range $1..n$, store the integers $1, \ldots, n$ in an array of $O(n)$ cells such that for all $i \in \{1, \ldots, n\}$, all elements of $\{j : 1 \leq j \leq n \text{ and } x_j = i\}$ occur together, separated only by empty cells;
- Integer chain-sorting: Given n integers $x_1, \ldots, x_n$ in the range $1..n$, construct a linked list containing the integers $1, \ldots, n$ such that for all $i, j \in \{1, \ldots, n\}$, if i precedes j in the list, then $x_i \leq x_j$.

Moreover, given slightly superlinear processor and space bounds, these problems or variations of them can be solved in constant time with high probability.

As a corollary of the integer chain-sorting result, it follows that n integers in the range $1..n$ can be sorted in $O(\log n / \log\log n)$ time with optimal speedup with high probability.

---

* Supported by the ESPRIT Basic Research Actions Program of the EC under contracts No. 3075 and 7141 (projects ALCOM and ALCOM II). A preliminary version of this paper was presented at the 23rd Annual ACM Symposium on Theory of Computing (STOC 91) under the title "Constant-Time Parallel Integer Sorting".

1 Introduction

This paper studies a number of problems that are of fundamental importance in parallel computing. Most of these have traditional, "exact" variants that are known not to possess fast parallel solutions. More precisely, computing the parity of n bits reduces to instances of these problems of size n, which, therefore, by the lower bound of Beame and Hastad (1989), cannot be solved faster than in $\Theta(\log n/\log\log n)$ time on a PRAM with any polynomial number of processors. Relaxing the problem definitions to allow approximate solutions, however, we are able to obtain very fast algorithms that run with optimal speedup on a randomized CRCW PRAM.

The first problem studied is that of *space allocation*, which we formalize as the *interval allocation* problem. Imagine that we are presented with n simultaneous requests, each of which is for a certain number of consecutive memory cells. Note that a request is not for specific memory cells, but merely indicates the number of cells needed. Abstractly speaking, such requests might originate with a collection of concurrently executing tasks, each of which needs a certain amount of working space for its computation. The present paper provides several concrete examples of situations where such requests arise naturally; many more can be found in the papers cited below. Given the set of requests, the goal is to satisfy each request, i.e., to supply the requesting agent with a private block of memory of the requested size.

We may view the allocated blocks as nonoverlapping subarrays of a single base array. The exact version of the interval allocation problem requires the size of the base array to exactly equal the sum of all requested sizes, and is clearly subject to the lower bound mentioned above. We must therefore relax this requirement, but still want to insist that not too much space be wasted. For reasons similar to those that motivate the use of the O-notation, we require the size of the base array to be at most a constant factor larger than the sum of the requested sizes. With this convention, we are able to solve interval allocation problems of size n in $O(\log^* n)$ time with optimal speedup with high probability. As shown by MacKenzie (1992), this is as fast as possible for any algorithm that uses no more than n processors. A variant of interval allocation called *interval marking* is a natural formalization of the (vaguely defined) *processor allocation* problem, which adds to the importance of the interval allocation problem.

The second problem studied is that of *profiling*. We are here given an array of n keys, and the task is to determine the multiplicity (i.e., the number of occurrences) of each value represented among the keys. The exact version of the problem asks for the exact multiplicities and, again, is clearly subject to the lower bound of $\Omega(\log n/\log\log n)$. We therefore content ourselves with approximate counts. We assume that the values represented among the keys are integers in a range $1\,..\,m$, where m is at most linear in n.

If this is not the case initially, static hashing can frequently be used to map the original values injectively to a sufficiently small range of integers, after which approximate counts can be computed using our algorithms and associated with the original values (see (Bast and Hagerup, 1991) and (Bast *et al.*, 1992) for the best known results on static hashing). The output hence takes the form of m nonnegative integers $\widehat{b}_1, \ldots, \widehat{b}_m$, where $\widehat{b}_i$ is an estimate of the number b_i of occurrences of the value i, for $i = 1, \ldots, m$.

We study three different variants of the profiling problem. In the first of these, the number of different values is assumed to be much smaller than the number of keys; specifically, $m = O(n^{1-\delta})$, for some fixed $\delta > 0$. In these circumstances, constant time and n processors suffice, with high probability, to compute what we call a *fine-profile*, i.e., a sequence $\widehat{b}_1, \ldots, \widehat{b}_m$ with $b_i \leq \widehat{b}_i \leq Kb_i$, for some constant $K \geq 1$ and for $i = 1, \ldots, m$. This simple result furnishes a basic tool used in many of our other algorithms.

The second and third variant of the profiling problem are concerned with the case $m = n$, which is of special interest and importance (see below). The second variant requires the estimates $\widehat{b}_1, \ldots, \widehat{b}_n$ to be independent random variables (note that $\widehat{b}_1, \ldots, \widehat{b}_m$ are random variables because the execution of one of our (randomized) algorithms constitutes a random experiment; the input is considered to be fixed). Given the difficulty of analysis often caused by a lack of independence, this is a reasonable property for which to ask. As concerns the accuracy of the estimates, we require on the one hand that $\sum_{i=1}^{n} \widehat{b}_i = O(n)$, a natural condition, and on the other hand that $\Pr(b_i > a\widehat{b}_i) \leq 2^{-a}$, for $i = 1, \ldots, n$ and for all $a \geq 1$ (i.e., the probability of an estimate being a times too small decreases exponentially in a), a less natural condition that represents a compromise between what we would ideally like and what we can easily compute. We show that estimates $\widehat{b}_1, \ldots, \widehat{b}_n$ with this property, called a *coarse-profile*, can be computed in $O((\log^* n)^2)$ time with optimal speedup with high probability. In the third variant of profiling, we give up on the independence of $\widehat{b}_1, \ldots, \widehat{b}_n$ and instead try to compute the estimates faster and to obtain more accurate estimates.

The third problem studied is that of *semisorting* (the term was taken from (Valiant, 1990)). To semisort a sequence of objects, each with a distinguished key, is to rearrange the objects so that all objects with a common key occur together. We assume that the keys are integers in the range $1..n$; as above, static hashing can often be used to enforce this condition if it is not satisfied initially. The lower bound of $\Omega(\log n / \log\log n)$ applies to semisorting, as defined so far, so we relax the definition by allowing the output to be given in the form of a *padded sequence* of size $O(n)$, i.e., $O(n)$ special *null* objects are allowed to intervene in arbitrary positions between the n objects that form the actual semisorted sequence. Our result is that semisorting problems of size n can be solved in $O(\log^* n)$ time with optimal speedup with high probability. The proof is quite involved and makes crucial use of the results obtained for the third variant of profiling — the

condition $\Pr(b_i > a\widehat{b}_i) \leq 2^{-a}$ or, rather, a variation of it turns out to be exactly what is needed. We extend the semisorting result to *strong semisorting*, which requires that all occurrences of a key of multiplicity b appear in a subarray of the output array of size $O(b)$.

Semisorting has many and diverse uses. Our result on strong semisorting directly provides one of our best profiling results, namely a fine-profile for the case $m = n$. Another immediate application is to integer *chain-sorting*. General chain-sorting takes as input n keys drawn from a totally ordered universe and makes each key point to the next larger key, if any (with an arbitrary total order imposed by the algorithm on each set of keys of a common value), i.e., the keys are stored in sorted order in a linked list. We consider the chain-sorting problem with integer input keys drawn from the range $1..n$. In contrast with what is the case for profiling and semisorting, a preprocessing based on hashing, which is a nonmonotonic operation, does not enable our integer chain-sorting algorithm to cope with more general input keys; allowing only keys in the range $1..n$ therefore is a true restriction. In recognition of this fact, we continue to use the term "integer chain-sorting", rather than simply "chain-sorting". Note also that the lower bound of Beame and Hastad does not apply to chain-sorting, even with no restriction on key values. On the other hand, the well-known lower bound of $\Omega(n \log n)$ for (randomized) comparison-based sequential sorting, which holds also for chain-sorting, implies that our result, $O(\log^* n)$ time with optimal speedup with high probability, does not extend from integer chain-sorting to general chain-sorting. As a rather trivial by-product of our fast chain-sorting algorithm, we are able to improve the best previous result on (standard) randomized sorting of n integers in the range $1..n$: We show that this problem can be solved in $O(\log n / \log\log n)$ time with optimal speedup with high probability.

More substantial applications of our semisorting result were reported elsewhere. Semisorting is used in (Hagerup, 1992a, 1992b) to simulate stronger PRAM variants on the weaker TOLERANT PRAM; semisorting there serves to bring together all write requests pertaining to a common memory cell. Hagerup and Katajainen (1993) employ semisorting in the construction of the Voronoi diagram of n random sites drawn independently from the uniform distribution over the unit square; a grid divides the unit square into approximately n cells, and the set of sites in each cell is computed by means of semisorting. Our result also allows a significant simplification of the hashing scheme of (Bast and Hagerup, 1991). In (Hagerup and Raman, 1992), finally, semisorting is used for a variety of different purposes.

From a different point of view, the present paper explores the power flowing from a combination of three new techniques in algorithm design and analysis: First, the "log-star" technique introduced by Raman (1990) and developed further by Matias and Vishkin (1991). Second, randomized "scattering" procedures for estimating various

quantities crudely, but rapidly. And third, the analysis of randomized algorithms using martingale theory, which is not new, but which in the past has not been used as often as it deserves. A less detailed and more accessible account of most of the material in this paper can be found in (Hagerup, 1992b); the reader may want to study that paper before taking on the present one.

The structure of the paper is as follows: After some preliminaries in Section 2, Section 3 introduces various concepts under the general heading of "scattering" and lists some of their basic properties. Section 4 deals with a special case of interval allocation called *compaction*, and Section 6 extends this to so-called *colored compaction*. Section 5 presents first results for the fine-profiling problem, and Section 7 uses the results of Sections 5 and 6 to solve the interval allocation problem. Sections 8 and 9 are devoted to coarse-profiling and semisorting, respectively, and Section 10 describes applications of semisorting. Section 11, finally, studies the consequences of allowing slightly superlinear processor and space bounds. Every section uses essentially all sections before it, so that it is difficult to read sections out of context.

2 Preliminaries

A *CRCW PRAM* (concurrent-read concurrent-write parallel random access machine) is a synchronous parallel machine with processors numbered $1, 2, \ldots$ and with a global memory that supports concurrent (i.e., simultaneous) access to a single cell by arbitrary sets of processors. The semantics of concurrent writing can be defined in many ways. Accordingly, many different variants of the CRCW PRAM, each distinguished by a different rule for the resolution of write conflicts, have been introduced; see, e.g., (Chlebus *et al.*, 1989; Hagerup and Radzik, 1990; Hagerup, 1992a) for definitions of many of these models and for discussion of the relationships between them. The following two write conflict resolution rules and corresponding variants are relevant to the present paper:

ARBITRARY (Shiloach and Vishkin, 1982): If two or more processors attempt to write to a given cell in a given step, then one of them succeeds, but there is no rule assumed to govern the selection of the successful processor;

TOLERANT (Grolmusz and Ragde, 1987): If two or more processors attempt to write to a given cell in a given step, then the contents of that cell do not change.

It is easy to see that the ARBITRARY PRAM is (not necessarily strictly) stronger than the TOLERANT PRAM in the sense that one step of a TOLERANT PRAM can be simulated by a constant number of steps on an ARBITRARY PRAM with the same number of processors and memory cells. In fact, most CRCW PRAM models commonly considered are stronger than the TOLERANT PRAM in this sense. We employ the TOLERANT model throughout

the paper, with the sole exception of Lemmas 3.4(b) and 3.5 and Theorem 11.6 and their proofs, which use the ARBITRARY model. The most direct implementation of some of our other algorithms, however, assumes the ARBITRARY model, and we have to put in an extra effort in order to derive a solution for the weaker TOLERANT PRAM. Since we expect the distinction between different variants of the CRCW PRAM to be of little concern to many readers, we try to make it possible to skip material that deals only with the translation between models.

Consider the following assertion: "Every problem that can be solved in τ time steps with p processors can also, for every given $k \in \mathbb{N}$, be solved in $O(k\tau)$ time with $\lceil p/k \rceil$ processors". A simple but important simulation shows the assertion to hold for the ARBITRARY PRAM: Each physical processor simulates up to k virtual processors in a step-by-step fashion. We express this by saying that the ARBITRARY PRAM is *self-simulating* and sometimes use the word "processor" to denote a virtual processor in the sense of this simulation. The number of *operations* executed by a parallel algorithm that uses τ time steps and p processors is defined to be its time-processor product $p\tau$. By the above simulation, we always have $p\tau = \Omega(T)$, where T is the sequential complexity of the problem solved by the algorithm. Accordingly, the parallel algorithm is said to have *optimal speedup* or to be *optimal* if $p\tau = O(T)$. Because of the self-simulating property, if a problem can be solved on an ARBITRARY PRAM using t time steps and q operations, then it can also, for every given $\tau \geq t$, be solved in $\Theta(\tau)$ time using $O(q + \tau)$ operations, i.e., the algorithm can be slowed down without loss. This makes it convenient to express the performance of the algorithm by giving the pair (t, q) of minimum computation time and number of operations. In contrast with all other commonly considered PRAM variants, the TOLERANT PRAM is not known to be self-simulating. Since it is still important to know the extent to which a particular algorithm can be slowed down (see below), we are forced to indicate this explicitly, typically in a statement of the form "For all $\tau \geq \log^* n$, $O(\tau)$ time and $\lceil n/\tau \rceil$ processors suffice to...". We advise the reader to interpret such a statement as "The time is $O(\log^* n)$, and the algorithm is optimal and can be slowed down". Note that if an algorithm consists of l parts with (minimal time, number of operations) performance pairs $(t_1, q_1), \ldots, (t_l, q_l)$ and if each part can be slowed down, in the above sense, then the whole algorithm has a performance pair (t, q), where $t = O(\sum_{i=1}^{l} t_i)$ and $q = O(\sum_{i=1}^{l} q_i)$. It is also easy to see that when $k \in \mathbb{N}$ is a constant, we can always reduce the number of processors from p to $\lceil p/k \rceil$, even on the TOLERANT PRAM, without increasing the processing time by more than a constant factor and the space requirements by more than $O(p)$. We shall freely use this observation, which was also made in (Gil, 1990).

The majority of our algorithms are randomized. Randomized algorithms are customarily divided into *Monte Carlo* algorithms, which may occasionally err, and *Las*

Vegas algorithms, which never err, but which may either take a long time to produce a (correct) result, or finish on time without producing any result — it is easy to transform any Las Vegas algorithm from one of these modes of operation to the other. In all cases, the analysis of a randomized algorithm bounds the probability of the relevant undesirable behavior, which we call the *failure probability* (for a Las Vegas algorithm, this is a slight misnomer).

A Las Vegas algorithm is clearly more desirable than a Monte Carlo algorithm, since it is trivial to run a Las Vegas algorithm as a Monte Carlo algorithm: If the algorithm has not produced any result within a suitable response time, abort it (if it is still running) and output an arbitrary value. The distinction between Monte Carlo algorithms and Las Vegas algorithms will be crucial at one point of our exposition. At any rate, although this is not always done, be believe that it is important to classify each randomized algorithm as either Monte Carlo or Las Vegas. We will do so by appending the appropriate one of "(Monte Carlo)" and "(Las Vegas)" to the bound on the failure probability given for each algorithm.

Informally, a randomized algorithm can be formulated as a Las Vegas algorithm whenever its output can be verified using negligible resources, either after the fact or by the algorithm itself — this is because the algorithm can be executed repeatedly until some output passes the verification. Whenever we classify an algorithm as a Las Vegas algorithm, it will be easy to see that such verification is possible, and we will not demonstrate this explicitly.

As usual, $E_1 = O(E_2)$, where E_1 and E_2 are expressions, means that there is a constant $c > 0$ such that $E_1 \leq cE_2$. Note that we require this relation to hold for all legal values of the parameters occurring in E_1 and E_2, not just for sufficiently large values of these parameters. The constant c is independent of all other parameters, except that it may depend on quantities that are explicitly qualified as "fixed" (in the present paper, such quantities are always denoted by the symbols δ and μ). The meaning of $E_1 = \Omega(E_2)$ is defined analogously.

In order to make many proofs more readable, we make extensive use of the notion of a *negligible probability*. What constitutes a negligible probability depends on the context. Is the goal, e.g., to show that some event occurs with probability $2^{-n^{\Omega(1)}}$, then in the proof we can ignore any polynomial (in n) number of probabilities of the form $2^{-\Omega(n^{\epsilon})}$, for arbitrary $\epsilon > 0$, since for sufficiently large values of n the sum of such probabilities will be bounded by $2^{-n^{\epsilon'}}$, for suitable $\epsilon' > 0$; we here rely on the fact that all problems considered in the paper become trivial if the problem size n is bounded by a constant. An event that occurs *with high probability* is the complement of an event of negligible probability. We often tacitly assume that such events always occur. Whenever we speak of "choosing at random", we mean choosing from the uniform distribution and

independently of other such choices. We assume processors to be equipped with unit-time operations for integer addition, subtraction, multiplication and division with remainder, for computing 2^s, for every given $s \in \mathbb{N}$, and for choosing a random integer from the set $\{1, \ldots, s\}$, for every given $s \in \mathbb{N}$. As a minimum, we assume that the available operations can be executed in constant time for integer operands and results of absolute value bounded by $n + m + p$, where n is the input size, m is the largest absolute value of an input number, and p is the number of processors of the machine under consideration, i.e., a (standard) logarithmic word length suffices.

When nothing else is stated, arrays are assumed to be one-dimensional. Given an array A, we denote by $|A|$ its *size*, i.e., the number of cells in A. Although, in principle, a memory cell contains a single integer, we often find it convenient to pretend that a cell can contain an entire record of an arbitrary, but constant, number of (integer) fields; achieving this is simply a matter of considering a constant number of cells as a unit, also called a cell. When we state that a problem can be solved by a certain number of processors or speak of allocating a certain number of processors to some task, we always assume the processors to be consecutively numbered. Without stating this explicitly, we also assume that each processor "comes equipped with" a cell indexed by its processor number in a suitable array shared by all processors, which can be used for coordination between processors working on a common task. One consequence of this is that our processor bounds are always dominated by our space bounds.

We use "log" to denote the binary logarithm function. For $k = 0, 1, \ldots$, $\log^{(k)}$ denotes k-fold application of the function $x \mapsto \max\{\log x, 1\}$, i.e., $\log^{(0)} x = x$ and $\log^{(k)} x = \max\{\log \log^{(k-1)} x, 1\}$, for all $x > 0$ and all $k \in \mathbb{N}$. For $n \in \mathbb{N}$, $\log^* n = \min\{k \in \mathbb{N} : \log^{(k)} n = 1\}$. Although extracting logarithms is not one of our standard operations, we will assume that the function $x \mapsto \lfloor \log x \rfloor$ can be evaluated in constant time by a single processor, for $x \in \{1, \ldots, n\}$. This is justified by an observation of Hagerup and Radzik (1990), who show that a table realizing this function on the domain in question can be constructed in constant time with n processors. As a consequence, $\lfloor \log^{(k)} n \rfloor$ can be computed in $O(k)$ time by a single processor, for arbitrary given $k \in \mathbb{N}$. It is also easy to see that for any given rational number q, the function $x \mapsto \lfloor x^q \rfloor$ can be evaluated in constant time with n processors, for $x \in \{1, \ldots, n\}$ (details are given in (Hagerup, 1992d)).

Some of the constants appearing in our proofs are very large. This is not evidence of a true problem, but merely reflects a decision never to add to the complexity of an argument in order to obtain smaller constants. In particular, we make frequent use of the very crude estimates $2^x \geq x$, for all $x \geq 0$, and $\lceil x \rceil \leq 2\lfloor x \rfloor$, for all $x \geq 1$. We expect that a less generous analysis would yield reasonable constant factors.

Our probabilistic analysis is based on the two lemmas below. Lemma 2.1 states

various inequalities commonly known as Chernoff bounds.

Lemma 2.1: For every binomially distributed random variable S, the following holds:
(a) For all $z \geq 2E(S)$, $\Pr(S \geq z) \leq e^{-z/6}$;
(b) $\Pr(S \leq E(S)/2) \leq e^{-E(S)/8}$;
(c) For all $z > 0$, $\Pr(S \geq z) \leq \left(\frac{eE(S)}{z}\right)^z$.

Proof: Part (a) with $z = 2E(S)$ as well as parts (b) and (c) are well-known and proved, e.g., in (Hagerup and Rüb, 1990). In order to show the general form of part (a), assume that S is the number of heads in m independent tosses of a coin with probability p of heads and, without loss of generality, that $z \leq m$. Let $Z_1, \ldots, Z_m$ be independent random variables with range $\{0,1,2\}$ and with $\Pr(Z_i = 1) = p$ and $\Pr(Z_i \in \{1,2\}) = z/(2m)$ $(\geq p)$, for $i = 1, \ldots, m$. Then $S_1 = |\{i : 1 \leq i \leq m \text{ and } Z_i = 1\}|$ is distributed as S, $S_2 = |\{i : 1 \leq i \leq m \text{ and } Z_i \geq 1\}|$ is binomially distributed with $E(S_2) = z/2$, and $S_2 \geq S_1$. But then, by the special case of part (a) already established, $\Pr(S \geq z) = \Pr(S_1 \geq z) \leq \Pr(S_2 \geq z) = \Pr(S_2 \geq 2E(S_2)) \leq e^{-z/6}$. ∎

The following fact is implied by Azuma's inequality (see (Bollobás, 1987) or (McDiarmid, 1989)). Corollary 2.3 expresses the special form of Lemma 2.2(a) that we shall most often use.

Lemma 2.2: Let $m \in \mathbb{N}$, let $Z_1, \ldots, Z_m$ be independent random variables with finite ranges, and let S be an arbitrary real function of $Z_1, \ldots, Z_m$ with $E(S) \geq 0$. Assume that S changes by at most c in response to an arbitrary change in a single Z_i. Then
(a) For every $z \geq 2E(S)$, $\Pr(S \geq z) \leq e^{-z^2/(8c^2m)}$;
(b) $\Pr(S \leq E(S)/2) \leq e^{-(E(S))^2/(8c^2m)}$.

Corollary 2.3: Under the assumptions of Lemma 2.2,

$$\Pr(S \geq \max\{2E(S), 4cm^{1/2+\epsilon}\}) \leq 2^{-m^\epsilon}$$

for all $\epsilon > 0$.

In later applications, we write "by a Chernoff bound" instead of "by Lemma 2.1", and "by a martingale argument" rather than "by Corollary 2.3".

When we state that an algorithm makes at most m random choices, a change in one of which affects some real quantity S by at most c, what we mean is that S is determined somehow by an execution of the algorithm, and that the execution is deterministically given by the input, except that it may also be influenced by at most m *independent* random quantities computed by the algorithm, a change in one of which (with the other random quantities kept fixed) causes S to change by an amount of at most c. A statement

to the effect that a change in a single random choice affects at most a certain number of output variables is to be interpreted in a similar manner.

The algorithms implied by the results listed below are needed as basic subroutines in many places. For all $n \in \mathbb{N}$, the *integer prefix summation problem* of size n is, given n integers $a_1, \ldots, a_n$, to compute the prefix sums $\sum_{j=1}^{i} a_j$, for $i = 1, \ldots, n$. Lemma 2.4 combines many previous results by various authors by giving the optimal prefix summation time for any combination of three independent parameters.

Lemma 2.4 (Hagerup, 1992c): For all given integers $n, m, p \geq 4$, the prefix sums of n integers, each of absolute size at most m, can be computed on a TOLERANT PRAM using

$$O\left(\frac{n}{p} + \frac{\log n}{\log\log p} + \max\left\{\log\frac{\log m}{\log p}, 1\right\}\right)$$

time, p processors and $O(n+p)$ space.

Corollary 2.5: For every fixed $\delta > 0$ and for all given integers $n, \tau \geq 2$, the prefix sums of $(\log n)^{O(1)}$ integers, each of absolute size polynomial in n, can be computed on a TOLERANT PRAM using $O(\tau)$ time, $O(\lceil n^\delta/\tau \rceil)$ processors and $O(n^\delta)$ space.

Lemmas 2.6 and 2.7 provide algorithms for the TOLERANT PRAM for problems that are trivial on certain stronger PRAM variants. Lemma 2.6 is due to Alon and Megiddo, who describe a constant-time algorithm for the more general problem of linear programming in fixed dimension. Specialized to maximum-finding, their algorithm can be viewed as a PRAM implementation of an algorithm for the parallel comparison-tree given by Reischuk (1985).

Lemma 2.6 (Alon and Megiddo, 1990): There is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$, the maximum of n integers can be computed on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Lemma 2.7 (Fich *et al.*, 1988b, Theorem 1): For all given $n, \tau \in \mathbb{N}$, the following problem can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space: Given n bits $x_1, \ldots, x_n$, compute $\max(\{j : 1 \leq j \leq n \text{ and } x_j = 1\} \cup \{0\})$.

When dealing with groups of consecutively numbered processors, the *nearest preceding element* problem defined below formalizes the task of broadcasting information from the lowest-numbered processor in each group to the remaining group members.

Definition: For all $n \in \mathbb{N}$, the *nearest preceding element* problem of size n is, given n bits $x_1, \ldots, x_n$, to compute $y_1, \ldots, y_n$, where $y_i = \max(\{j : 1 \leq j < i \text{ and } x_j = 1\} \cup \{0\})$, for $i = 1, \ldots, n$.

Part (a) of Lemma 2.8 below is due to Berkman and Vishkin (1989) and Ragde (1990), who in fact prove a slightly stronger claim than the one cited here. Part (b) was essentially shown by Raman (1990). His algorithm solves only the less general *ordered chaining* problem, which requires the computation of y_i only if $x_i = 1$, and he considers only the case of constant k, but an extension to the form given here is straightforward.

Lemma 2.8: For all given $n \in \mathbb{N}$, nearest preceding element problems of size n can be solved on a TOLERANT PRAM
(a) in $O(\tau)$ time using $\lceil n/\tau \rceil$ processors and $O(n)$ space, for all given integers $\tau \geq \log^* n$;
(b) in $O(\tau)$ time using $\lceil kn/\tau \rceil$ processors and $O(n)$ space, for all given $k, \tau \in \mathbb{N}$ with $\tau \geq k$, provided that the number of nonzero input bits is $O(n/\log^{(k)} n)$.

Lemma 2.9, finally, states that small integers can be sorted fast with optimal speedup.

Lemma 2.9 (Cole and Vishkin, 1989, Section 3.2): For every fixed $\delta > 0$ and for all given integers $n, m, \tau \geq 4$ with $\tau \geq \log n / \log\log n + m^\delta$, n integers in the range $0\,..\,m$ can be sorted on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space.

3 Scattering

The fundamental intuitive meaning of a *scattering* is that each of a number of objects is placed randomly and independently of other objects in one of a number of cells placed in a row. In this paper we are frequently interested in the resulting *fullness* of the row, i.e., in the ratio of occupied cells to the total number of cells. Since this is clearly a random variable that tends to take on larger values if more objects are scattered, it provides a (very crude) basis for estimating the number of scattered objects. By letting each object participate in the scattering with some suitable probability instead of with probability 1 as above (a *conditional scattering*), we can adjust the "region of sensitivity" of a scattering according to need. A *graduated conditional scattering* (GCS) takes this idea one step further by providing a whole sequence of conditional scatterings, each with a different associated probability of participation, which gives us a way to make more substantial statements about the number of scattered objects. Graduated conditional scatterings were introduced in (Hagerup and Radzik, 1990), although not for the purpose of estimation.

Our analysis of the outcome of a GCS is limited to determining the last scattering in the sequence whose predecessor scatterings all satisfy a certain property. Two properties are relevant to us: The row of a scattering being *full* (all cells are occupied), and the row

being roughly half full. It turns out that testing according to full rows is computationally easier, but leads to less accurate estimates. We now provide the technical details.

Definition: For all $s \in \mathbb{N}$ and $0 \leq p \leq 1$, a *conditional scattering* with probability p and of width s is a random experiment carried out by a set U as follows: Each element $u \in U$, independently of other elements, chooses a random number Z_u with $\Pr(Z_u = 0) = 1 - p$ and $\Pr(Z_u = i) = p/s$, for $i = 1, \ldots, s$. An element $u \in U$ is said to *occupy* the value i if $Z_u = i$, for $i = 1, \ldots, s$, and the *fullness* of the conditional scattering is k/s, where $k = |\{Z_u : u \in U\} \backslash \{0\}| = |\{i : 1 \leq i \leq s \text{ and } i \text{ is occupied by at least one element of } U\}|$. Two distinct elements in U *collide* if they occupy the same (nonzero) value. A scattering of width s is a conditional scattering with probability 1 and of width s.

Note that the value 0 plays a special role in the above definition. An element of U that chooses the random value 0, informally, is one that chooses not to participate in the conditional scattering; p is hence the probability of participating.

Lemma 3.1: Let $m, s \in \mathbb{N}$ and $0 \leq p \leq 1$, and let N be the number of occupied values in a conditional scattering with probability p and of width s carried out by a set of m elements. Then

(a) For all $k \in \{0, \ldots, s\}$, $\Pr(N \leq k) \leq \binom{s}{k} \cdot 2^{mp(k/s-1)}$;

(b) $\Pr(N \leq s/2) \leq 2^{s-mp/2}$;

(c) $\Pr(N < s) \leq s \cdot 2^{-mp/s}$;

(d) For all $z > 0$, $\Pr(N \geq z) \leq (mpe/z)^z$.

Proof: For all $k \in \{0, \ldots, s\}$, the probability that a fixed element occupies a value outside a fixed set $D \subseteq \{1, \ldots, s\}$ of size k is $p(s-k)/s$. Hence the probability that all occupied values belong to D is $(1 - p(s-k)/s)^m \leq e^{-mp(s-k)/s} \leq 2^{mp(k/s-1)}$. Part (a) now follows by observing that D can be chosen in $\binom{s}{k}$ ways. Parts (b) and (c) are special cases of (a), and part (d) is implied by Chernoff bound (c). ∎

Definition: For all $r, s \in \mathbb{N}$, a *graduated conditional scattering* (GCS) with parameters $r \times s$ is a sequence $\mathcal{S} = (S_1, \ldots, S_r)$, where S_i, called the ith *row* of $\mathcal{S}$, is a conditional scattering with probability 2^{-i} and of width s, for $i = 1, \ldots, r$. For $0 \leq f \leq 1$, define the *f-row* of $\mathcal{S}$ as 0 if S_1 has fullness $< f$, and otherwise as the largest integer $i \in \{1, \ldots, r\}$ such that S_j has fullness $\geq f$, for $j = 1, \ldots, i$.

Lemma 3.2: For every $z > 0$, $\min\{1, \sum_{i=0}^{\infty} 2^{-2^i z}\} \leq 2^{1-z}$.

Proof: $\sum_{i=0}^{\infty} 2^{-2^i z} \leq \sum_{i=0}^{\infty} 2^{-z(i+1)} = \frac{2^{-z}}{1-2^{-z}}$. If $z \leq 1$, then $1 \leq 2^{1-z}$. On the other hand, if $z > 1$, then $\frac{2^{-z}}{1-2^{-z}} \leq \frac{2^{-z}}{1/2} = 2^{1-z}$. ∎

Lemma 3.3: Let $m, r, s \in \mathbb{N}$ and $a > 0$ and let L be the 1-row of a GCS of a set of m elements with parameters $r \times s$. Then if $M = 2^L s$,

(a) $\Pr(L = r) \leq (2^{-r} em/s)^s$;

(b) $\Pr(M > \max\{s, am\}) \leq (2e/a)^s$;

(c) $\Pr(L < r \text{ and } m > aM) \leq s \cdot 2^{1-a/2}$;

(d) $\Pr(L < r) \leq s \cdot 2^{1-m/(2s \cdot 2^r)}$.

Proof: (a) follows immediately from Lemma 3.1(d). If $L > 0$, row L is full. Hence by Lemma 3.1(d), for every $l \geq 0$,

$$\Pr(L > l) \leq \sum_{i=\lceil l \rceil}^{\infty} \left(\frac{2^{-i} em}{s}\right)^s \leq \left(\frac{2^{1-l} em}{s}\right)^s. \tag{1}$$

Likewise, if $L < r$, row $L + 1$ is not full, and Lemmas 3.1(c) and 3.2 imply that for every $l \leq r$,

$$\begin{aligned} \Pr(L < l) &\leq \min\{1, \sum_{i=-\infty}^{\lfloor l \rfloor} s \cdot 2^{-m \cdot 2^{-i-1}/s}\} \\ &\leq s \cdot \min\{1, \sum_{i=0}^{\infty} 2^{-2^i m \cdot 2^{-\lfloor l \rfloor - 1}/s}\} \leq s \cdot 2^{1-m \cdot 2^{-l-1}/s}. \end{aligned} \tag{2}$$

To show (b), apply (1) with $l = \log(\max\{s, am\}/s) \geq 0$. This yields

$$\Pr(M > \max\{s, am\}) = \Pr(L > l) \leq \left(\frac{2^{1-l} em}{s}\right)^s \leq (2e/a)^s.$$

To show (c), apply (2) with $l = \min\{\log(m/(as)), r\} \leq r$ to obtain

$$\Pr(L < r \text{ and } m > aM) = \Pr(L < l) \leq s \cdot 2^{1-m \cdot 2^{-l-1}/s} \leq s \cdot 2^{1-a/2}.$$

(d) follows from (c) by taking $a = m/(s \cdot 2^r)$. ∎

The following lemmas investigate the computational aspects of graduated conditional scattering. Note that part (b) of Lemma 3.4 and Lemma 3.5 are needed only for the proof of Theorem 11.6, which is not part of the main development.

Lemma 3.4: Let $r, s \in \mathbb{N}$ be given and suppose that a processor is associated with each element of some set U. Then the 1-row of a GCS of U with parameters $r \times s$ can be determined in constant time

(a) on a TOLERANT PRAM using rs additional processors and $O(rs)$ space;

(b) on an ARBITRARY PRAM using one additional processor and $O(rs)$ space.

Proof: (a) Let A be an $r \times s$ array. Each processor associated with an element of U chooses a random cell $A[I, J]$, where $\Pr(I = i, J = j) = 2^{-i}/s$, for $i = 1, \ldots, r$ and

$j = 1, \ldots, s$, and distinct processors act independently (with whatever probability is left, processors do nothing). In other words, each processor chooses a random row, row i being chosen with probability 2^{-i}, and then picks a cell at random from the chosen row. Next the value 1 is stored in each cell of A that was chosen by at least one processor associated with an element of U, and the value 0 is stored in each of the remaining cells of A. On the ARBITRARY PRAM, this would be trivial; on the TOLERANT PRAM, we proceed as follows: Use the rs additional processors to associate one processor, called a *guard*, with each cell of A and let each guard begin by storing the value 1 in its associated cell. Subsequently let each processor associated with an element of U attempt to write (an arbitrary value) to its chosen cell in A; simultaneously, each guard attempts to change the value stored in its associated cell from 1 to 0. By definition of the TOLERANT PRAM, this will succeed if and only if the cell was not chosen by any processor associated with an element of U, i.e., afterwards the cell contains the desired value. This technique, which we call *guarded writing*, was first used by Grolmusz and Ragde (1987) and appears to be a fundamental technique for programming the TOLERANT PRAM. Once the cells of A have been marked with zeros and ones as described above, it is easy to use the algorithm of Lemma 2.7 to compute the conjunction w_i of $A[i,1], \ldots, A[i,s]$, for $i = 1, \ldots, r$. Finally Lemma 2.7 is used again to determine the smallest $i \in \{1, \ldots, r\}$, if any, with $w_i = 0$.

(b) Our algorithm centers around a solution to a variant of a problem known as the *leftmost prisoner problem*. The leftmost prisoner problem, introduced by Fich *et al.* (1988a), is unusual in that an instance of the problem is not given by an input in a traditional sense; rather, the instance is defined by the processors available for its solution themselves. In more detail, an instance of the leftmost prisoner problem of size n is given by a set of processors numbered $1, \ldots, n$, each of which is either *active* or *inactive*. At least one processor is active, and the task is to compute the smallest processor number of an active processor, whereby inactive processors do not participate in the computation in any way. The latter restriction is essential — without it, the problem could be solved very easily using the algorithm of Lemma 2.7. The complexity of leftmost prisoner problems of size $n \geq 4$ on the ARBITRARY PRAM was shown to be $\Theta(\log \log n)$ by Chlebus *et al.* (1988) and Grolmusz (1991). Here we are interested in a variant of the problem called the *leftmost empty prison cell* problem. The setup is exactly as for the leftmost prisoner problem, but we want to compute the smallest processor number of an *inactive* processor, or an indication of the fact that all processors are active. In Lemma 3.5 below we show that leftmost empty prison cell problems can be solved in constant time on an ARBITRARY PRAM. Here we will take this result for granted and describe its application to graduated conditional scattering.

We use an $r \times s$ array A and begin by letting each processor associated with an element of U choose a random cell in A exactly as in the proof of part (a). For each row

of A, we now wish to associate a processor with the row if and only if each cell in the row was chosen by at least one processor. To this end we view each row as defining an instance of the leftmost empty prison cell problem. For $j = 1, \ldots, s$, each processor having chosen the jth cell in the row temporarily adopts j as its processor number and represents an active processor in the sense of the leftmost empty prison cell problem; the fact that several processors may have chosen the same cell in A leads to no problem, since they will all carry out the same computation. For $j = 1, \ldots, s$, if the jth cell in the row under consideration was not chosen by any processor, we associate with it a fictitious inactive processor with processor number j. We can now use an algorithm for the leftmost empty prison cell problem to determine whether any processor is inactive, i.e., whether some cell in the row was not chosen by any processor. If and only if all cells were chosen, we associate one (or all) of the processors having chosen a cell in the row with the row; note that in the special case in which no processor chose a cell in the row, no processor will be associated with the row, as desired.

We now view the processors associated with some of the rows of A as defining an instance of the leftmost empty prison cell problem in a similar way and observe that solving this problem produces the desired result. The special case in which no row of A has an associated processor can be handled by the single processor dedicated to the GCS. ∎

Lemma 3.5: Leftmost empty prison cell problems can be solved in constant time on an ARBITRARY PRAM.

Proof: In the algorithm described below we shall frequently want to mark cells that we may not have been able to initialize. This is problematic, because an "undefined" value present from the outset in a cell that is not marked may happen to coincide with the value that would have been written there had the cell been marked. We avoid this difficulty by means of what we call *dynamic marking*: A processor marks a cell by first writing 0 and subsequently 1 (say) to the cell. Any processor wishing to know whether the cell is marked reads its contents both between the two writes and after the second write and deems the cell marked if and only if it observes a change from 0 to 1. Although, in this scheme, the writing and reading of a mark takes place in an interleaved fashion, in the description below we will pretend that the writing precedes the reading.

We can assume that n is a power of 2 and that at least one processor is inactive, since both requirements can be satisfied by adding a suitable number of fictitious inactive processors (an answer larger than the number of original processors should then be interpreted as an indication that the original processors are all active). Assume that the processors are ordered linearly from left to right by increasing processor numbers. If the leftmost processor is inactive, all active processors can discover this fact through dynamic

marking and output the correct answer (namely 1); assume therefore that this is not the case.

Starting from the left, divide the processors into groups of sizes $1, 1, 2, 4, \ldots, n/2$ and call a group *complete* if all processors in the group are active, and *incomplete* otherwise. A first part of the computation serves to let each active processor know whether its group is complete. This can be done as follows: Using dynamic marking, each processor determines whether its left neighbor is active, whereby the left neighbor of the leftmost processor in a group is taken to be the rightmost processor in the group (i.e., the ordering within each group is cyclic). Then a cell associated with each group is initialized to 1 by all active processors in the group and subsequently set to 0 by all active processors in the group whose left neighbors are inactive. It is easy to see that if at least one processor in the group is active, the value of the cell remains 1 if and only if the group is complete.

The processors in incomplete groups do not participate in the remaining computation. The processors in a complete group of size m, on the other hand, use dynamic marking and the algorithm of Lemma 2.7 to solve the leftmost empty prison cell problem defined by the $m' = \min\{4m, n\}$ leftmost processors and output the result if and only if at least one of the m' leftmost processors is inactive.

Any output produced by a complete group clearly is the desired answer. On the other hand, if the processor number of the leftmost inactive processor P is k, the group to the left of P's group exists (by assumption) and is complete and of size at least $k/4$, so that an output will be produced at least by this complete group. ∎

Whereas graduated conditional scatterings were introduced for the purpose of estimation, we also employ a different kind of scattering, called *v-scattering* or (with implicit v) *multi-scattering*, for the task of placing elements in distinct cells of a destination array. Because of the more operational use, the definition below is formulated in algorithmic terms.

Definition: For all $v, s \in \mathbb{N}$, to *v-scatter* a set U over an array A of s cells is to execute the following algorithm: If $v > s$, do nothing. Otherwise divide A into v disjoint subarrays of size at least $\lfloor s/v \rfloor$ each and create v copies of each element in U. Then let the set of ith copies use the ith subarray to carry out a scattering of width $\lfloor s/v \rfloor$ and identify the set of noncolliding copies, for $i = 1, \ldots, v$. An element in U is said to be *successful* if it has at least one noncolliding copy; in particular, if $v > s$ we consider all elements of U to be unsuccessful. For each successful element $u \in U$, let i and j be, respectively, the number of a noncolliding copy of u and the value occupied by that copy, and place u in the jth cell of the ith subarray of A; note that this never places distinct elements in the same cell. The *density* of the v-scattering is the quantity $|U|v/s$.

Using Lemma 2.7, it is easy to see that if each element of a set U has an associated group of v processors, then U can be v-scattered over an arbitrary array A in constant time. Lemma 3.6 quantifies the efficiency of this procedure as a means of placing the elements of U in A. The proof of Lemma 3.6 applies a martingale argument in a situation where, a priori, the number of random choices made is too large for such an application. We overcome this difficulty by *fixing* most of these choices in advance, i.e., by considering a restricted probability space. If we can show that some event occurs with probability at most q independently of how the random choices are fixed, then the event occurs with probability at most q even in the actual experiment, where random choices in fact are not fixed. The same principle will be used again later.

Lemma 3.6: Let $m, s, v \in \mathbb{N}$, denote by D the set of unsuccessful elements in a v-scattering of a set U of size m over an array of size s and let $p = mv/s$ be the density of the v-scattering. Then

(a) For all $u \in U$, $\Pr(u \in D) \leq p^v$;

(b) For every fixed subset R of U and for all $z \geq 2|R|p^v$,

$$\Pr(|R \cap D| \geq z) \leq e^{-z^2/(32|R|v)}.$$

Proof: For part (a), it suffices to show for $v \leq s$ that if U carries out a scattering of width $\lfloor s/v \rfloor$, then the probability that a fixed element $u \in U$ collides is at most mv/s. If $m \geq s/v$, this is certainly true. Otherwise the probability under consideration is at most

$$\frac{m-1}{\lfloor s/v \rfloor} \leq \frac{m-1}{s/v-1} \leq \frac{mv}{s}.$$

For part (b), let R be a fixed subset of U with $|R| = r$ and consider the random choices made by copies of elements not in R to be fixed in an arbitrary way. As in the proof of part (a), a fixed element in R is unsuccessful with probability at most p^v, so that $E(|R \cap D|) \leq rp^v$. A moment's thought reveals that a change in a single random choice (namely that of a single copy) can change $|R \cap D|$ by at most 2. Since there are altogether rv such choices, an application of Lemma 2.2(a) now shows that for $z \geq 2rp^v$, $|R \cap D| \geq z$ with probability at most $e^{-z^2/(32rv)}$. ∎

Section 6 extends the basic multi-scattering algorithm above to *colored multi-scattering*, where the set U to be multi-scattered is partitioned into *color classes* $U_1, \ldots, U_m$ and U_i is multi-scattered over a separate array A_i, for $i = 1, \ldots, m$. It is easy to see that if the density of the multi-scattering of U_i over A_i is bounded by p, for $i = 1, \ldots, m$, then the assertions of Lemma 3.6 carry over to the more general situation. This agrees well with intuition, since the coloring of elements only helps the multi-scattering algorithm to distribute copies evenly and avoid collisions.

4 Compaction

This section studies the *compaction* problem, which occurs as a base case of the more general interval allocation problem considered in Section 7. Roughly speaking, the compaction problem is to move a number of objects, scattered over a large *source* array, to distinct cells in a smaller *destination* array, possibly with a small number of objects, said to be *unlucky*, left behind in the source array. Our formalization of the problem abstracts away the identities of the objects to be moved and simply takes the input to be a sequence $x_1, \ldots, x_n$ of n bits, where n is the size of the source array; $x_j = 1$ signifies the presence and $x_j = 0$ the absence of an object in the jth cell of the source array, for $j = 1, \ldots, n$. The output takes the form of n nonnegative integers $y_1, \ldots, y_n$. If $x_j = 1$, the object in the jth cell of the source array can be moved to the y_jth cell of the destination array, for $j = 1, \ldots, n$, except that by convention $y_j = 0$ signals that the corresponding object is unlucky. If $x_j = 0$, the value of y_j is immaterial and may as well be set to zero (condition (1) below).

Definition: For all $n \in I\!N$ and $s \geq 0$, an *incomplete placement* with *bound* s for n bits $x_1, \ldots, x_n$ is a sequence $y_1, \ldots, y_n$ of n nonnegative integers such that
(1) For $j = 1, \ldots, n$, if $x_j = 0$, then $y_j = 0$;
(2) For $1 \leq i < j \leq n$, if $y_i \neq 0$, then $y_i \neq y_j$;
(3) $\max\{y_j : 1 \leq j \leq n\} \leq s$.
The set $\{j : 1 \leq j \leq n,\ x_j = 1 \text{ and } y_j = 0\}$ is called the *residue set* of the incomplete placement. If the residue set is empty, the placement is said to be *complete*.

Condition (2) in the above definition expresses that distinct objects may not be placed in the same destination cell, and condition (3) states that size $\lfloor s \rfloor$ suffices for the destination array. The residue set is the set of indices of objects that are not placed in the destination array.

Most of the computational problems introduced in this paper take as (part of) their input a sequence $x_1, \ldots, x_n$. Although, formally, $x_1, \ldots, x_n$ are integers (sometimes restricted further to be single bits), informally they represent objects of additional internal structure. In particular, if $i \neq j$, the objects represented by x_i and x_j are distinct, even if it happens that $x_i = x_j$. This is mirrored closely by what happens in our algorithms for solving such problems. They typically begin by transforming the input $x_1, \ldots, x_n$ to n records $X_1, \ldots, X_n$ that are subsequently manipulated instead of $x_1, \ldots, x_n$. For $j = 1, \ldots, n$, fields in the record X_j contain the integer x_j, called the *value* of X_j, the integer j, called its *index*, as well as any other attributes that the algorithms may need. Usually we shall not describe our algorithms at the level of such programming detail; note, however, that the symbols $X_1, \ldots, X_n$ will be used in

the above sense throughout the paper. When we speak of the jth *input element*, for $j = 1, \ldots, n$, we usually mean the record X_j, and $\mathcal{X} = \{X_1, \ldots, X_n\}$ is called the *input set*. In particular, for $i \neq j$, the ith and jth input elements are distinct.

For reasons of convenience, we will occasionally state that some algorithm is applied to a subset of $\mathcal{X}$. What we really mean in such a case is that the algorithm is applied to the corresponding subsequence of $x_1, \ldots, x_n$, usually permuted in some way and augmented with a number of suitable dummy elements, neither of which affects the problem in an essential way. Furthermore, we assume that enough additional information is kept to interpret the output of the algorithm in terms of the original sequence $x_1, \ldots, x_n$.

In the context of compaction, an *active element* is an input element of nonzero value that has not yet been placed in the destination array. Once successfully placed, we say that it has been *deactivated* or that it has become *inactive*. The incomplete compaction problem with parameters $d_1 \underset{s}{\longrightarrow} d_2$, defined below, is, given at most d_1 active elements, to move all except at most d_2 of these to a destination array of size at most s.

Definition: For all $n \in \mathbb{N}$ and $d_1, d_2, s \geq 0$, the *incomplete compaction* problem of size n and with parameters $d_1 \underset{s}{\longrightarrow} d_2$ is the following: Given n bits $x_1, \ldots, x_n$ with $\sum_{j=1}^{n} x_j \leq d_1$, compute an incomplete placement for $x_1, \ldots, x_n$ with bound s whose residue set is of size at most d_2. If $d_2 = 0$, we speak of *complete* rather than incomplete compaction.

Lemma 4.1: For all given $n, d \in \mathbb{N}$, complete compaction problems of size n and with parameters $d \underset{d^4}{\longrightarrow} 0$ can be solved on a (deterministic) TOLERANT PRAM using constant time, n processors and $O(n)$ space.

Proof: The result was proved by Ragde for the stronger ARBITRARY PRAM (Ragde, 1990, Theorem 1). Using Lemma 2.7, it is easy to translate Ragde's algorithm to the TOLERANT PRAM. ∎

Lemma 4.1 works in constant time, but places the active elements in an array with many more cells than the number of active elements. The far more important complete linear compaction problem, with inessential differences also known as the *linear approximate compaction* or *LAC* problem (Matias and Vishkin, 1991), requires the size of the destination array to be within a constant factor of the bound on the number of active elements.

Definition: For all $n, d \in \mathbb{N}$, the *complete linear compaction* problem of size n and with *limit* d is the complete compaction problem of size n and with parameters $d \underset{O(d)}{\longrightarrow} 0$.

We next show that complete linear compaction problems can be solved in constant time using a superlinear number of processors. Our algorithm first multi-scatters the active elements over an auxiliary array in order to distribute them approximately evenly. The auxiliary array is then divided into segments of a fixed size chosen so large as to make it unlikely that any segment contains more than c times the average number of active elements, for a suitable constant $c > 1$. If a destination array c times larger than the number of active elements is now divided evenly among the segments, all that remains is to distribute the destination cells allocated to each segment within the segment, i.e., among the active elements stored there. This can be done via brute-force prefix summation (Corollary 2.5) following a "loose" compaction of the active elements within the segment (Lemma 4.1). The details follow.

Lemma 4.2: For every fixed $\delta > 0$, there is a constant $\epsilon > 0$ such that for all given $n, d \in \mathbb{N}$, complete linear compaction problems of size n and with limit d can be solved on a TOLERANT PRAM using constant time, $O(n^{1+\delta})$ processors and $O(n^{1+\delta})$ space with probability at least $1 - 2^{-n^{\epsilon}}$ (Las Vegas).

Proof: Let $v = \lceil 8/\delta \rceil$. Without loss of generality we can assume that δ is rational (so that we can easily compute with δ, cf. Section 2) and that $n \geq 4$ and $d \leq n/v$ (since otherwise the compaction problem is trivial). It suffices to describe an algorithm that uses constant time, $O(n^{1+\delta/2})$ processors and $O(n^{1+\delta/2})$ space and that fails with probability at most $1/2$, since we can execute such an algorithm $n^{\Omega(1)}$ times in parallel and select as our output the outcome of any successful execution. If $d < \log n$, the problem can be solved by first using the algorithm of Lemma 4.1 to move the active elements to an array of $(\log n)^{O(1)}$ cells and subsequently compacting them exactly, i.e., numbering them consecutively, using prefix summation (Corollary 2.5). Assume hence that $d \geq \log n$.

Let A be an array of size vs, where s is chosen as a multiple of $r = \lceil d/\lfloor \log n \rfloor \rceil$ with $s \geq n^{1+\delta/4}$, but $s = O(n^{1+\delta/4})$. Then v-scatter the active elements in the source array over A; by Lemma 3.6(a), the probability that some element cannot be placed in A is at most $n \cdot (d/s)^v \leq n \cdot (n^{-\delta/4})^{8/\delta} = 1/n$.

Divide A into vr disjoint *segments* of size s/r each. The number S of active elements placed in a fixed segment in the above v-scattering is clearly bounded by the number S' of copies of elements choosing a cell in the segment in the v-scattering (S may be smaller than S' because copies choosing a cell in the segment can collide, and still smaller because elements with a noncolliding copy placed in the segment may be moved to the position of another noncolliding copy). Since the v-scattering partitions A into v subarrays of r segments each and at most d copies choose cells in the subarray containing the segment under consideration, S' is binomially distributed with expected value $d/r \leq \log n$; Chernoff bound (a) therefore implies that $S' \geq 12 \log n$ with probability

at most $e^{-2\log n} \leq n^{-2}$. It follows that except with probability at most $n \cdot n^{-2} = 1/n$, no segment contains more than $12\log n$ active elements.

Since we have $n^{\Omega(1)}$ processors per segment, we can now use the algorithms of Lemma 4.1 and Corollary 2.5 as in the beginning of the proof to attempt to place the active elements in each segment in an array of size $12\lceil\log n\rceil$ (the attempt fails only in the unlikely event that some segment contains more than $12\lceil\log n\rceil$ elements). Assigning to each segment a subarray of size $12\lceil\log n\rceil$ of a common destination array and moving each active element to the appropriate cell in the destination array completes the compaction. The total size of the destination array is $12vr\lceil\log n\rceil = O(d)$, and the probability that the algorithm fails is at most $2/n \leq 1/2$. ∎

Corollary 4.3: For every fixed $\delta > 0$ there is a constant $\epsilon > 0$ such that for all given $n, d, \tau \in \mathbb{N}$, complete linear compaction problems of size n and with limit $d = O(n^{1-\delta})$ can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau\rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^{\epsilon}}$ (Las Vegas).

Proof: Without loss of generality assume that δ is rational and that $\delta \leq 1$. Since the problem is easily solved using standard prefix summation (Lemma 2.4) if $\tau = n^{\Omega(1)}$, we can further assume the availability of at least $n^{1-\delta/4}$ processors. Divide the input set $\mathcal{X}$ into $O(n^{1-\delta/2})$ *clusters* of $O(n^{\delta/2})$ input elements each and call a cluster *nonempty* if it contains at least one active element. There are obviously at most d nonempty clusters, so the algorithm of Lemma 4.2 can be used to place the indices of these in an array of size $O(d)$. This implicitly places all active elements in an array of size $O(d \cdot n^{\delta/2}) = O(n^{1-\delta/2})$, after which the compaction can be completed via a second application of the algorithm of Lemma 4.2. ∎

Matias and Vishkin (1991) showed that complete linear compaction problems of size n and with limit d can be solved in $O(\log^* d)$ time with n processors by successively solving $O(\log^* d)$ incomplete compaction problems. The basic idea is that the gradual deactivation of elements frees resources that can be used to speed up the rate of deactivation, thus leading to the fast convergence of the algorithm. In more detail, Matias and Vishkin show that if the number of active elements has already dropped to d/v^c, for a suitable constant $c > 0$ and for some $v \in \mathbb{N}$, then in constant time it can be decreased further to $d/2^{cv}$. The algorithm of (Matias and Vishkin, 1991) realizing this claim is reasonably complicated and relies crucially on Lemma 4.1. We give a trivial algorithm for the same task whose use of Lemma 4.1 is inessential and easy to avoid, and whose complete analysis is much simpler than what would be required for the algorithm of Matias and Vishkin. As concerns the claim of simplicity, observe below that the appeal to Corollary 4.3 is needed only to deal with a special case that was not even considered in (Matias and Vishkin, 1991).

Our algorithm for incomplete compaction inputs at most d/v^3 active elements stored in a source array of size n and places all except $d/2^{3v}$ of these in a destination array of size $O(d/v)$. The basic idea is to $5v$-scatter the active elements over an array of size $10d/v^2$. Since, by assumption, the density of this $5v$-scattering is at most $1/2$, a fixed active element remains active with probability at most 2^{-5v} (Lemma 3.6(a)), which allows us to conclude that with high probability the size of the residue set will be bounded by $d/2^{3v}$. The only problem with this approach is that we do not know how to allocate the $5v$ processors per active element necessary to carry out the $5v$-scattering in constant time. Similarly as in the proof of Corollary 4.3, we therefore divide the input set $\mathcal{X}$ into *clusters* of size v each and $5v$-scatter not the active elements themselves, but instead (the indices of) the *nonempty* clusters, i.e., those clusters that contain at least one active element. Since the number of nonempty clusters is obviously bounded by the number of active elements, the density of the modified $5v$-scattering is also at most $1/2$. Furthermore, the allocation of $5v$ processors to each cluster is trivial, and placing the nonempty clusters in an array of size $O(d/v^2)$ implicitly places the active elements in an array of size $O(d/v)$.

Lemma 4.4: There is a constant $\epsilon > 0$ such that for all given $n, d, v, \tau \in \mathbb{N}$ with $d \leq n$, incomplete compaction problems of size n and with parameters

$$\frac{d}{v^3} \xrightarrow[O(d/v)]{} \frac{d}{2^{3v}}$$

can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo).

Proof: We give the proof for $\tau = 1$, leaving the easy extension to general values of τ to the reader (informally, the observation needed is that executing a multi-scattering over several steps rather than in one step can only cause more elements to be successful). We can obviously assume that $v^3 \leq d \leq n$ (otherwise we start with no active elements).

Consider the following algorithm:

Step 1: Divide $\mathcal{X}$ into $l = \lceil n/v \rceil$ *clusters* $\mathcal{X}_1, \ldots, \mathcal{X}_l$ of at most v input elements each and use the algorithm of Lemma 2.7 to compute a bit representation of the set $I = \{i : 1 \leq i \leq l \text{ and } \mathcal{X}_i \text{ contains at least one active element}\}$.

Step 2: Associate $5v$ processors with each element of I and $5v$-scatter I over an array of size $\lceil 10d/v^2 \rceil$; let $I' \subseteq I$ denote the set of unsuccessful indices. Use the outcome of the $5v$-scattering to place all active elements in $\bigcup_{i \in I \setminus I'} \mathcal{X}_i$ in an array of size $v\lceil 10d/v^2 \rceil = O(d/v)$.

The algorithm clearly runs on a TOLERANT PRAM within the desired resource bounds. A fixed active element remains active exactly if the index of its cluster is

unsuccessful in the $5v$-scattering in Step 2. By Lemma 3.6(b), the number of such unsuccessful cluster indices is bounded by $\max\{2d/2^{5v}, n^{5/9}\}$, except with probability at most $e^{-\zeta}$, where $\zeta = (n^{5/9})^2/(32(d/v^3)\cdot 5v) = \Omega(n^{1/9})$. With high probability the number of active elements therefore decreases to at most $v\cdot\max\{d/2^{4v}, n^{5/9}\} \le \max\{d/2^{3v}, n^{8/9}\}$. If this is more than $d/2^{3v}$, at most $n^{8/9}$ elements remain active, and these can be deactivated via an application of the algorithm of Corollary 4.3. ∎

Corollary 4.5: There is a constant $\epsilon > 0$ such that for all given $n, d \in \mathbb{N}$, complete linear compaction problems of size n and with limit d can be solved on a TOLERANT PRAM using $O(\log^* d)$ time, n processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Assume that $d \le n$ and apply the algorithm of Lemma 4.4 at most $\log^* d$ times in successive *stages*, starting with $v = 1$. Each stage after the first attempts to place the unlucky elements of the previous stage in a new but smaller array. Schematically,

$$\frac{d}{1^3} \xrightarrow[O(d/1)]{} \frac{d}{2^3} \xrightarrow[O(d/2)]{} \frac{d}{4^3} \xrightarrow[O(d/4)]{} \frac{d}{16^3} \xrightarrow[O(d/16)]{} \frac{d}{(2^{16})^3} \longrightarrow \cdots \longrightarrow 0.$$

The total size of the destination arrays is

$$O\Big(d\Big(1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{16} + \cdots\Big)\Big) = O(d). \quad ∎$$

As mentioned above, a weaker form of Corollary 4.5 was first proved by Matias and Vishkin (1991), who also noted that it has applications to processor scheduling as per Brent's principle. We next describe an improved algorithm that achieves optimal speedup. A similar result was derived in a somewhat different way by Goodrich (1991).

Theorem 4.6: There is a constant $\epsilon > 0$ such that for all given $n, d, \tau \in \mathbb{N}$ with $\tau \ge \log^* d$, complete linear compaction problems of size n and with limit d can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Assume that $\tau \le (\log n)/32$, since otherwise the compaction can be carried out using prefix summation (Lemma 2.4), and that $d \le n$. We describe a preprocessing stage that reduces the problem size from n to $O(n/\tau)$. Divide $\mathcal{X}$ into $\lceil n/\tau \rceil$ *clusters* of at most τ input elements each and associate a processor with each cluster. Using a global array A of size $8d$, the processors now execute 2τ *rounds*. In each round, each processor chooses an active element in its cluster, if any are left, and attempts to place the chosen element in a random cell of A. If the cell is not already occupied and there is no collision, the element is placed and becomes inactive. It is easy to see that each such trial fails

with probability at most $1/8$, even conditionally on any pattern of failures in previous rounds. As a consequence, all of τ fixed trials by a fixed processor fail with probability at most $(1/8)^\tau$ (if the processor runs out of active elements, let it subsequently execute dummy trials that always succeed). But if a fixed processor has any active elements left after 2τ rounds (call such a processor *busy*), at least τ of its trials must have failed, which, by the above, happens with probability at most $\binom{2\tau}{\tau}(1/8)^\tau \leq 2^{2\tau} \cdot 2^{-3\tau} = 2^{-\tau}$. The expected number of busy processors therefore is $O(n/2^\tau)$. Our intent is to use a martingale argument to show that with high probability, the actual number of busy processors is $O(n/\tau^2)$, which requires us to bound the effect on the number of busy processors of a change in a single random choice. A change in a single random choice here is the choice by some processor of a different cell in A in some round. Say that a processor P is *affected* (by the change under consideration) in a given round if the change influences the success of P's trial in the given round or in some earlier round. At most two processors are affected in the first round, and it is easy to see that the number of affected processors at most triples from one round to the next — an affected processor can "affect" at most two other processors in each later round. Therefore the total number of affected processors after 2τ rounds is at most $3^{2\tau} \leq 2^{4\tau} \leq n^{1/8}$; this is an upper bound on the change in the number of busy processors caused by a change in a single random choice. Since the algorithm makes a total of $O(n)$ random choices, a martingale argument now shows that with high probability, the actual number of busy processors is $O(n/2^\tau + n^{1/8} \cdot n^{5/8}) = O(n/\tau^2)$. But then the algorithm of Corollary 4.5 can be used to place (the processor numbers of) the busy processors in an array of size $O(n/\tau^2)$. This implicitly places the remaining active elements in an array of size $O(n/\tau)$, and the compaction can be completed via another application of the algorithm of Corollary 4.5. ∎

5 Fine-Profiling

The present paper studies several different kinds of profiling problems. In general terms, the task is, given an array containing occurrences of several different values, to estimate the *multiplicity* of each value, i.e., the number of occurrences of that value. We now introduce convenient notation and terminology that will be used throughout the remainder of the paper. In the context of an input consisting of n integers $x_1, \ldots, x_n$ in the range $0..m$, for $n, m \in \mathbb{N}$, we take $\mathcal{B}_i = \{X_j : 1 \leq j \leq n \text{ and } x_j = i\}$ and $b_i = |\mathcal{B}_i|$, for $i = 1, \ldots, m$. For $i = 1, \ldots, m$, the integer i will also be called a *color*, $\mathcal{B}_i$ is a *color class*, and b_i is called the *multiplicity* of the color i. For $1 \leq i < j \leq m$, we consider $\mathcal{B}_i$ and $\mathcal{B}_j$ to be distinct even if they happen to contain the same elements (this is possible only if $\mathcal{B}_i = \mathcal{B}_j = \emptyset$). When a color class $\mathcal{B}_i$ is manipulated as a single object by

some algorithm, it is represented by its index i. Note that elements of value 0 are not considered to belong to any color class. They are just "dummy elements" that represent the absence of a true element. Whenever we have dealt with certain color classes in some algorithm, we can "remove" the elements of these color classes by setting their values to 0, which allows us to focus on the remaining color classes. This will be used on several occasions. Finally let $\mathcal{B} = \bigcup_{i=1}^{m} \mathcal{B}_i$ be the set of input elements with nonzero values.

Given n integers $x_1, \ldots, x_n$ in the range $0..m$, for $n, m \in \mathbb{N}$, an *m-color profile* for $x_1, \ldots, x_n$ is a sequence $\widehat{b}_1, \ldots, \widehat{b}_m$ of m nonnegative integers, the idea being that $\widehat{b}_i$ is an estimate of b_i, for $i = 1, \ldots, m$. A *fine-profile*, defined below, provides estimates that are correct up to a constant factor; for reasons of convenience we also require each estimate to be no smaller than the true multiplicity. For our purposes, having such estimates usually is as good as knowing the exact multiplicities.

Definition: Let $n, m \in \mathbb{N}$ and let $x_1, \ldots, x_n$ be n integers in the range $0..m$. For $i = 1, \ldots, m$, take $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$. An *$m$-color fine-profile* for $x_1, \ldots, x_n$ is a sequence $\widehat{b}_1, \ldots, \widehat{b}_m$ of m nonnegative integers such that $b_i \leq \widehat{b}_i \leq K b_i$, for $i = 1, \ldots, m$ and for some constant $K \geq 1$. If additionally $\widehat{b}_1, \ldots, \widehat{b}_m$ are independent, the sequence $\widehat{b}_1, \ldots, \widehat{b}_m$ is called a *strong fine-profile* for $x_1, \ldots, x_n$. The m-color (strong) fine-profiling problem of size n is, given n and m, to compute an m-color fine-profile (composed of independent estimates) for n given integers in the range $0..m$.

Note that the quantity K in the above definition is a "true constant" (such as 10). In fact, we could fix K at a particular value that can be deduced from the proof of Theorem 5.3. We define a *linear overestimate* for a quantity b as a quantity $\widehat{b}$ with $b \leq \widehat{b} \leq Kb$, for some $K \geq 1$ that is a true constant in this sense. An m-color fine-profile for $x_1, \ldots, x_n$ may therefore also be characterized as a sequence of linear overestimates for $b_1, \ldots, b_m$.

A statement quite similar to Lemma 5.1 below can be derived by combining results of (Stockmeyer, 1983) and (Ajtai and Ben-Or, 1984) with the obvious simulation of unbounded fan-in circuits by CRCW PRAMs. We give a somewhat different proof, which in the context of PRAMs seems more direct.

Lemma 5.1: For every fixed $\delta > 0$ there is a constant $\epsilon > 0$ such that for all given $n \in \mathbb{N}$, the following problem can be solved on a TOLERANT PRAM using constant time, $O(n^\delta)$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo): Given n bits $x_1, \ldots, x_n$, compute a bit y such that

(1) $\sum_{j=1}^{n} x_j \geq n/2 \;\Rightarrow\; y = 1$;

(2) $\sum_{j=1}^{n} x_j \leq n/8 \;\Rightarrow\; y = 0$.

Proof: The idea of the proof, which the reader may appreciate better after the first reading, is to "amplify" a constant-factor difference to a "polynomial" difference, which can then easily be detected using Ragde's lemma (Lemma 4.1).

Assume that δ is rational, that $\delta \leq 1$ and that $n \geq 16$, take $h = 4\lfloor n^{\delta/2} \rfloor \leq n$ and let $t = 32\lceil \log n \rceil$. Begin by determining the number of ones in each of h random samples of t input bits each, i.e., choose ht independent random numbers $z_{1,1}, \ldots, z_{1,t}, z_{2,1}, \ldots, z_{2,t}, \ldots, z_{h,1}, \ldots, z_{h,t}$ from the uniform distribution over $\{1, \ldots, n\}$ and use the algorithm of Corollary 2.5 to compute $S_i = \sum_{j=1}^{t} x_{z_{i,j}}$, for $i = 1, \ldots, h$. The random variables $S_1, \ldots, S_h$ are independent and binomially distributed with expected value tb/n, where $b = \sum_{j=1}^{n} x_j$. Hence by Lemma 2.1, the following holds for $i = 1, \ldots, h$: If $b \geq n/2$, then $\Pr(S_i \leq 8\lceil \log n \rceil) \leq n^{-1}$, while if $b \leq n/8$, then $\Pr(S_i > 8\lceil \log n \rceil) \leq n^{-1}$. For $i = 1, \ldots, h$, take $A[i] = 1$ if $S_i > 8\lceil \log n \rceil$, and let $A[i] = 0$ otherwise. The remaining problem is, assuming that the vast majority of $A[1], \ldots, A[h]$ has a common value (0 or 1), to find that value. Do this by attempting, using the algorithm of Lemma 4.1 with $d = \lfloor (h/4)^{1/4} \rfloor$, to move the set of ones in A to an array of size $h/4$. Set $y = 1$ if and only if this fails.

In order to analyze the last part of the algorithm, note that $S = \sum_{i=1}^{h} A[i]$ is binomially distributed, and that the preceding discussion implies that $E(S) \geq h/2$ if $b \geq n/2$, while $E(S) \leq 1$ if $b \leq n/8$. By another application of Lemma 2.1, the following happens with high probability: $S > h/4$ if $b \geq n/2$, while $S \leq (h/4)^{1/4}$ if $b \leq n/8$. In the first case, the compaction using the algorithm of Lemma 4.1 surely fails, while in the second case it will succeed. In either case y receives the correct value. ∎

When using the algorithm of Lemma 5.1 to analyze the outcome of a GCS $\mathcal{S} = \{S_1, \ldots, S_r\}$ below, we apply the algorithm separately to each row of $\mathcal{S}$ and define a row to be *almost-full* if the algorithm assigns the value 1 to the bit y associated with the row. The *threshold* of $\mathcal{S}$ is 0 if S_1 is not almost-full, and otherwise is the largest integer $i \in \{1, \ldots, r\}$ such that S_j is almost-full, for $j = 1, \ldots, i$.

In loose analogy with the definition of the f-row of a GCS and motivated by Lemma 5.1, define an (f_1, f_2)-*row* of a GCS $\mathcal{S} = (S_1, \ldots, S_r)$, for $0 \leq f_1 \leq f_2 \leq 1$, as 0 if none of $S_1, \ldots, S_r$ has fullness $> f_1$, and otherwise as any integer $i \in \{1, \ldots, r\}$ such that S_i has fullness $> f_1$, while either $i = r$ or S_{i+1} has fullness $< f_2$.

Lemma 5.2: Let $m, r, s \in \mathbb{N}$ and let L be a $(\frac{1}{8}, \frac{1}{2})$-row of a GCS of a set of m elements with parameters $r \times s$. Let $M = 2^L s$ and take $c_1 = 1/(2^{12}e)$ and $c_2 = 12$. Then

(a) If $m \geq c_1 s$, then $\Pr(m < c_1 M) \leq 2^{-s}$;

(b) $\Pr(L > 0 \text{ and } m < c_1 M) \leq 2^{-s}$;

(c) If $r \geq \lfloor \log m \rfloor$, then $\Pr(m > c_2 M) \leq 2^{-s}$.

Proof: We proceed as in the proof of Lemma 3.3. If $L > 0$, the fullness of row L is at least $1/8$. Hence by Lemma 3.1(d), for every $l \geq 0$,

$$\Pr(L > l) \leq \sum_{i=\lceil l \rceil}^{\infty} \left(\frac{8em \cdot 2^{-i}}{s}\right)^{\lceil s/8 \rceil} \leq 2\left(\frac{8em \cdot 2^{-l}}{s}\right)^{\lceil s/8 \rceil} \leq \left(\frac{2^{4-l}em}{s}\right)^{\lceil s/8 \rceil}. \quad (3)$$

Likewise, if $L < r$, the fullness of row $L+1$ is at most $1/2$. Hence by Lemmas 3.1(b) and 3.2, for every $l \leq r$,

$$\begin{aligned} \Pr(L < l) &\leq \min\{1, \sum_{i=-\infty}^{\lfloor l \rfloor} 2^{s-m\cdot 2^{-i-2}}\} \\ &\leq 2^s \cdot \min\{1, \sum_{i=0}^{\infty} 2^{-2^i m \cdot 2^{-\lfloor l \rfloor - 2}}\} \leq 2^{s+1-m\cdot 2^{-l-2}}. \end{aligned} \quad (4)$$

To verify (a), apply (3) with $l = \log(m/(c_1 s)) \geq 0$ to obtain

$$\Pr(m < c_1 M) = \Pr(L > l) \leq \left(\frac{2^{4-l}em}{s}\right)^{\lceil s/8 \rceil} \leq 2^{-s}.$$

(b) follows immediately from (a) and the observation that $L > 0$ implies $m \geq s/8 \geq c_1 s$. To verify (c), apply (4) with $l = \log(m/(c_2 s)) \leq r$ to obtain

$$\Pr(m > c_2 M) = \Pr(L < l) \leq 2^{s+1-m\cdot 2^{-l-2}} = 2^{s+1-c_2 s/4} = 2^{1-2s} \leq 2^{-s}. \quad \blacksquare$$

The algorithm described in the theorem below outputs a sequence of independent integers, except that it may fail and not produce any output at all. As regards the independence, the precise statement is that for each input and conditionally on the event that any output is produced, the integers output by the algorithm are independent. Similar interpretations should be imposed on other results in the sequel concerning randomized algorithms that are claimed to output independent random numbers. A simpler proof of a statement similar to Theorem 5.3 was indicated by Goodrich (1991).

Theorem 5.3: For every fixed $\delta > 0$ there is a constant $\epsilon > 0$ such that for all given $n, m, \tau \in I\!N$ with $m = O(n^{1-\delta})$, m-color strong fine-profiling problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo).

Proof: The idea of the algorithm is simple: If the size of a color class is $n^{\Omega(1)}$, it can be reliably estimated using a GCS; otherwise the color class can be blown up by a factor of $n^{\Theta(1)}$ and its size estimated in the same way. We now provide the details.

Without loss of generality assume that δ is rational and that $n \geq 2$. Since computing exact multiplicities reduces to sorting, the given problem is easily solved using the

algorithm of Lemma 2.9 if $\tau = n^{\Omega(1)}$; we can therefore also assume the availability of at least $n^{1-\delta/3}$ processors. Let $s = \lceil n^{\delta/3} \rceil$. For $i = 1, \ldots, m$, use guarded writing to set $\widehat{b}_i = 0$ if $b_i = 0$, and otherwise carry out the following procedure, where c_1 and c_2 are as defined in Lemma 5.2:

Step 1: Using the algorithm of Lemma 5.1, execute a GCS $\mathcal{S}_i$ of $\mathcal{B}_i$ with parameters $\lfloor \log n \rfloor \times s$ and compute l_i as the threshold of $\mathcal{S}_i$. Note that we need a guard processor for each cell of each GCS, and that by assumption sufficiently many processors are available.

Step 2: If $l_i > 0$, take $\widehat{b}_i = c_2 \cdot 2^{l_i} s$. Otherwise use the algorithm of Corollary 4.3 to allocate s processors to each element of $\mathcal{B}_i$, let these processors execute a GCS $\mathcal{S}_i'$ with parameters $\lfloor \log(sn) \rfloor \times s$ and take $\widehat{b}_i = c_2 \cdot 2^{l_i'}$, where l_i' is the threshold of $\mathcal{S}_i'$.

It is easy to see that the space needed by the algorithm is $O(n)$. To see that the same holds for the number of operations, note that by Lemma 5.2(c) and condition (1) of Lemma 5.1, no processors are allocated to a fixed color class $\mathcal{B}_i$ with $b_i > c_2 s$, except with negligible probability.

Now fix $i \in \{1, \ldots, m\}$. The following happens with high probability: If $l_i > 0$, we have $b_i \leq \widehat{b}_i \leq (c_2/c_1) b_i$ (by Lemma 5.2, parts (b) and (c)). If $l_i = 0$, then $s b_i \leq c_2 \cdot 2^{l_i'} s \leq (c_2/c_1) s b_i$, i.e., $b_i \leq \widehat{b}_i \leq (c_2/c_1) b_i$ (by Lemma 5.2, parts (a) and (c)). In either case, $\widehat{b}_i$ is a linear overestimate for b_i.

The estimates produced by the algorithm are clearly independent unless the processor allocation according to Corollary 4.3 fails, in which case the algorithm can report failure and refrain from producing any output. ∎

Remark: The above algorithm for fine-profiling is Monte Carlo, i.e., we cannot detect if the estimates computed are off by more than the allowed constant factor. In Section 10 we derive a Las Vegas fine-profiling algorithm (Corollary 10.5).

6 Colored Compaction

It is essential for the application to interval allocation described in Section 7 as well as for other reasons to generalize the compaction problem studied in Section 4 to *colored compaction*, where objects of different *colors*, initially placed in a single source array, are to be moved to distinct destination arrays, one for each color. As before, an object that cannot be placed is called unlucky, and we will not distinguish between an object and the input element representing it. In the formal definition below, the value of an element represents its color, the special value 0 still representing the absence of an object. This is in agreement with our terminology concerning $\mathcal{B}_1, \ldots, \mathcal{B}_m$.

Definition: Given $n, m \in \mathbb{N}$ and $d_1, \ldots, d_m \geq 0$ as well as n integers $x_1, \ldots, x_n$ in the range $0 \,..\, m$, an *incomplete placement* for $x_1, \ldots, x_n$ with bounds $d_1, \ldots, d_m$ is a sequence $y_1, \ldots, y_n$ of n nonnegative integers such that

(1) For $j = 1, \ldots, n$, if $x_j = 0$, then $y_j = 0$;

(2) For $1 \leq i < j \leq n$, if $x_i = x_j$ and $y_i \neq 0$, then $y_i \neq y_j$;

(3) For $i = 1, \ldots, m$, $\max(\{y_j : 1 \leq j \leq n \text{ and } x_j = i\} \cup \{0\}) \leq d_i$.

The set $\{j : 1 \leq j \leq n,\ x_j \neq 0 \text{ and } y_j = 0\}$ is called the *residue set* of the incomplete placement. If the residue set is empty, the placement is *complete.*

Condition (2) ensures that distinct elements of the same color are not placed in the same destination cell, and condition (3) states that the lucky elements of $\mathcal{B}_i$ fit into an array of size $\lfloor d_i \rfloor$, for $i = 1, \ldots, m$. If there is only a single color, i.e., for $m = 1$, the above definition reduces to our earlier definition of an incomplete placement.

The compaction problems introduced in Section 4 generalize in a natural way to colored compaction. Our next goal is to extend the compaction algorithms given for the special case of a single color to the case of several colors. Recall that in a first approximation, the algorithm of Lemma 4.4 multi-scatters the active elements over an array of size s, for suitably chosen s. A straightforward generalization to the case of m colors would be to multi-scatter the elements of $\mathcal{B}_i$ over an array associated with $\mathcal{B}_i$ and of a suitably chosen size s_i, for $i = 1, \ldots, m$. Attempting this, we are faced with a somewhat extraneous problem, namely that we do not know how to allocate m disjoint arrays of prescribed sizes $s_1, \ldots, s_m$ sufficiently fast without wasting too much space. Lemma 6.1 below therefore assumes "pre-allocated" such arrays to be supplied by any "user" of the lemma; the sizes of these arrays simultaneously serve as the bounds $d_1, \ldots, d_m$. As is rather obvious, the compaction of a particular color class will not be very successful unless the size of the array provided for that color class is considerably larger than the size of the color class — we later define such color classes to be *well-supplied.* Lemma 6.1 therefore identifies the set of (indices of) elements in well-supplied color classes, and its assertions apply only to such elements.

A more interesting complication in the generalization of the algorithm of Lemma 4.4 to the case of several colors lies in the fact that in Step 2 of the algorithm, several elements forming a cluster are multi-scattered together. While this works fine in the case of a single color, it is not appropriate if the elements in a cluster have different colors, i.e., are to be placed in different destination arrays. In order to solve this problem, recall that the clusterwise scattering was motivated by efficiency considerations: multi-scattering single elements works just as well, but requires several processors standing by each active element. The idea now is first to compact the active elements as though they were all of the same color — we already know how to do that — but to use the outcome of this compaction exclusively to allocate the necessary processors to each active element, after

which the colored compaction can be completed in the simple way described above. As suggested by this description, an initial part of our algorithm for colored compaction is quite similar to the corresponding algorithm for (uncolored) compaction. Because of the slight differences and since we want to extend the analysis of the algorithm given earlier, however, we essentially reproduce it as Steps 1 and 2 of the algorithm of Lemma 6.1 below.

Before we state the lemma, recall that for $c > 0$, a real-valued function S defined on a set M equipped with a metric ϕ is said to satisfy a *Lipschitz condition* with constant c if for all $x, y \in M$, we have $|S(x) - S(y)| \leq c \cdot \phi(x, y)$. The only metric space relevant to the present paper, and the one implicitly intended in every reference to a Lipschitz condition, is the set of subsets of $\{1, \ldots, n\}$, equipped with the metric ϕ with $\phi(U, V) = |U \triangle V|$, for all $U, V \subseteq \{1, \ldots, n\}$, where $U \triangle V$ denotes the symmetric difference of U and V, i.e., $U \triangle V = (U \backslash V) \cup (V \backslash U)$.

Lemma 6.1: There is a constant $\epsilon > 0$ such that the following holds: Let $n, m, v, \tau \in \mathbb{N}$ be given, suppose that $x_1, \ldots, x_n$ are n given integers in the range $0..m$ and let $A_1, \ldots, A_m$ be m given nonoverlapping arrays. Take $B_i = \{j : 1 \leq j \leq n \text{ and } x_j = i\}$ and $b_i = |B_i|$, for $i = 1, \ldots, n$, and define $J = \{i : 1 \leq i \leq m \text{ and } |A_i| \geq 6vb_i\}$, $B' = \bigcup_{i \in J} B_i$ and $b = \sum_{i=1}^{m} b_i$. Then an incomplete placement for $x_1, \ldots, x_n$ with bounds $|A_1|, \ldots, |A_m|$ can be computed on a TOLERANT PRAM using $O(\tau)$ time, $\lceil (n + v^3 b)/\tau \rceil$ processors and $O(n + v^3 b)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo), such that the residue set D of the placement satisfies conditions (a)–(c) below.

(a) For every $j \in B'$, $\Pr(j \in D) \leq 2^{-v}$;

(b) For every fixed subset R of B' and for all $z \geq |R|/2^v$,

$$\Pr(|R \cap D| \geq z) \leq 2e^{-z^2/(2^9 |R| v^3)};$$

(c) For every nonnegative real function S of D that satisfies a Lipschitz condition with constant c,

$$S = O(E(S) + cv^3 n^{5/8})$$

with probability at least $1 - 2^{-n^{1/8}}$.

Remark: As anticipated above, a color class B_i is called *well-supplied*, in the context of Lemma 6.1, if $|A_i| \geq 6vb_i$, so that B' is the set of indices of elements in well-supplied color classes. Part (a) of the lemma says that any fixed element of a well-supplied color class is unlikely to remain active, part (b) extends this property from single elements to arbitrary sets of elements in well-supplied color classes, and part (c) states that any function of the residue set that satisfies a Lipschitz condition with a constant that is not too large with high probability does not significantly exceed its expected value.

Proof: With the same justification as in the proof of Lemma 4.4, we give the proof only for $\tau = 1$. Start by using the algorithm of Theorem 5.3 to compute an estimate $\widehat{b}$ of b and assume that $\widehat{b}$ is indeed a linear overestimate for b. Then execute the following:

Step 1: Divide $\mathcal{X}$ into $l = \lceil n/v \rceil$ *clusters* $\mathcal{X}_1, \ldots, \mathcal{X}_l$ of at most v input elements each and use the algorithm of Lemma 2.7 to compute a bit representation of the set $I = \{j : 1 \leq j \leq l \text{ and } \mathcal{X}_j \text{ contains at least one active element}\}$.

Step 2: Associate $4v$ processors with each element of I and $4v$-scatter I over an array A of size $8v\widehat{b}$; let $I' \subseteq I$ denote the set of unsuccessful indices.

Step 3: Associate $3v^2$ processors with each cell of A and use these to $3v$-scatter the active elements in $\mathcal{B}_i \cap \bigcup_{j \in I \setminus I'} \mathcal{X}_j$ over A_i, for $i = 1, \ldots, m$.

The algorithm clearly runs in constant time on a TOLERANT PRAM using $O(n + v^3 b)$ processors and $O(n + v^3 b)$ memory cells. Since $|I| \leq b \leq \widehat{b}$, the density of the multi-scattering in Step 2 is bounded by $1/2$ and, by definition, the densities of the multi-scatterings of well-supplied color classes in Step 3 are also bounded by $1/2$. Let us agree to call a cluster $\mathcal{X}_i$ with $i \in I'$ (i.e., i was unsuccessful in Step 2) *unsuccessful.* If we take D_1 for the index set of active elements in unsuccessful clusters and denote by D_2 the index set of elements that are unsuccessful in Step 3, clearly $D = D_1 \cup D_2$.

For the proof of part (a), fix an arbitrary active element X_j in a well-supplied color class (i.e., $j \in B'$) and suppose that $X_j \in \mathcal{X}_i$. As argued above, $j \in D$ (i.e., X_j is unlucky) if and only if either $i \in I'$ (i.e., i is unsuccessful in Step 2) or $j \in D_2$ (i.e., X_j participates in Step 3, but is unsuccessful). But by Lemma 3.6(a), applied twice with $p \leq 1/2$, $\Pr(i \in I') \leq 2^{-4v}$ and $\Pr(j \in D_2) \leq 2^{-3v}$; hence $\Pr(j \in D) \leq 2^{-4v} + 2^{-3v} \leq 2^{-v}$.

For part (b), let R be a fixed subset of B', take $r = |R|$ and let $z \geq r/2^v$. Clearly $|R \cap D| \geq z$ only if either $|R \cap D_1| \geq z/2$ or $|R \cap D_2| \geq z/2$; we will consider these events separately and show each to be unlikely. First consider the multi-scattering in Step 2 and let $R' \subseteq I$ be the index set of clusters containing at least one active element whose index belongs to R; obviously $|R'| \leq r$. If $|R \cap D_1| \geq z/2$, i.e., at least $z/2$ elements of R are indices of elements in unsuccessful clusters, there must be at least $z/(2v)$ unsuccessful clusters containing elements with indices in R, i.e., $|R' \cap I'| \geq z/(2v)$. In other words, $\Pr(|R \cap D_1| \geq z/2) \leq \Pr(|R' \cap I'| \geq z/(2v))$. Since $z/(2v) \geq 2r \cdot 2^{-4v}$, we can apply Lemma 3.6(b) to bound the latter probability by $e^{-\zeta}$, where $\zeta = (z/(2v))^2/(32|R'| \cdot 4v) \geq z^2/(2^9 r v^3)$. Since also $z/2 \geq 2r \cdot 2^{-3v}$, another application of Lemma 3.6(b) shows that $\Pr(|R \cap D_2| \geq z/2) \leq e^{-\zeta}$, where $\zeta = (z/2)^2/(32|R| \cdot 3v) \geq z^2/(2^9 r v)$. It follows that $|R \cap D| \geq z$ with probability at most $e^{-z^2/(2^9 r v^3)} + e^{-z^2/(2^9 r v)} \leq 2e^{-z^2/(2^9 r v^3)}$.

For part (c), note that S can be considered as a function of all the random choices made by the algorithm. Now, a change in a single random choice made in Step 2 affects at

most 2 clusters. Each of the at most $2v$ elements in the affected clusters in turn affects at most $3v$ other elements in Step 3. Hence altogether at most $2v + 6v^2 \leq 8v^2$ of the output variables $y_1, \ldots, y_n$ are affected, and it is easy to see that no more output variables are affected by a change in a single random choice in Step 3. In other words, if D changes to D' in response to a change in a single random choice, we always have that $|D \triangle D'| \leq 8v^2$; by the Lipschitz condition imposed on S, this means that S changes by at most $8cv^2$. Since the algorithm makes at most $7vb \leq 7vn$ random choices altogether, an application of Corollary 2.3 yields that $S \leq \max\{2E(S), 4 \cdot 8cv^2 \cdot (7vn)^{5/8}\} = O(E(S) + cv^3n^{5/8})$ with probability at least $1 - 2^{-n^{1/8}}$. ∎

In order to simplify the applications of Lemma 6.1 in the following, we discuss a generic application in detail at this point and introduce a convenient shorthand that will be used in later applications.

Assume that we are given a set $\mathcal{B}$ of active elements with colors in $\{1, \ldots, m\}$ and stored in an array Q of size n. For $i = 1, \ldots, m$, let $\mathcal{B}_i$ be the set of elements in $\mathcal{B}$ with color i. Further assume that we are given m disjoint arrays $A_1, \ldots, A_m$. Our goal is to place most of the elements of $\mathcal{B}_i$ in A_i, for $i = 1, \ldots, m$. To this end let x_j be the color of the element X_j in the jth cell of Q, for $j = 1, \ldots, n$, with $x_j = 0$ if the jth cell of Q does not contain any active element, and let v be a suitable positive integer (the choice of v will depend on our (application-specific) knowledge of the ratios $|A_i|/|\mathcal{B}_i|$). Then apply the algorithm of Lemma 6.1 to compute an incomplete placement $y_1, \ldots, y_n$ for $x_1, \ldots, x_n$ with bounds $|A_1|, \ldots, |A_m|$. Finally, for all $j \in \{1, \ldots, n\}$ with $x_j \neq 0$ and $y_j \neq 0$, actually place X_j in the y_jth cell of A_{x_j} (the algorithm of Lemma 6.1 already places X_j as just described; this is not specified in the statement of the lemma, however, so we repeat the operation here).

In what follows, an application of Lemma 6.1 as above will be called simply "v-compacting $\mathcal{B}_i$ to A_i, for $i = 1, \ldots, m$". Lemma 6.1 guarantees that a fixed element of a well-supplied color class will be unlucky with probability at most 2^{-v}. Furthermore, if the number of nonzero input numbers is $O(n/v^3)$, as in most applications of Lemma 6.1 in the present paper, then the algorithm uses $O(n)$ operations and $O(n)$ space.

In the remainder of the section we extend the definition of complete linear compaction to the case of several colors and prove a result corresponding to Theorem 4.6.

Definition: For all $n, m, d_1, \ldots, d_m \in \mathbb{N}$, the *complete linear colored compaction* problem of size n and with limits $d_1, \ldots, d_m$ is, given n integers $x_1, \ldots, x_n$ in the range $0..m$ such that $|\{j : 1 \leq j \leq n \text{ and } x_j = i\}| \leq d_i$, for $i = 1, \ldots, m$, to compute a complete placement for $x_1, \ldots, x_n$ with bounds $O(d_1), \ldots, O(d_m)$.

The problem, discussed before the statement of Lemma 6.1, of allocating m disjoint arrays $A_1, \ldots, A_m$ is easy in a special case, namely when m is so small that the allocation

can be done by means of brute-force prefix summation (Corollary 2.5). This leads to the following result.

Theorem 6.2: There is a constant $\epsilon > 0$ such that for all given $n, m, \tau, d_1, \ldots, d_m \in \mathbb{N}$ with $m = (\log n)^{O(1)}$ and $\tau \geq \log^* n$, complete linear colored compaction problems of size n and with limits $d_1, \ldots, d_m$ can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n + \sum_{i=1}^{m} d_i)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: We first devise a nonoptimal algorithm that solves the problem in $O(\log^* n)$ time using n processors and afterwards show that optimality can be achieved essentially as in the proof of Theorem 4.6. The basic idea is to place the active elements in $\log^* n$ successive stages, similarly as in the algorithm of Corollary 4.5. With a view towards a future application (Lemma 9.7), we first consider a more general setting, in which we ignore the difficulties of space allocation discussed above, but in return show how to tolerate $m = O(n^{1-\delta})$ different colors, for arbitrary fixed $\delta > 0$.

Without loss of generality assume that δ is rational and that $d_i \leq n$, for $i = 1, \ldots, n$ (otherwise the elements in $\mathcal{B}_i$ can be deactivated in a trivial manner). Say that a color class $\mathcal{B}_i$ is *small* if $b_i \leq n^{\delta/2}$, and *large* otherwise, and note that the total number of elements in small color classes is $O(n^{1-\delta} \cdot n^{\delta/2}) = O(n^{1-\delta/2})$. We begin by reducing the number of active elements in large color classes to a similar level, but in a balanced way (i.e., each large color class loses most of its elements).

Define the *active fraction* as zero if all color classes are small, and otherwise as the maximum, over all large color classes $\mathcal{B}_i$, of the ratio of the number of (currently) active elements in $\mathcal{B}_i$ to b_i. We aim at decreasing the active fraction to at most $n^{-\delta/8}$, after which the total number of remaining active elements will be $O(n^{1-\delta/2} + n^{1-\delta/8}) = O(n^{1-\delta/8})$. We first show that if the active fraction has been reduced to at most v^{-3}, for some given $v \in \mathbb{N}$, then in constant time it can be reduced further to at most $\max\{2^{-3v}, n^{-\delta/8}\}$ with high probability. If $v \geq n^{\delta/24}$, there is nothing to show. Otherwise use the algorithm of Lemma 6.1 to $3v$-compact the remaining active elements in $\mathcal{B}_i$ to an array of size $\lceil 6 \cdot 3d_i/v^2 \rceil$, for $i = 1, \ldots, m$. Assuming that the active fraction is at most v^{-3}, Lemma 6.1 shows that the $3v$-compaction can be carried out in constant time using $O(n + v^3(n^{1-\delta/2} + \sum_{i=1}^{m}(b_i/v^3))) = O(n)$ processors and that for each fixed large color class $\mathcal{B}_i$ (which, by assumption, is well-supplied), the number of active elements in $\mathcal{B}_i$ after the $3v$-compaction is bounded by $\max\{b_i/2^{3v}, b_i^{3/4}\}$, except with probability at most $2e^{-\zeta}$, where $\zeta = (b_i^{3/4})^2/(2^9(b_i/v^3) \cdot (3v)^3) = \Omega(b_i^{1/2}) = \Omega(n^{\delta/4})$. With high probability, the fraction of active elements left in each large color class $\mathcal{B}_i$ after the $3v$-compaction therefore is at most $\max\{2^{-3v}, b_i^{-1/4}\} \leq \max\{2^{-3v}, n^{-\delta/8}\}$, for $i = 1, \ldots, m$, i.e., the active fraction has been reduced to the same level. Applying this procedure at most $\log^* n$ times with $v = 1, 2, 2^2, 2^{2^2}, \ldots$ with high probability reduces the

active fraction to at most $n^{-\delta/8}$, as desired. Note that the successive arrays used by a color class can be taken as tightly packed subarrays of a single array.

In order to complete the compaction, we first use the algorithm of Corollary 4.3 to move the remaining $O(n^{1-\delta/8})$ active elements to an array of size $O(n^{1-\delta/8})$, after which constant time and n processors suffice to $\lfloor n^{\delta/8} \rfloor$-scatter the active elements of $\mathcal{B}_i$ over an array of size $2 \cdot \max\{d_i, \lceil n^{3\delta/4} \rceil\}$, for $i = 1, \ldots, m$. Since the number of active elements in a large color class $\mathcal{B}_i$ is at most $b_i/n^{\delta/8} \leq d_i/n^{\delta/8}$, the density of each of these multi-scatterings is bounded by $1/2$, so that Lemma 3.6(a) ensures that with high probability all elements are successful. At this point, for $i = 1, \ldots, m$, the elements of $\mathcal{B}_i$ are stored in an array of size $\lceil 18d_i/1^2 \rceil + \lceil 18d_i/2^2 \rceil + \lceil 18d_i/4^2 \rceil + \cdots + 2 \cdot \max\{d_i, \lceil n^{3\delta/4} \rceil\} = \Theta(d_i + n^{3\delta/4})$. For all color classes $\mathcal{B}_i$ with $d_i \geq n^{3\delta/4}$, this compaction is sufficiently tight, while the remaining color classes can be compacted into linear space using the algorithm of Corollary 4.3 (recall that we have $\Theta(n^{\delta})$ processors for each color class).

Under the restriction $m = (\log n)^{O(1)}$ of the theorem, the allocation of an array to each color class, which was ignored above, can clearly be done in constant time using the algorithm of Corollary 2.5.

The algorithm described so far can be excuted in $O(\log^* n)$ time with n processors. To achieve optimal speedup, it suffices, in the light of Theorem 4.6, to show that the number of active elements can be reduced to $O(n/\tau)$ in $O(\tau)$ time. To this end assume that $\tau \leq n^{1/8}$ (otherwise sort the input numbers using the algorithm of Lemma 2.9), divide $\mathcal{X}$ into $\lceil n/\tau \rceil$ *clusters* of at most τ input elements each and associate a processor with each cluster. Then use the algorithm of Corollary 2.5 to allocate an array A_i of size $8d_i$ to $\mathcal{B}_i$, for $i = 1, \ldots, m$. Similarly as in the proof of Theorem 4.6, each processor now attempts in 2τ *rounds* to place the active elements of its cluster in the arrays corresponding to their colors. The argument in the proof of Theorem 4.6 shows that with high probability the number of active elements left after the last round is $O(n/\tau)$. ∎

7 Interval Allocation

While the compaction problem asks that unit intervals be placed in a base segment, the interval allocation problem, defined below, specifies intervals of varying length to be placed. Viewed another way, each input element is a request for a *block* of consecutive indices of a size given by the value of the request. Informally, condition (2) means that blocks do not overlap, and (3) means that the allocated blocks are optimally packed, except for a constant factor. Another difference between complete linear compaction and interval allocation is that in the case of compaction, an upper bound on the number of elements present (the limit d) is provided as part of the input, while in the case of interval allocation the choice of an appropriate size for the base segment is left to the

algorithm. Since a suitable value for d could actually be computed using the algorithm of Theorem 5.3 as in the proof of Lemma 6.1, this difference is of no particular significance, but merely convenient for the exposition.

Definition: For all $n \in \mathbb{N}$, the (complete) *interval allocation* problem of size n is the following: Given n nonnegative integers $x_1, \ldots, x_n$, compute n nonnegative integers $y_1, \ldots, y_n$ such that

(1) For $j = 1, \ldots, n$, $x_j = 0 \Leftrightarrow y_j = 0$;

(2) For $1 \leq i < j \leq n$, if $0 \notin \{x_i, x_j\}$, then $\{y_i, \ldots, y_i + x_i - 1\} \cap \{y_j, \ldots, y_j + x_j - 1\} = \emptyset$;

(3) $\max\{y_j : 1 \leq j \leq n\} = O(\sum_{j=1}^{n} x_j)$.

A natural extension of the interval allocation problem would augment the solution by an appropriate size for the base segment, i.e., by an integer s with $s \geq \max\{y_j + x_j - 1 : 1 \leq j \leq n\}$, but $s = O(\sum_{j=1}^{n} x_j)$. Such a quantity is usually needed in applications of interval allocation, and our algorithm for interval allocation essentially generates it internally. By Lemma 2.6, however, a suitable choice for s (namely, $\max\{y_j + x_j - 1 : 1 \leq j \leq n\}$) can be computed from the output of interval allocation, as defined here, for which reason we have refrained from including it in the problem definition.

While interval allocation is a generalization of compaction, we will now show that interval allocation reduces to colored compaction. First note that if all nonzero requests (i.e., requests of nonzero value) are of value 1, i.e., if $x_j \in \{0, 1\}$, for $j = 1, \ldots, n$, then we can indeed first use the algorithm of Theorem 5.3 to compute a linear overestimate $\widehat{b}$ for the number b of nonzero requests, and subsequently solve the complete linear compaction problem with input $x_1, \ldots, x_n$ and limit $\widehat{b}$. Informally, what happens is that nonzero requests are interpreted as active elements in the usual sense; once the correct size of the destination array A has been established, a unit block (i.e., a single index) is associated with each cell of A, the active elements are placed in A by the compaction algorithm, and the block associated with a cell in A is allocated to the element, if any, placed in that cell (i.e., the requests are satisfied). The same approach works as long as all nonzero requests are of a common value l — we simply associate a block of l consecutive indices with each cell of A, rather than a single index.

If there are nonzero requests of m distinct values, we clearly have to use m distinct destination arrays $A_1, \ldots, A_m$, each associated with blocks of a different size, i.e., we have to resort to colored compaction. A number of difficulties have to be tackled in this generalization. First, we need to estimate the sizes of several color classes simultaneously; this can still be done by the algorithm of Theorem 5.3 if $m = O(n^{1-\delta})$, for some fixed $\delta > 0$. Second, the colored compaction can be carried out using the algorithm of Theorem 6.2, but only under the (more stringent) restriction $m = (\log n)^{O(1)}$. Third,

the blocks of indices associated with $A_1, \ldots, A_m$ must be allocated from an appropriate base array. More precisely, with each array A_i we allocate a *segment* consisting of $|A_i|$ blocks of the appropriate size, i.e., of the size associated with the color i, after which the association of a single block with each cell in A_i is trivial. Since this allocation reduces to prefix summation, we aim to carry it out using the algorithm of Corollary 2.5. This requires, on the one hand, that $m = (\log n)^{O(1)}$, as above, and, on the other hand, that the sizes of the blocks to be allocated is polynomial in n. On the outset, these requirements are not satisfied: The numbers $x_1, \ldots, x_n$ could all be distinct, which would give us as many as n color classes, and they could be arbitrarily large. However, the input can be scaled and rounded to satisfy the requirements, as we show next.

First observe that if M denotes the maximum request value, i.e., $M = \max\{x_j : 1 \leq j \leq n\}$, then replacing each nonzero request value x_j by the nearest multiple of $u = \lceil M/n \rceil$ no smaller than x_j increases the sum W of all request values by at most $ub \leq M + b$, where b is the number of nonzero requests, i.e., at most triples W. As a result of this transformation, we can consider all request values to be integers in the range $0..n$ (simply measure requests and blocks in units of size u). If at this point we replace each nonzero (modified) request value by the nearest larger power of 2, W at most doubles, and only $O(\log n)$ different request values remain, i.e., the compaction-based algorithm sketched above becomes applicable with $m = O(\log n)$ color classes. Every modified request value is at least as large as the corresponding original value; since the modified values sum to at most 6 times the original sum, however, solving the interval allocation problem defined by the modified request values produces a solution to the original interval allocation problem.

The maximum request value M can be computed using the algorithm of Lemma 2.6. In fact, for the applications of Theorem 7.1 in the present paper, we will frequently have $M = O(n)$, in which case it is not necessary to actually compute M (because the procedure described above, with trivial modifications, can be employed with $u = 1$ whenever M is polynomial in n). Summing up, we have seen that interval allocation reduces to complete linear colored compaction with a logarithmic number of colors.

Theorem 7.1: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in I\!N$ with $\tau \geq \log^* n$, interval allocation problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo).

Proof: By the above discussion and Theorem 6.2. In particular, note that since resources (indices) are divided optimally between color classes and with a constant-factor "waste" within color classes, the overall "waste factor" is bounded by a constant, as required by condition (3) in the definition of interval allocation. ∎

Remark: Part of the development of Theorem 7.1 took place in a dialog with

Joseph Gil. An earlier version of the present paper achieved running times of $O(\log\log n \log^* n / \log\log\log n)$, the bottleneck being compaction. After receiving a preliminary sketch of the algorithm of Theorem 7.1 geared towards this running time, Gil informed us of the results of Matias and Vishkin (1991), unpublished at the time, and observed their applicability in the context of the algorithm, which allowed him to derive a first interval allocation algorithm with a running time of $O(\log^* n)$. We improved his result by giving an algorithm with optimal speedup and a lower failure probability and by implementing the algorithm on the weaker TOLERANT PRAM, after which the communication with Gil ceased. Theorem 7.1 states the last result mentioned above, except for a still smaller failure probability. A slightly weaker result was published in (Gil *et al.*, 1991). A time bound of $O(\log\log n)$ for a less general *load balancing* problem was shown by Gil (1990, 1991).

Remark: At this point we can give only a Monte Carlo algorithm for interval allocation, and Theorem 7.1 is formulated accordingly. This is due to the use of the Monte Carlo fine-profiling algorithm of Theorem 5.3. Using Corollary 10.5 instead of Theorem 5.3, however, allows us to obtain a Las Vegas algorithm for interval allocation. The same remark applies to Theorem 7.2 below.

Whereas the use of Theorem 7.1 in memory allocation is obvious, one additional observation is needed for its application to the allocation of processors. The reason is that a processor is an active device that needs to know about the task that it is to execute. Theorem 7.1 can be used to communicate this information to the first processor in each *team*, i.e., in each group of consecutively numbered processors allocated to a common task, but the information must subsequently be broadcast to the remaining processors in each team. In recognition of this fact, we consider a slight variation of the interval allocation problem called the *interval marking* problem. In the definition below, informally, $x_1, \ldots, x_n$ are the sizes of n requests for processors. The output consists of a size indicator s together with s integers $z_1, \ldots, z_s$ and specifies the allocation of s virtual processors $P_1, \ldots, P_s$ as follows: For $j = 1, \ldots, s$, the meaning of $z_j = i \in \{1, \ldots, n\}$ is that P_j is a member of the team allocated to the ith request; the meaning of $z_j = 0$ is that P_j is not part of any team. Condition (1) requires the processors in each team to be consecutively numbered, (2) expresses that the number of virtual processors in the team allocated to the ith request is indeed exactly x_i, for $i = 1, \ldots, n$, and (3) states that the total number of virtual processors exceeds the number of requested processors by at most a constant factor; this allows the allocated processors to be simulated without loss, up to a constant factor, by any number of available physical processors.

Definition: For all $n \in \mathbb{N}$, the *interval marking* problem of size n is the following: Given n nonnegative integers $x_1, \ldots, x_n$, compute nonnegative integers $s, z_1, \ldots, z_s$ such that

(1) For all integers i, j, k with $1 \leq i \leq j \leq k \leq s$, if $z_i = z_k \neq 0$, then $z_j = z_i$;
(2) For $i = 1, \ldots, n$, $|\{j : 1 \leq j \leq s \text{ and } z_j = i\}| = x_i$;
(3) $s = O(\sum_{j=1}^{n} x_j)$.

In most applications of interval marking in the present paper we will have $\sum_{j=1}^{n} x_j = O(n)$. We next show that under this restriction, we can solve the interval marking problem with input $x_1, \ldots, x_n$ in constant deterministic time with n processors after solving the interval allocation problem with input $x_1, \ldots, x_n$ using the algorithm of Theorem 7.1. The reader may think of this as a reduction of interval marking to interval allocation; this is not quite exact, however, since we will use a special property of the solution produced by Theorem 7.1 (conversely, under the restriction $\sum_{j=1}^{n} x_j = O(n)$, it is easy to show that interval allocation reduces to interval marking).

We view interval allocation with input $x_1, \ldots, x_n$ as allocating disjoint subarrays $A_1, \ldots, A_n$ of sizes $x_1, \ldots, x_n$ from a base array and note that it is trivial to mark the first cell of A_i with the integer i, for $i = 1, \ldots, n$. As already discussed above, the corresponding interval marking problem can essentially be solved by copying the integer stored in the first cell of A_i to the remaining cells of A_i, for $i = 1, \ldots, n$. In other words, it suffices to provide each cell of a subarray with a pointer to the beginning of the subarray. Now recall that our algorithm for interval allocation actually allocates all subarrays from $O(\log n)$ segments, each of which consists of tightly packed subarrays of the same size. If we store the subarray size of each segment in the first cell of the segment, which is easy to do, it suffices to provide each cell of a segment with a pointer to the beginning of the segment, since with this pointer and the relevant subarray size it can easily compute the beginning of its subarray. We are now left with a problem that can be viewed as an instance of the nearest preceding element problem defined in Section 2; each beginning of a segment corresponds to one nonzero input bit. Since the number of segments is $O(\log n)$, we can appeal to part (b) of Lemma 2.8 and obtain the following result.

Theorem 7.2: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$ with $\tau \geq \log^* n$, interval marking problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil (n + W)/\tau \rceil$ processors and $O(n + W)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo), where W is the sum of the input numbers.

Proof: By the above discussion, Lemma 2.8(b) and Theorem 7.1. The dependence of the resource bounds on W is due to the fact that the size of the output is $O(W + 1)$. ∎

The proof of Theorem 7.2 shows how to allocate processors to requesting tasks, each of which requests one or more processors. A situation frequently encountered is that many tasks are so small that they do not require "an entire processor", while at the same time the number of tasks is so large that we cannot afford to allocate a processor to each.

Theorem 7.2 easily extends to cover this situation as well. Let $n, \tau \in \mathbb{N}$, assume that we are given a collection $\mathcal{T}_1, \ldots, \mathcal{T}_n$ of n tasks, and suppose that for $j = 1, \ldots, n$, we know an integer $q_j \in \mathbb{N}$ such that $\mathcal{T}_j$ can be executed in $O(q_j)$ time with one processor, or in $O(\tau)$ time with $\lceil q_j/\tau \rceil$ processors. Then the n tasks can be executed in $O(\tau)$ time with $\lceil W/\tau \rceil$ processors, where $W = \sum_{j=1}^{n} q_j \geq n$. To see this, define $\mathcal{T}_j$ to be *small* if $q_j \leq \tau$, and *large* otherwise, for $j = 1, \ldots, n$. Begin by allocating $\lfloor q_j/\tau \rfloor$ processors to $\mathcal{T}_j$, for $j = 1, \ldots, n$, clearly a total of at most W/τ processors. For $j = 1, \ldots, n$, if $\mathcal{T}_j$ is large, then $\lfloor q_j/\tau \rfloor \geq \frac{1}{2}\lceil q_j/\tau \rceil$, so that the processors allocated to $\mathcal{T}_j$ suffice to execute $\mathcal{T}_j$ in $O(\tau)$ time. What remains is to execute the small tasks. Partition these into $m = \lceil n/\tau \rceil$ *groups* $G_1, \ldots, G_m$ of at most τ tasks each. For $i = 1, \ldots, m$, compute the *total length* Q_i of G_i as the sum of the lengths of the tasks in G_i. Then allocate $\lceil Q_i/\tau \rceil$ processors to G_i, for $i = 1, \ldots, m$, a total of at most $m + W/\tau \leq 2\lceil W/\tau \rceil$ processors. It is not difficult to see that using sequential prefix summation, the tasks in each group can be distributed among the processors allocated to the group in such a way that each processor receives tasks of total length $O(\tau)$. All that remains is to let each processor execute the tasks given to it sequentially in $O(\tau)$ time.

When invoking the above principle, we will speak of "operation allocation" rather than processor allocation.

While Theorems 7.1 and 7.2 are our main results concerning the interval allocation and interval marking problems, we also need a more technical lemma (Lemma 7.3 below) that parallels Lemma 6.1 and allows us to perform what we call *incomplete allocation* in constant time. Just as Lemma 6.1 claims efficient deactivation only of elements in well-supplied color classes, those for which the available array is at least $6v$ times larger than the number of elements to be placed there, Lemma 7.3 is wasteful in a sense that we make explicit through the introduction of a so-called *slack parameter.*

Definition: For all $n \in \mathbb{N}$ and $\lambda \geq 1$, an *incomplete interval placement* with *slack* λ for n nonnegative integers $x_1, \ldots, x_n$ is a sequence $y_1, \ldots, y_n$ of n nonnegative integers such that

(1) For $1 \leq i < j \leq n$, if $0 \notin \{y_i, y_j\}$, then $\{y_i, \ldots, y_i + x_i - 1\} \cap \{y_j, \ldots, y_j + x_j - 1\} = \emptyset$;

(2) $\max\{y_j : 1 \leq j \leq n\} = O(\lambda \cdot \sum_{j=1}^{n} x_j)$.

The set $\{j : 1 \leq j \leq n,\ x_j \neq 0 \text{ and } y_j = 0\}$ is called the *residue set* of the incomplete interval placement. If the residue set is empty, the interval placement is *complete.*

Contrasted with the definition of (complete) interval allocation, the above definition does not require a block to be allocated to every request, and blocks may be allocated from a range λ times as large. An algorithm that computes complete interval placements with constant slack performs standard interval allocation.

Lemma 7.3: For all given $n, v, \tau \in \mathbb{N}$, an incomplete interval placement for n given nonnegative integers $x_1, \ldots, x_n$ with slack v can be computed on a TOLERANT PRAM using $O(\tau)$ time, $\lceil (n + v^3 W)/\tau \rceil$ processors and $O(n + v^3 W)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo), where $W = \sum_{j=1}^{n} x_j$, such that the residue set D of the placement satisfies the following:

(a) For $j = 1, \ldots, n$, $\Pr(j \in D) \leq 2^{-v}$;

(b) For every fixed subset R of $\{1, \ldots, n\}$ and for all $z \geq |R|/2^v$,

$$\Pr(|R \cap D| \geq z) \leq 2e^{-z^2/(2^9 |R| v^3)};$$

(c) For every nonnegative real function S of D that satisfies a Lipschitz condition with constant c,

$$S = O(E(S) + cv^3 n^{5/8})$$

with probability at least $1 - 2^{-n^{1/8}}$.

Proof: The reduction of interval allocation to complete linear colored compaction with a logarithmic number of colors extends to incomplete interval allocation and incomplete colored compaction in a straightforward way. Once each request has been marked with its color, an integer in the range $1..m$, we can use the algorithm of Theorem 5.3 to compute a linear overestimate $\widehat{b}_i$ for the size of $\mathcal{B}_i$, for $i = 1, \ldots, m$, and that of Corollary 2.5 to allocate arrays $A_1, \ldots, A_m$ with $|A_i| = 6v\widehat{b}_i$, for $i = 1, \ldots, m$, and their associated blocks. Since $\sum_{i=1}^{m} \widehat{b}_i = O(W)$, both the arrays and the blocks can be allocated from a base array of size $O(vW)$. Lemma 7.3 now follows easily from Lemma 6.1; in particular, note that since the number of nonzero input numbers is bounded by W, the number of operations and memory cells needed is $O(n + v^3 W)$. ∎

Because of the near-equivalence of interval allocation and interval marking, Lemma 7.3 can be used for incomplete allocation of processors as well as of memory cells. Extending our earlier terminology, we call an input element *unlucky* in an application of Lemma 7.3 if its index belongs to the residue set, and *lucky* otherwise. Just as we introduced the concept of v-compaction to facilitate the application of Lemma 6.1, let us agree to use the term "v-allocation", for $v \in \mathbb{N}$, to denote an application of Lemma 7.3 with slack v, followed by the actual allocation of memory cells or processors to the lucky elements. By Lemma 7.3(a), the probability that a fixed element is unlucky in a v-allocation is at most 2^{-v}, and if the "total resource demand" $W = \sum_{j=1}^{n} x_j$ is $O(n/v^3)$, then the v-allocation uses $O(n)$ operations and $O(n)$ space.

8 Coarse-Profiling

While many applications call for the profiling of sequences of values in the range $1..n$ stored in an array of size n, our best strong fine-profiling result (Theorem 5.3)

allows only $O(n^{1-\delta})$ different values, for fixed $\delta > 0$. It is hence necessary to relax the requirements imposed on a profile. Whereas the definition of fine-profiling seems quite natural, it is not obvious how to define a computationally more tractable profiling problem. The following definition of a *coarse-profile*, which at first glance may seem rather artificial, turns out to be useful.

Definition: Let $n, m \in \mathbb{N}$, and let $x_1, \ldots, x_n$ be n integers in the range $0..m$. For $i = 1, \ldots, m$, take $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$. An *$m$-color coarse-profile* for $x_1, \ldots, x_n$ is a sequence of m independent nonnegative integer random variables $\widehat{b}_1, \ldots, \widehat{b}_m$ such that

(A) $\sum_{i=1}^{m} \widehat{b}_i = O(n)$;

(B) For $i = 1, \ldots, m$ and for all $a \geq 1$, $\Pr(b_i > a\widehat{b}_i) \leq 2^{-a}$.

For $n, m \in \mathbb{N}$, the m-color coarse-profiling problem of size n is to compute an m-color coarse-profile for n given integers in the range $0..m$.

We will refer to condition (A) in the above definition as the *linear-sum condition*. As in the case of fine-profiling, input elements of value 0 are "dummy elements" that do not take part in the profiling.

In the following, we show that n-color coarse-profiling problems of size n can be solved with optimal speedup in $O((\log^* n)^2)$ time. We first explain the main ideas in the context of an algorithm that uses n processors and later indicate how to achieve optimal speedup. We begin by tackling a simpler problem, that of estimating just the large multiplicities. Our approach is to extrapolate from a fine-profile for a random sample of size $n^{1-\gamma}$, for some suitably chosen constant $\gamma > 0$. One small complication is that although the sample certainly contains at most $n^{1-\gamma}$ distinct values, these are spread out over the entire range of size n, whereas for an application of Theorem 5.3 a much smaller range is required. The following technical lemma, which will be used frequently in the following, allows us to rename the sample values as required, i.e., to compute an injective function (namely $i \mapsto y_i$) from the set of original sample values to a range of size $O(n^{1-\gamma})$.

Lemma 8.1: For all given $n, d, s, \tau \in \mathbb{N}$ with $d = O(n/\tau)$, the following problem reduces, using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space on a TOLERANT PRAM, to a complete compaction problem of size n and with parameters $d \underset{s}{\longrightarrow} 0$: Given d integers $x_1, \ldots, x_d$ in the range $0..n$, compute n nonnegative integers $y_1, \ldots, y_n$ such that

(1) For $i = 1, \ldots, n$, $y_i \neq 0 \Leftrightarrow i \in \{x_j : 1 \leq j \leq d\}$;

(2) For $1 \leq i < j \leq n$, if $y_i \neq 0$, then $y_i \neq y_j$;

(3) $\max\{y_i : 1 \leq i \leq n\} \leq s$.

Proof: Just as we associate a record X_j with value x_j with each input variable x_j, let us associate an *output record* Y_i with value y_i with the output variable y_i, for $i = 1, \ldots, n$.

The problem is simply to mark those output records that are to receive nonzero values, since afterwards the problem can be solved by compacting the marked records. Whereas the marking would be trivial on the ARBITRARY PRAM, on the TOLERANT PRAM several occurrences of a value i might prevent the marking of Y_i. Our solution is to use guarded writing of a new kind that we call *inverted guarded writing*. For $\tau > 1$ we do not have a (physical) guard processor for each output record, which is the reason why guarded writing of the kind employed in previous sections cannot be used. Instead note that by assumption, we can associate a processor with each element $j \in \{1, \ldots, d\}$. Let this processor continuously write some value to Y_{x_j}. At the same time associate a virtual guard processor with each output record and let $\lceil n/\tau \rceil$ physical processors simulate the virtual guard processors in $O(\tau)$ time. If each virtual guard processor attempts to modify the value stored in its associated output record and marks the record if and only if this fails, the desired marking will result. ∎

In the theorem below, condition (1) says that every nonzero estimate is a linear overestimate for the multiplicity that it estimates. Condition (2) ensures that nonzero estimates are in fact obtained at least for the large multiplicities.

Theorem 8.2: For every fixed $\delta > 0$ there is a constant $\epsilon > 0$ such that the following holds: Let $n, \tau \in \mathbb{N}$ be given, let $x_1, \ldots, x_n$ be n given integers in the range $0..n$ and take $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, n$. Then it is possible, with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo) and on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space, to compute n independent nonnegative integer random variables $\widehat{b}_1, \ldots, \widehat{b}_n$ such that conditions (1) and (2) below hold for some constant $K \geq 1$ and for each $i \in \{1, \ldots, n\}$.

(1) $\widehat{b}_i > 0 \Rightarrow b_i \leq \widehat{b}_i \leq K b_i$;

(2) $b_i \geq n^\delta \Rightarrow \widehat{b}_i > 0$.

Proof: Without loss of generality assume that δ is rational. By Lemma 2.9, we can also assume that $\tau \leq n^{\delta/4}$, so that at least $n^{1-\delta/4}$ processors are available. Let $h = \lceil n^{\delta/4} \rceil$ and carry out the following algorithm:

Step 1: Draw a random sample $\mathcal{Y}$ of $\mathcal{X}$ by including each input element in $\mathcal{Y}$ independently of other elements and with probability $1/h$. For $i = 1, \ldots, n$, let $b_i^{\mathcal{Y}} = |\mathcal{B}_i \cap \mathcal{Y}|$.

Step 2: Use the algorithm of Theorem 5.3 to estimate $b_i^{\mathcal{Y}}$, for $i = 1, \ldots, n$. First place the elements of $\mathcal{Y}$ in an array of size $O(n^{1-\delta/4})$. Since $|\mathcal{Y}| = O(n^{1-\delta/4})$ with high probability by Chernoff bound (a), this can be done by the algorithm of Corollary 4.3. Then use the algorithms of Lemma 8.1 and Corollary 4.3 to replace the values represented among the

elements in $\mathcal{Y}$ by values in a range of size $O(n^{1-\delta/4})$. The algorithm of Theorem 5.3 now provides estimates $\widehat{b}_1^{\mathcal{Y}}, \ldots, \widehat{b}_n^{\mathcal{Y}}$ such that with high probability, $b_i^{\mathcal{Y}} \leq \widehat{b}_i^{\mathcal{Y}} \leq K' b_i^{\mathcal{Y}}$, for $i = 1, \ldots, n$ and for some constant K' (take $\widehat{b}_i^{\mathcal{Y}} = 0$ for all $i \in \{1, \ldots, n\}$ with $b_i^{\mathcal{Y}} = 0$).

Step 3: For $i = 1, \ldots, n$, if $\widehat{b}_i^{\mathcal{Y}} \geq n^{\delta/2}$, then take $\widehat{b}_i = 2h\widehat{b}_i^{\mathcal{Y}}$; otherwise take $\widehat{b}_i = 0$.

In order to analyze the algorithm, fix $i \in \{1, \ldots, n\}$. If $b_i \geq n^{\delta/2}$, then by Lemma 2.1, with high probability $b_i/(2h) \leq b_i^{\mathcal{Y}} \leq 2b_i/h$ and hence $b_i/(2h) \leq \widehat{b}_i^{\mathcal{Y}} \leq 2K'b_i/h$, from which follows that either $\widehat{b}_i = 0$ or $b_i \leq \widehat{b}_i \leq 4K'b_i$. If $b_i \geq n^{\delta}$, clearly with high probability $\widehat{b}_i > 0$. On the other hand, if $b_i < n^{\delta/2}$, then with high probability $b_i^{\mathcal{Y}} < n^{\delta/2}/K'$, hence $\widehat{b}_i^{\mathcal{Y}} < n^{\delta/2}$ and $\widehat{b}_i = 0$. ∎

Lemma 8.3: There is a constant $\epsilon > 0$ such that for all given $n \in \mathbb{N}$, n-color coarse-profiling problems of size n can be solved on a TOLERANT PRAM using $O((\log^* n)^2)$ time, n processors and $O(n)$ space with probability at least $1 - 2^{-n^{\epsilon}}$ (Monte Carlo).

We begin by describing the algorithm informally. As in the previous profiling algorithms, the basic idea is that a GCS for each color can be used to estimate the multiplicity of that color. A fact of prime importance for the design of the algorithm, however, is that a GCS is suitable for estimating multiplicities in a certain range only: Very large color classes are likely to fill every row of their respective graduated conditional scatterings, which therefore offer no means of distinguishing between them. For a GCS with parameters $r \times s$, this happens for multiplicities above approximately 2^r, so that 2^r is the upper limit of multiplicities that can be estimated. On the other hand, color classes of sizes significantly below s obviously leave all rows practically empty so that, again, it is not possible to distinguish between them. The range of multiplicities that can be meaningfully estimated using a GCS with parameters $r \times s$ therefore is roughly from s to 2^r. As can be seen from Lemmas 3.3 and 5.2, increasing s has the effect of increasing the reliability of the estimate obtained from the GCS. This, however, turns out not to be needed for the algorithm discussed here, so that we can fix s at a constant value, which, for proof-technical reasons, we choose to be 6; this avoids the exclusion of small multiplicities. A more significant trade-off concerns the parameter r. Increasing r extends the range of multiplicities that can be estimated, but obviously also increases the amount of space needed for the GCS. In order to estimate the size of a large color class, we therefore have to allocate a large amount of space for its GCS; on the other hand, we plan to get by with $O(n)$ space altogether, which prevents us from allocating so much space to every color class, independently of its size. In a sense, therefore, we need an estimate of a multiplicity before we can use a GCS to estimate it, a seemingly hopeless situation. What makes our day is the exponential difference between the space needed by a GCS with parameters $r \times 6$, which is $O(r)$, and the upper limit of the range

of multiplicities that it can estimate, which is roughly 2^r, together with the fact that a color class whose size is too large to be estimated at least bears witness to this fact by filling every row of its GCS.

We begin by carrying out a GCS with parameters $r_1 \times 6$ for each of the n colors, where r_1 is a constant. This gives us an estimate of the size of each color class that does not fill every row of its GCS. No estimate is obtained for the remaining color classes, but most of these can be assumed to be of size at least $r_2 \approx 2^{r_1}$. We therefore subsequently carry out a GCS with parameters $r_2 \times 6$ for each "unresolved" color, without violating the $O(n)$ space bound. This allows us to estimate multiplicities up to $r_3 \approx 2^{r_2}$, and we continue in the same manner. The resulting profiling algorithm works in $O(\log^* n)$ stages. A rather intricate analysis is needed to show that the output of the algorithm is indeed a coarse-profile. Note also that we have carefully matched the definition of a coarse-profile to what the algorithm actually produces.

We now provide a more formal description of the algorithm and begin by defining an integer sequence $v_1, \ldots, v_T$. Briefly let $f(z) = z - 4\log z$, for all $z > 0$. Then $f(2) < 0$, $f(16) = 0$, $f(z) \to \infty$ for $z \to \infty$, and f' has only one zero. It follows that $z \geq 4\log z$ for all $z \geq 16$. In particular, since $\lceil \log z \rceil < 4\log z$ for $z \geq 2$, the relation $0 \leq \lceil \log z \rceil < z$ can be seen to hold for all $z \in \mathbb{N}$. Now consider the algorithm

$z := \lceil n^{1/20} \rceil$;
repeat
  write(z);
  $z := \lceil \log z \rceil$;
until $z = 0$;

The last relation derived above clearly implies that the output of the algorithm is a finite sequence $v_T, \ldots, v_1$ (note the reverse indexing) with $1 = v_1 < v_2 < \cdots < v_T = \lceil n^{1/20} \rceil$. Since $\lceil \log \lceil \log z \rceil \rceil \leq \log(2\log z) + 1 = \log(4\log z) \leq \log z$ for $z \geq 16$, where the last relation was derived above, we certainly have $T \leq 2\log^* n + 16 = O(\log^* n)$. The lemma below lists some basic properties of the sequence $v_1, \ldots, v_T$.

Lemma 8.4:

(a) For $t = 1, \ldots, T-1$, $v_{t+1} \leq 2^{v_t} < 2v_{t+1}$;

(b) For $t = 5, \ldots, T-1$, $v_{t+1} \geq (2v_t)^2$;

(c) For $t = 1, \ldots, T$, $v_t \geq 2^{t-2}$.

Proof:

(a) $\log v_{t+1} \leq \lceil \log v_{t+1} \rceil = v_t < \log v_{t+1} + 1$ and hence $v_{t+1} \leq 2^{v_t} < 2v_{t+1}$.

(b) We noted above that $2^z \geq z^4$ for all $z \geq 16$. Hence if $v_t \geq 16$ for some $t \leq T-1$, then $v_{t+1} \geq \frac{1}{2} \cdot 2^{v_t} \geq \frac{1}{2} \cdot (v_t)^4 \geq (2v_t)^2 \geq 16$. Repeated use of part (a) shows that $v_2 \geq 2$,

$v_3 \geq 3$, $v_4 \geq 5$ and $v_5 \geq 17$. By induction, the relation $v_{t+1} \geq (2v_t)^2$ therefore holds for $t = 5, \ldots, T-1$.

(c) The claim can be seen to hold for $t \leq 5$. For $t \geq 6$ it follows from part (b) by induction. ∎

Theorem 8.2 allows us to obtain accurate estimates for all color classes of size $n^{1/8}$ or more, whose elements can subsequently be replaced by "dummy elements" of value 0 (see the discussion preceding the definition of a fine-profile in Section 5). We can therefore assume without loss of generality that $b_i \leq n^{1/8}$, for $i = 1, \ldots, n$. Under this assumption, the algorithm first computes the sequence $v_1, \ldots, v_T$, which can clearly be done in $O(\log^* n)$ time (it can actually be done in constant time), and then executes the following steps, where $K = 144$.

```
(1)  for i ∈ {1,...,n} pardo let B_i be active;
(2)  for t := 1 to T do (* Stage t *)
(3)    for each active color class B_i pardo
(4)      begin
(5)        Allocate a 5v_t × 6 array and 30v_t processors to B_i;
(6)        Let the elements of B_i carry out a GCS S_i
(7)          with parameters 5v_t × 6;
(8)        l_i := the 1-row of S_i;
(9)        if l_i < 5v_t (* not entirely full *) then
(10)         begin
(11)           Make B_i inactive;
(12)           b̂_i := min{K · 2^{l_i}, ⌊n^{1/8}⌋};
(13)         end;
(14)     end;
```

For $t = 1, \ldots, T$, let *Stage* t be the tth execution of lines (3)–(14). For $i = 1, \ldots, n$, if a value is assigned to $\widehat{b}_i$ in line (12) in Stage t, for some $t \in \{1, \ldots, T\}$, we shall say that $\widehat{b}_i$ *becomes defined* in Stage t. Observe that the value assigned to $\widehat{b}_i$ is the minimum of what might, in the light of Lemma 3.3, be regarded as a reasonable estimate of b_i (namely $K \cdot 2^{l_i}$), and a value no smaller than b_i (namely $\lfloor n^{1/8} \rfloor$). Forming the minimum with $\lfloor n^{1/8} \rfloor$ clearly does not make any estimate less accurate; it is necessary to prevent an occasional violation of the linear-sum condition (the condition can still be violated, but the probability of this event is negligible).

Each allocation of processors and space in line (5) can be done in $O(\log^* n)$ time using the algorithms of Theorems 7.1 and 7.2. Provided that the allocated resources stay within the limits imposed by Lemma 8.3, it is not difficult to see that the whole algorithm

can be executed within the time, processor and space bounds stated in Lemma 8.3 (use Lemma 3.4(a)). The lemmas below show that the resources allocated are not excessive and that the sequence $\widehat{b}_1, \ldots, \widehat{b}_n$ is indeed a coarse-profile with the required probability.

Lemma 8.5: For $i = 1, \ldots, n$, with high probability $\widehat{b}_i$ becomes defined in some stage.

Proof: Fix $i \in \{1, \ldots, n\}$. If $\widehat{b}_i$ does not become defined in any stage, $l_i = 5v_T$ in Stage T. By Lemma 3.3(a), the probability of this is at most $(2^{-5v_T} en/6)^6$, i.e., negligible. ∎

In order to avoid undefined symbols in what follows, we will assume that $\widehat{b}_i = 0$ in the unlikely event that $\widehat{b}_i$ does not become defined in any stage, for $i = 1, \ldots, n$. The algorithm could easily be modified to obey this convention, but this would serve little purpose.

A *tardy* color class $\mathcal{B}_i$, defined below, intuitively is one that "should" have been deactivated in an earlier stage. Since the size of a color class may be highly overestimated if the color class becomes tardy, we must show this to be an unlikely event.

Definition: For $i = 1, \ldots, n$ and $t = 1, \ldots, T$, call the color class $\mathcal{B}_i$ *tardy* in Stage t if it is still active at the beginning of Stage t and $b_i < v_t^4$.

Lemma 8.6: For $i = 1, \ldots, n$ and $t = 2, \ldots, T$, $\Pr(\mathcal{B}_i \text{ is tardy in Stage } t) \leq v_t^{-6}$.

Proof: If $\mathcal{B}_i$ is tardy in Stage t, then $l_i = 5v_{t-1}$ in Stage $t-1$. By Lemmas 3.3(a) and 8.4(a), the probability of this is at most

$$\left(\frac{2^{-5v_{t-1}} e v_t^4}{6}\right)^6 \leq (v_t^{-5} v_t^4)^6 = v_t^{-6}. \quad \blacksquare$$

Lemma 8.7: With high probability, the algorithm uses $O(n)$ processors and $O(n)$ space.

Proof: We need $O(n)$ space for storing the input elements and color classes and their attributes and $O(n)$ processors standing by the input elements and color classes. The remaining resources used are those explicitly allocated in line (5) in the algorithm. Since the number of processors allocated equals the number of memory cells allocated, it suffices to bound the latter. Furthermore, by a martingale argument, it suffices to show that the expected amount of space allocated to $\mathcal{B}_i$ in Stage t is $O(b_i + 1)$, for arbitrary fixed $i \in \{1, \ldots, n\}$ and $t \in \{1, \ldots, T\}$ (recall that $v_t \leq \lceil n^{1/20} \rceil$). But if an array is allocated to $\mathcal{B}_i$ in Stage t (i.e., if $\mathcal{B}_i$ is active at the beginning of Stage t), it is of size $30v_t$, which is $O(1)$ if $t = 1$, and $O(b_i)$ if $\mathcal{B}_i$ is not tardy in Stage t. By Lemma 8.6, finally, $\mathcal{B}_i$ is tardy in Stage t with probability at most v_t^{-6} if $t \geq 2$, so that this case adds only $O(1)$ to the overall expected size. ∎

What remains is to demonstrate that the output of the algorithm satisfies conditions (A) and (B) in the definition of a coarse-profile. To this end we must show that $\widehat{b}_1, \ldots, \widehat{b}_n$ are not too large (Lemmas 8.8 and 8.9), and that they are not too small (Lemma 8.10).

Lemma 8.8: For all $i \in \{1, \ldots, n\}$ with $b_i > 0$ and for all $a \geq 1$, $\Pr(\widehat{b}_i > ab_i) \leq (K/a)^6$.

Proof: It suffices to show for all $t \in \{1, \ldots, T\}$ with $\Pr(D_t) > 0$, where D_t denotes the event that $\widehat{b}_i$ becomes defined in Stage t, that $\Pr(\widehat{b}_i > ab_i \mid D_t) \leq (K/a)^6$. To this end let $D_{\geq t}$ be the event that $\mathcal{B}_i$ is active at the beginning of Stage t, take $D_{>t} = D_{\geq t} \backslash D_t$ and denote by H the subevent of $D_{\geq t}$ in which $K \cdot 2^{l_i} > ab_i$ in Stage t. We are interested in the quantity $\Pr(H \mid D_t)$. The central fact to note is that either $H = \emptyset$, in which case there is nothing to show, or else $D_{>t} \subseteq H$. In the latter case

$$P(H \mid D_t) = \frac{\Pr(H \cap D_t)}{\Pr(D_t)} \leq \frac{\Pr(H \cap D_t) + \Pr(H \cap D_{>t})}{\Pr(D_t) + \Pr(H \cap D_{>t})} = \frac{\Pr(H \cap D_{\geq t})}{\Pr(D_{\geq t})} = \Pr(H \mid D_{\geq t}),$$

and since we can clearly assume that $ab_i \geq K \geq 6$, Lemma 3.3(b) directly implies that the right-hand probability is at most $(2eK/(6a))^6 \leq (K/a)^6$. ∎

Lemma 8.9: With high probability, $\sum_{i=1}^{n} \widehat{b}_i = O(n)$.

Proof: Fix $i \in \{1, \ldots, n\}$. Since $\widehat{b}_i \leq n^{1/8}$ for all $i \in \{1, \ldots, n\}$, by a martingale argument it suffices to show that $E(\widehat{b}_i) = O(b_i + 1)$. If $b_i = 0$, it is easy to see that $\widehat{b}_i = K$ (since all rows of any GCS of $\mathcal{B}_i$ are empty). Otherwise use Lemma 8.8 to find that

$$\begin{aligned}
E(\widehat{b}_i) &= \sum_{j=1}^{\infty} j \Pr(\widehat{b}_i = j) \\
&= \sum_{j=1}^{Kb_i} j \Pr(\widehat{b}_i = j) + \sum_{j=Kb_i+1}^{\infty} j \Pr(\widehat{b}_i = j) \\
&\leq Kb_i + \sum_{l=0}^{\infty} 2^{l+1} Kb_i \cdot \Pr(2^l Kb_i < \widehat{b}_i \leq 2^{l+1} Kb_i) \\
&\leq Kb_i + \sum_{l=0}^{\infty} 2^{l+1} Kb_i \cdot 2^{-6l} = O(b_i). \quad \blacksquare
\end{aligned}$$

In Lemma 8.10 below, the condition $a\widehat{b}_i > 0$ ensures precisely that $\widehat{b}_i$ becomes defined in some stage. By Lemma 8.5, the difference between $\Pr(b_i > a\widehat{b}_i)$ and $\Pr(b_i > a\widehat{b}_i > 0)$ is negligible, so that it suffices to bound the latter.

Lemma 8.10: For $i = 1, \ldots, n$ and for all $a \geq 1$, $\Pr(b_i > a\widehat{b}_i > 0) \leq 2^{-a}$.

Proof: Fix $i \in \{1, \ldots, n\}$ and for $t = 1, \ldots, T$, denote by D_t the event that $\widehat{b}_i$ becomes defined in Stage t. Without loss of generality assume that $b_i > 0$. We show that

$\sum_{t=1}^{T} \Pr(b_i > a\widehat{b}_i \text{ and } D_t) \leq 2^{-a}$. The analysis proceeds stage by stage and is divided into three cases depending on whether b_i is "too large", "of suitable size" or "too small" for being estimated in the stage under consideration. We show that the probability that $\widehat{b}_i$ becomes defined at all in a stage in which b_i is either "too large" or "too small" is negligible, and that the probability that it receives a far too small value in a stage in which b_i is "of suitable size" is sufficiently small per stage to sum to a negligible quantity over the relatively few stages of this kind. Let $t_2 = \min\{t : 1 \leq t \leq T \text{ and } 2^{9v_t} > b_i\}$ (note that the minimum is not formed over an empty set), $t_1 = \max\{1, t_2 - \lceil a \rceil - 1\}$, $t_3 = \min\{T, t_2 + \lceil a \rceil + 5\}$, $I_1 = \{t \in \mathbb{N} : 1 \leq t < t_1\}$, $I_2 = \{t \in \mathbb{N} : t_1 \leq t \leq t_3\}$ and $I_3 = \{t \in \mathbb{N} : t_3 < t \leq T\}$ and note that $I_1 \cup I_2 \cup I_3 = \{1, \ldots, T\}$.

Case 1 ("b_i is too large"): Assume that $I_1 \neq \emptyset$ and hence that $t_1 = t_2 - \lceil a \rceil - 1$. Lemma 3.3(d) implies that for $t = 1, \ldots, T$,

$$\Pr(D_t) \leq 6 \cdot 2^{1 - b_i/(12 \cdot 2^{5v_t})} \leq 12 \cdot 2^{-b_i \cdot 2^{-9v_t}}.$$

Hence by Lemma 3.2, $\sum_{t \in I_1} \Pr(D_t) \leq 24 \cdot 2^{-b_i \cdot 2^{-9v_{t_1}}}$. But by the choice of t_2, we have that $b_i \cdot 2^{-9v_{t_2-1}} \geq 1$ and therefore

$$b_i \cdot 2^{-9v_{t_1}} \geq b_i \cdot 2^{-9(v_{t_2-1} - \lceil a \rceil)} \geq 2^{9a} \geq 9a.$$

Thus $\sum_{t \in I_1} \Pr(D_t) \leq 24 \cdot 2^{-9a} \leq 2^{-4a}$.

Case 2 ("b_i is of suitable size"): By Lemma 3.3(c),

$$\Pr(b_i > a\widehat{b}_i \text{ and } D_t) \leq 6 \cdot 2^{1 - Ka/12} = 12 \cdot 2^{-12a} \leq 2^{-8a},$$

for $t = 1, \ldots, T$, and hence

$$\sum_{t \in I_2} \Pr(b_i > a\widehat{b}_i \text{ and } D_t) \leq |I_2| \cdot 2^{-8a} \leq (2\lceil a \rceil + 7) \cdot 2^{-8a} \leq 11a \cdot 2^{-8a} \leq 2^{-3a}.$$

Case 3 ("b_i is too small"): Assume that $I_3 \neq \emptyset$ and hence that $t_3 = t_2 + \lceil a \rceil + 5$. By Lemma 8.4 and since $t_2 \geq 1$, $b_i \leq 2^{9v_{t_2}} \leq (2v_{t_2+1})^9 < (2v_{t_2+4})^9 \leq v_{t_2+5}^5 \leq v_{t_2+6}^3 \leq v_{t_3}^4$, i.e., $\mathcal{B}_i$ is tardy in Stage t for all $t \in I_3$. Hence by Lemma 8.6, $\sum_{t \in I_3} \Pr(D_t) \leq \sum_{t \in I_3} v_t^{-6} \leq 2 \cdot v_{t_3}^{-6}$. By Lemma 8.4(c), $v_{t_3} \geq 2^{t_3 - 2} \geq 2^{a+4}$ and therefore $\sum_{t \in I_3} \Pr(D_t) \leq 2^{-6a}$.

The contributions of Cases 1–3 add up to at most $2^{-4a} + 2^{-3a} + 2^{-6a} \leq 2^{-2a} \leq 2^{-a}$. This ends the proofs of Lemmas 8.10 and 8.3. ∎

We now describe a simple procedure called *scattering in time* that will be used on five separate occasions. Since in each case we shall need different properties of the procedure, we believe that it serves little purpose to list at this point all the properties of the procedure that we shall ever need. After describing the procedure, we therefore

analyze it only with respect to its resource requirements; later we will refer back to the procedure and derive whatever properties are of interest. In four of the five cases, scattering in time is used as a "profile enhancer", i.e., informally, it inputs a profile $\widehat{b}_1, \ldots, \widehat{b}_n$ and produces a "better" profile $\widetilde{b}_1, \ldots, \widetilde{b}_n$. In the fourth case we input a very good profile and use it to semisort.

Scattering in time takes as input n input elements $X_1, \ldots, X_n$ (the *primary input*) with values in the range $0..n$ (input elements with a value of 0 are dummy elements signifying "no element") and n nonnegative integers $\widehat{b}_1, \ldots, \widehat{b}_n$ (the *profile input*) with $\sum_{i=1}^{n} \widehat{b}_i = O(n)$, as well as an integer τ (the *phase count*) with $1 \leq \tau \leq \sqrt{n}$. As usual, let $\mathcal{B}_i$ be the set of input elements of value i, for $i = 1, \ldots, n$. We begin by using the algorithm of Theorem 7.1 to allocate an array A_i of size $\widehat{b}_i$ to $\mathcal{B}_i$, for $i = 1, \ldots, n$, each cell of which contains a counter, initialized to zero, and a list header, initially denoting an empty list. For $i = 1, \ldots, n$, every element of $\mathcal{B}_i$ now chooses a random integer, called its *list number*, from the set $\{1, \ldots, \widehat{b}_i\}$, and another random integer, called its *phase number*, from the set $\{1, \ldots, \tau\}$. By Chernoff bound (a), with high probability the set $\mathcal{Y}_l$ of elements with phase number l is of size $O(n/\tau)$, for $l = 1, \ldots, \tau$. Using the algorithm of Theorem 6.2 if $\tau \leq \log n$ and that of Lemma 2.9 if $\tau > \log n$, we can therefore store the elements of $\mathcal{Y}_l$ in an array Q_l of size $O(n/\tau)$, for $l = 1, \ldots, \tau$.

The arrays $Q_1, \ldots, Q_\tau$ are next processed one by one. To process an array Q_l means to associate a processor with each cell in Q_l, and then to process all cells in Q_l simultaneously and in constant time. If a cell of Q_l is empty or contains a dummy element, the processor in charge of that cell does nothing. Otherwise, suppose that the element X stored in the cell belongs to $\mathcal{B}_i$ and chose j as its list number. The processor then attempts to increment the counter stored in $A_i[j]$ by 1. If this fails, i.e., if some other processor attempts to increment the same counter in the same time step, X is said to *collide*. Otherwise X is *noncolliding*, and the processor in charge of X proceeds to insert X in the list whose header is stored in $A_i[j]$. An important fact to note is that an element X of $\mathcal{B}_i$ collides exactly if some other element of $\mathcal{B}_i$ chooses both the same list number and the same phase number as X, for $i = 1, \ldots, n$, i.e., the scattering in time of $\mathcal{B}_i$ can be analyzed as a 1-scattering of $\mathcal{B}_i$ over an array of $\tau\widehat{b}_i$ cells; the advantage of scattering in time is that it uses less space and (therefore) fewer operations.

Once a first pass as described above has been completed, we shall sometimes carry out a second pass exactly like the first pass, except that the only elements taking part in the computation are those that collided in the first pass. In either case, we say that an element is *successful* if and only if there is some pass in which it does not collide. It is easy to see that one-pass or two-pass scattering in time can be carried out with high probability using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space. After the scattering in time, the set $\mathcal{X}'$ of successful elements as well as all counters and lists are available for

further processing. In particular, note that for all $i \in \{1, \ldots, n\}$ such that $\widehat{b}_i = \tau^{O(1)}$, we can compute $|\mathcal{B}_i \cap \mathcal{X}'|$ exactly in $O(\tau)$ time using operation allocation, as described in Section 7. It suffices to observe that if $\widehat{b}_i = \tau^{O(1)}$, Lemma 2.4 can be used to compute the sum of the counters stored in A_i, either in $O(\widehat{b}_i)$ time with one processor, or in $O(\tau)$ time with $\lceil \widehat{b}_i/\tau \rceil$ processors. Since $\sum_{i=1}^{n} \widehat{b}_i = O(n)$, the resource requirements remain as stated above.

Our first application of scattering in time is to the computation of a profile with the somewhat unnatural properties listed in Lemma 8.11 below. Once Lemma 8.11 has been established, a second application of scattering in time will allow us to obtain a coarse-profiling algorithm with optimal speedup. The proof of Lemma 8.11 is rather technical, but the main ideas behind it are as follows: We already have an n-processor coarse-profiling algorithm (Lemma 8.3). In the context of an algorithm with optimal speedup and a running time of $\Theta(\tau)$, we can allow ourselves to apply this nonoptimal algorithm to a random sample of the input set of size $\Theta(n/\tau)$. It turns out that this yields suitable estimates of multiplicities somewhat larger than τ, say, at least $\tau^{3/2}$. On the other hand, very small color classes are likely not to be represented in the sample at all, so that their sizes must be estimated in a different way. We here use the scattering in time described above, which enables us to estimate multiplicities up to roughly τ. Finally, in order to bridge the gap between τ and $\tau^{3/2}$, we use another scattering in time, but this time applied to a random sample of the input set of size $\Theta(n/\sqrt{\tau})$. The complete algorithm hence consists of three essentially independent subalgorithms, each of which "caters to" a different range of multiplicities.

Lemma 8.11: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$ with $\tau \geq (\log^* n)^2$, the following problem can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo): Given n integers $x_1, \ldots, x_n$ in the range $0..n$, compute independent nonnegative integer random variables $\widehat{b}_1, \ldots, \widehat{b}_n$ such that

(A) $\sum_{i=1}^{n} \widehat{b}_i = O(n)$;

(B) For $i = 1, \ldots, n$ and for all $a \geq 1$, $\Pr(b_i > a\widehat{b}_i) \leq 2^{-b_i/(8\tau)} + 2^{-2a}$;

(C) For $i = 1, \ldots, n$, $\Pr(b_i > \sqrt{\tau}\widehat{b}_i) \leq 2^{-\sqrt{\tau}}$,

where $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, n$.

Proof: Without loss of generality, we can assume that $2 \leq \tau \leq n^{1/2}$, since otherwise the problem is easily solved using the algorithm of Lemma 2.9. Let $K = 2^8$ and carry out the following algorithm:

Step 1: Apply one-pass scattering in time with phase count τ to the primary input $X_1, \ldots, X_n$ and the (trivial) profile input $8K, \ldots, 8K$ and let $\mathcal{X}'$ be the resulting set of noncolliding elements. For $i = 1, \ldots, n$, take $\widehat{b}_i^{(1)} = 4|\mathcal{B}_i \cap \mathcal{X}'|$ (since $8K = \tau^{O(1)}$, we

argued above that this quantity is readily available).

Step 2: Draw a random sample $\mathcal{Y} \subseteq \mathcal{X}$ by including each input element in $\mathcal{Y}$ independently of other elements and with probability $1/\lceil\sqrt{\tau}\rceil$. Repeat Step 1, but this time include only elements of $\mathcal{Y}$ in the primary input (i.e., replace each element not in $\mathcal{Y}$ by a dummy element with a value of 0), and let $\mathcal{Y}'$ be the resulting set of noncolliding elements. For $i = 1, \ldots, n$, take $\widehat{b}_i^{(2)} = 8\lceil\sqrt{\tau}\rceil|\mathcal{B}_i \cap \mathcal{Y}'|$.

Step 3: Draw a random sample $\mathcal{Z} \subseteq \mathcal{X}$ by including each input element in $\mathcal{Z}$ independently of other elements and with probability $1/\tau$. By Chernoff bound (a), we can assume that $|\mathcal{Z}| = O(n/\tau)$. Use the algorithms of Lemma 8.1 and Theorem 4.6 to store $\mathcal{Z}$ in an array of size $O(n/\tau)$ and to replace the values of the elements in $\mathcal{Z}$ by values in a range of size $O(n/\tau)$. Then apply the algorithm of Lemma 8.3 to $\mathcal{Z}$ to obtain a profile $\widehat{b}_1^{\mathcal{Z}}, \ldots, \widehat{b}_n^{\mathcal{Z}}$ (take $\widehat{b}_i^{\mathcal{Z}} = 0$ for each $i \in \{1, \ldots, n\}$ with $\mathcal{B}_i \cap \mathcal{Z} = \emptyset$). For $i = 1, \ldots, n$, let $\widehat{b}_i^{(3)} = 4\tau\widehat{b}_i^{\mathcal{Z}}$.

Step 4: For $i = 1, \ldots, n$, compute the final estimate of b_i as $\widehat{b}_i = \max\{\widehat{b}_i^{(1)}, \widehat{b}_i^{(2)}, \widehat{b}_i^{(3)}, K\}$.

For $i = 1, \ldots, n$, let $b_i^{\mathcal{Y}} = |\mathcal{B}_i \cap \mathcal{Y}|$ and $b_i^{\mathcal{Z}} = |\mathcal{B}_i \cap \mathcal{Z}|$. We can assume that $\widehat{b}_1^{\mathcal{Z}}, \ldots, \widehat{b}_n^{\mathcal{Z}}$ is indeed a coarse-profile for (the sequence of values of elements in) $\mathcal{Z}$. It is easy to see that with high probability, the resource requirements of the algorithm are as stated in the lemma. The correctness of the algorithm is demonstrated in the lemmas below, each of which shows one of the properties (A)–(C).

Lemma 8.12: With high probability, $\sum_{i=1}^n \widehat{b}_i = O(n)$.

Proof: $\sum_{i=1}^n \widehat{b}_i^{(1)} = 4\sum_{i=1}^n |\mathcal{B}_i \cap \mathcal{X}'| \leq 4\sum_{i=1}^n b_i = 4n$. In the same way, $\sum_{i=1}^n \widehat{b}_i^{(2)} \leq 8\lceil\sqrt{\tau}\rceil|\mathcal{Y}|$, and $|\mathcal{Y}| = O(n/\sqrt{\tau})$ with high probability by Chernoff bound (a). Finally, by the linear-sum condition, $\sum_{i=1}^n \widehat{b}_i^{(3)} = 4\tau\sum_{i=1}^n \widehat{b}_i^{\mathcal{Z}} = O(\tau|\mathcal{Z}|) = O(n)$. ∎

Lemma 8.13: For $i = 1, \ldots, n$ and for all $a \geq 1$, $\Pr(b_i > a\widehat{b}_i) \leq 2^{-b_i/(8\tau)} + 2^{-2a}$.

Proof: Clearly, $b_i^{\mathcal{Z}}$ is binomially distributed with expected value b_i/τ. Hence by Chernoff bound (b), $\Pr(b_i^{\mathcal{Z}} < b_i/(2\tau)) \leq 2^{-b_i/(8\tau)}$. Furthermore, by property (B) of a coarse-profile, $\Pr(b_i^{\mathcal{Z}} > 2a\widehat{b}_i^{\mathcal{Z}}) \leq 2^{-2a}$. But $b_i^{\mathcal{Z}} \geq b_i/(2\tau)$ and $b_i^{\mathcal{Z}} \leq 2a\widehat{b}_i^{\mathcal{Z}}$ together imply $b_i \leq 2\tau b_i^{\mathcal{Z}} \leq 4a\tau\widehat{b}_i^{\mathcal{Z}} = a\widehat{b}_i^{(3)}$. Hence $\Pr(b_i > a\widehat{b}_i) \leq \Pr(b_i > a\widehat{b}_i^{(3)}) \leq 2^{-b_i/(8\tau)} + 2^{-2a}$. ∎

Lemma 8.14: For $i = 1, \ldots, n$, $\Pr(b_i > \sqrt{\tau}\widehat{b}_i) \leq 2^{-\sqrt{\tau}}$.

Proof: Without loss of generality assume that $b_i > K\sqrt{\tau}$. By the definition of $\widehat{b}_i$, if $b_i > \sqrt{\tau}\widehat{b}_i$, then $b_i > \sqrt{\tau}\widehat{b}_i^{(l)}$, for $l = 1, 2, 3$, so that for each $i \in \{1, \ldots, n\}$ we can show the event $b_i > \sqrt{\tau}\widehat{b}_i$ to be unlikely in any of three ways. Correspondingly, we consider three cases. If $\tau < 4$, Case 2 disappears, and Cases 1 and 3 overlap; the argument remains valid, however.

Case 1: $b_i \le 2K\tau$. If $b_i > \sqrt{\tau}\widehat{b}_i^{(1)}$, then $|\mathcal{B}_i \cap \mathcal{X}'| = \widehat{b}_i^{(1)}/4 \le b_i/(4\sqrt{\tau}) \le b_i/2$, which implies that at least $b_i/2$ elements of $\mathcal{B}_i$ collide in Step 1. But since $b_i \le 2K\tau$, the density of the scattering in time of $\mathcal{B}_i$ is $b_i/(8K\tau) \le 1/4$; hence, by Lemma 3.6(b), the probability that at least $b_i/2$ elements of $\mathcal{B}_i$ collide is bounded by $e^{-\zeta}$, where $\zeta = (b_i/2)^2/(32b_i \cdot 1) = b_i/2^7 \ge K\sqrt{\tau}/2^7 = 2\sqrt{\tau}$, from which the claim follows. Note that we actually showed the stronger relation $\Pr(b_i \ge 4\sqrt{\tau}|\mathcal{B}_i \cap \mathcal{X}'|) \le 2^{-2\sqrt{\tau}}$, which will be used in Case 2 below.

Case 2: $2K\tau < b_i \le K\tau^{3/2}$. If $b_i > \sqrt{\tau}\widehat{b}_i^{(2)}$, then $|\mathcal{B}_i \cap \mathcal{Y}'| = \widehat{b}_i^{(2)}/(8\lceil\sqrt{\tau}\rceil) \le b_i/(\sqrt{\tau} \cdot 8\lceil\sqrt{\tau}\rceil)$, which happens only if either $b_i^{\mathcal{Y}} \le b_i/(2\lceil\sqrt{\tau}\rceil)$ (the sample is small) or $|\mathcal{B}_i \cap \mathcal{Y}'| \le b_i^{\mathcal{Y}}/(4\sqrt{\tau})$ (many elements collide). Since $b_i^{\mathcal{Y}}$ is binomially distributed with expected value $b_i/\lceil\sqrt{\tau}\rceil$, Chernoff bound (b) implies that $\Pr(b_i^{\mathcal{Y}} \le b_i/(2\lceil\sqrt{\tau}\rceil)) \le e^{-b_i/(8\lceil\sqrt{\tau}\rceil)}$. On the other hand, we know from Case 1 that $\Pr(b_i^{\mathcal{Y}} \ge 4\sqrt{\tau}|\mathcal{B}_i \cap \mathcal{Y}'| \mid b_i^{\mathcal{Y}} \le 2K\tau) \le 2^{-2\sqrt{\tau}}$, and by Chernoff bound (a), $\Pr(b_i^{\mathcal{Y}} > 2K\tau) \le \Pr(b_i^{\mathcal{Y}} > 2b_i/\lceil\sqrt{\tau}\rceil) \le e^{-b_i/(3\lceil\sqrt{\tau}\rceil)}$. Using that $b_i/\lceil\sqrt{\tau}\rceil \ge (2K\tau)/(2\sqrt{\tau}) = K\sqrt{\tau}$, we finally obtain that $\Pr(b_i > \sqrt{\tau}\widehat{b}_i^{(2)}) \le e^{-b_i/(8\lceil\sqrt{\tau}\rceil)} + 2^{-2\sqrt{\tau}} + e^{-b_i/(3\lceil\sqrt{\tau}\rceil)} \le 2^{-K\sqrt{\tau}/8} + 2^{-2\sqrt{\tau}} + 2^{-K\sqrt{\tau}/3} \le 2^{-\sqrt{\tau}}$.

Case 3: $b_i > K\tau^{3/2}$. By Lemma 8.13, $\Pr(b_i > \sqrt{\tau}\widehat{b}_i) \le 2^{-b_i/(8\tau)} + 2^{-2\sqrt{\tau}} \le 2^{-K\sqrt{\tau}/8} + 2^{-2\sqrt{\tau}} \le 2^{-\sqrt{\tau}}$. ∎

Theorem 8.15: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$ with $\tau \ge (\log^* n)^2$, n-color coarse-profiling problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo).

Proof: Assume without loss of generality that $\tau \le n^{1/4}$ and that $b_i \le n^{1/8}$, for $i = 1, \ldots, n$. Begin by computing a profile $\widehat{b}_1, \ldots, \widehat{b}_n$ for $x_1, \ldots, x_n$ with the properties described in Lemma 8.11. Then apply two-pass scattering in time with phase count τ to the primary input $X_1, \ldots, X_n$ and the profile input $\widehat{b}_1, \ldots, \widehat{b}_n$. For $i = 1, \ldots, n$, call $\mathcal{B}_i$ *well-estimated* if $b_i \le \sqrt{\tau}\widehat{b}_i$, and call each element of $\mathcal{B}_i$ *good* if $\mathcal{B}_i$ is well-estimated, and *bad* otherwise. By property (C) of Lemma 8.11, the expected number of bad elements in $\mathcal{B}$ is at most $n \cdot 2^{-\sqrt{\tau}}$, and a martingale argument shows the actual number of bad elements to be $O(n/\tau + n^{3/4}) = O(n/\tau)$ with high probability. It is easy to see that the probability that a good element collides in one pass of a scattering in time is at most $1/\sqrt{\tau}$. Since we actually execute two passes, the probability that a good element is unsuccessful is at most $1/\tau$, so a martingale argument shows that the number of (good or bad) unsuccessful elements is $O(n/\tau)$ with high probability. Let $\mathcal{X}'$ be the set of successful elements.

Consider the situation after the scattering in time. For $i = 1, \ldots, n$, call $\mathcal{B}_i$ *resolved* if every element of $\mathcal{B}_i$ was successful in the scattering in time, and $\widehat{b}_i \le 16\tau$. The total number of unsuccessful elements being $O(n/\tau)$, we can use inverted guarded writing as

in the proof of Lemma 8.1 to determine the set of resolved color classes. The important observation is that if $\mathcal{B}_i$ is resolved, then we can compute b_i exactly as $|\mathcal{B}_i \cap \mathcal{X}'|$, for $i = 1, \ldots, n$. Hence for $i = 1, \ldots, n$, do the following: If $\mathcal{B}_i$ is resolved, replace (the estimate) $\widehat{b}_i$ by (the exact value) b_i; otherwise replace $\widehat{b}_i$ by $\max\{\widehat{b}_i, 16\tau\}$. Since all except $O(n/\tau)$ color classes are resolved, these changes preserve the linear-sum condition.

We must finally show that $\Pr(b_i > a\widehat{b}_i) \leq 2^{-a}$, for $i = 1, \ldots, n$ and for all $a \geq 1$. Since this is obvious if $\mathcal{B}_i$ is resolved, let us assume that this is not the case. But then $\widehat{b}_i \geq 16\tau$, so we can assume without loss of generality that $b_i > 16\tau a$. Now by property (B) of Lemma 8.11, $\Pr(b_i > a\widehat{b}_i) \leq 2^{-b_i/(8\tau)} + 2^{-2a} \leq 2^{-2a} + 2^{-2a} \leq 2^{-a}$. ∎

9 Semisorting

We first define the semisorting problem precisely and then outline the rest of the section.

Informally, the m-color semisorting problem inputs n elements with values in the range $0..m$ (elements with a value of 0 being dummy elements) and places these in an output array of size $O(n)$ such that all elements with a given color (i.e., nonzero value) occur together, separated only by empty cells. As usual, we model the input as n integers $x_1, \ldots, x_n$ in the range $0..m$ and the output as n nonnegative integers $y_1, \ldots, y_n$, where y_j should be thought of as the position in the output array of the jth input element, for $j = 1, \ldots, n$. Condition (1) below means that distinct (real) elements are not placed in the same output cell, condition (2) says that no element of a different color intervenes between two elements of the same color in the output array, and condition (3) requires the output array to be of size $O(n)$.

Definition: For all $n, m \in \mathbb{N}$, the *m-color semisorting* problem of size n is the following: Given n integers $x_1, \ldots, x_n$ in the range $0..m$, compute n nonnegative integers $y_1, \ldots, y_n$ such that

(1) For $1 \leq i < j \leq n$, if $x_i \neq 0$, then $y_i \neq y_j$;

(2) For all $i, j, k \in \{1, \ldots, n\}$, if $y_i < y_j < y_k$ and $x_i = x_k$, then $x_j = x_i$;

(3) $\max\{y_j : 1 \leq j \leq n\} = O(n)$.

The *m-color strong semisorting* problem is identical, except for the additional requirement (4) below, where $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, m$.

(4) For all $j, k \in \{1, \ldots, n\}$, if $x_j = x_k = i \neq 0$, then $|y_j - y_k| = O(b_i)$.

It is perhaps instructive to compare the definition of semisorting with that of complete linear colored compaction given in Section 6. Informally, complete linear colored compaction requires upper bounds $d_1, \ldots, d_m$ on the sizes of the color classes to

be specified as part of the input, and the input elements are placed in arrays $A_1, \ldots, A_m$ of sizes $O(d_1), \ldots, O(d_m)$, respectively, each of which is indexed starting at 1. In order to use an algorithm for complete linear colored compaction to semisort, one could therefore first compute estimates $\widehat{b}_1, \ldots, \widehat{b}_m$ with $\widehat{b}_i \geq b_i$, for $i = 1, \ldots, m$, but $\sum_{i=1}^{m} \widehat{b}_i = O(n)$, then use the given compaction subroutine with limits $\widehat{b}_1, \ldots, \widehat{b}_m$, and finally place the arrays $A_1, \ldots, A_m$ together in a base array of size $O(n)$.

The present section culminates in a proof that n-color semisorting problems of size n can be solved in $O(\log^* n)$ time with optimal speedup (with high probability). As mentioned in the introduction, this leads to an algorithm with optimal speedup for computing n-color fine-profiles in $O(\log^* n)$ time (Corollary 10.5). On the other hand, our path to optimal semisorting takes us via no fewer than four auxiliary profilers (i.e., algorithms that compute profiles) of different types, all of which are finally subsumed by the n-color fine-profiler.

We now give the plan of this section in more detail. In this overview only, we will use the terms "Type–1"–"Type–4" to refer to the auxiliary profilers mentioned above. We characterize these profilers using our standard terminology: The input consists of n integers $x_1, \ldots, x_n$ in the range $0..n$, we take $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, n$, and the output of a profiler is a sequence $\widehat{b}_1, \ldots, \widehat{b}_n$ of n nonnegative integers. We always require the linear-sum condition $\sum_{i=1}^{n} \widehat{b}_i = O(n)$ to be satisfied, for which reason we do not repeat this condition below.

A Type–1 profile should ideally have the properties of a coarse-profile. In order to be able to compute a Type–1 profile in $O(\log^* n)$ time, however, we are forced to allow a certain limited dependence between $\widehat{b}_1, \ldots, \widehat{b}_n$. We begin by showing (Lemma 9.1) that the proof of Lemma 8.3 can be modified to yield an algorithm for computing Type–1 profiles in $O(\log^* n)$ time with n processors. We then describe an algorithm that, guided by a Type–1 profile, can semisort in $O(\log^* n)$ time (Lemma 9.7); this algorithm, however, uses n processors and hence is not optimal. As a corollary, we obtain a nonoptimal Type–2 profiler, whose output $\widehat{b}_1, \ldots, \widehat{b}_n$ overestimates all multiplicities, in the sense that $\widehat{b}_i \geq b_i$, for $i = 1, \ldots, n$ (Corollary 9.18). Scattering in time allows us to derive from this nonoptimal Type–2 profiler an optimal Type–3 profiler, whose output $\widehat{b}_1, \ldots, \widehat{b}_n$ is such that only $O(n/\log^* n)$ elements belong to (badly estimated) color classes $\mathcal{B}_i$ with $b_i > \sqrt{\log^* n} \cdot \widehat{b}_i$ (Lemma 9.19). By means of another scattering in time that uses both a nonoptimal Type–2 profiler and an optimal Type–3 profiler as subroutines, we are then able to derive an optimal Type–4 profiler, whose output $\widetilde{b}_1, \ldots, \widetilde{b}_n$ is such that only $O(n/\log^* n)$ elements belong to color classes $\mathcal{B}_i$ with $b_i > \widetilde{b}_i$ (Lemma 9.20). Scattering in time using this last result as well as the nonoptimal semisorting algorithm finally yields an optimal semisorting algorithm (Theorem 9.21).

As mentioned above, our first goal is to demonstrate that a profile similar to a

coarse-profile can be computed in $O(\log^* n)$ time. Recall that the algorithm of Lemma 8.3 works in $O(\log^* n)$ stages, each of which carries out a number of simultaneous graduated conditional scatterings after allocating the necessary resources. The bottleneck in the algorithm clearly is the resource allocation (line (5)): Whereas the remaining parts of a stage run in constant time, the application of the algorithms of Theorems 7.1 and 7.2 uses $\Theta(\log^* n)$ time in every stage, totalling $\Theta((\log^* n)^2)$ time overall.

The idea for reducing the running time to $O(\log^* n)$ is to replace the $\Theta(\log^* n)$-time complete allocation by the constant-time incomplete allocation of Lemma 7.3. As a result of this change, a color class $\mathcal{B}_i$ may, in a particular stage, lack the resources needed to carry out a GCS and possibly obtain an estimate $\widehat{b}_i$; in this case the elements of $\mathcal{B}_i$ do nothing, and $\mathcal{B}_i$ remains active, just as if it had filled all rows of its GCS. We must show that even with this additional source of uncertainty, the algorithm produces the desired output without exceeding its resource bounds; note that the resource requirements tend to go up, because a color class that remains active longer demands (significantly) more resources.

In stating condition (B) of the lemma below, which corresponds to condition (B) in the definition of a coarse-profile, we use the convenient shorthand $\sum_{b_i > a\widehat{b}_i}$ to denote summation over the set of indices of those color classes $\mathcal{B}_i$ with $b_i > a\widehat{b}_i$. Similar notation will be employed later without comment.

Lemma 9.1: There is a constant $\epsilon > 0$ such that for all given $n \in \mathbb{N}$, the following problem can be solved on a TOLERANT PRAM using $O(\log^* n)$ time, n processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo): Given n integers $x_1, \ldots, x_n$ in the range $0\,..\,n$, compute nonnegative integers $\widehat{b}_1, \ldots, \widehat{b}_n$ such that
(A) $\sum_{i=1}^{n} \widehat{b}_i = O(n)$;
(B) for all $a \geq 1$, $\sum_{b_i > a\widehat{b}_i} (b_i + \widehat{b}_i) = O(n/2^a + n^{3/4})$,
where $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, n$.

Proof: As in the proof of Lemma 8.3, we can assume that $b_i \leq n^{1/8}$, for $i = 1, \ldots, n$. We now use the algorithm of Lemma 8.3, except that we replace the complete allocation of processors and space in line (5) by $(5v_t)$-allocation according to Lemma 7.3. The resulting algorithm clearly works in $O(\log^* n)$ time. Borrowing the definition of a tardy color class from the proof of Lemma 8.3 and again taking $\widehat{b}_i = 0$ if $\widehat{b}_i$ does not become defined in any stage, for $i = 1, \ldots, n$, we establish Lemma 9.1 in Lemmas 9.2–9.6 below, which essentially parallel Lemmas 8.5–8.10.

Lemma 9.2: For $i = 1, \ldots, n$, with high probability $\widehat{b}_i$ becomes defined in some stage.

Proof: Fix $i \in \{1, \ldots, n\}$. If $\widehat{b}_i$ does not become defined in any stage, either $l_i = 5v_T$ in Stage T, or $\mathcal{B}_i$ is unlucky (i.e., it does receive any processors and space in the incomplete

allocation in line (5)) in Stage T. By the proof of Lemma 8.5, the former happens with negligible probability, while, by Lemma 7.3(a), the latter happens with probability at most 2^{-5v_T}, which is also negligible. ∎

Lemma 9.3: For $t = 1, \ldots, T$, with high probability the number of color classes that are tardy in Stage t is $O(n/v_t^5)$.

Proof: The claim is obvious for $t = 1$ since $v_1 = O(1)$, so fix $t \in \{2, \ldots, T\}$ and recall that for $i = 1, \ldots, n$, $\mathcal{B}_i$ is tardy in Stage t exactly if it is still active at the beginning of Stage t and $b_i < v_t^4$. This can happen only if either $l_i = 5v_{t-1}$ in Stage $t-1$, or $\mathcal{B}_i$ is unlucky in Stage $t-1$. By the proof of Lemma 8.6, the former happens with probability at most v_t^{-6}; a martingale argument hence shows that the number of color classes falling into this category is at most $O(n/v_t^6 + n^{3/4}) = O(n/v_t^5)$. On the other hand, Lemma 7.3(b) implies that at most $n/2^{5v_{t-1}} = O(n/v_t^5)$ color classes are unlucky in Stage $t-1$, except with probability $2e^{-\zeta}$, where $\zeta = (n/2^{5v_{t-1}})^2/(2^9 \cdot n \cdot (5v_{t-1})^3) = \Omega(n/v_t^{13}) = \Omega(n^{1/4})$. The claim follows. ∎

Lemma 9.4: With high probability, the algorithm uses $O(n)$ processors and $O(n)$ space.

Proof: Fix $t \in \{1, \ldots, T\}$. The following happens with high probability: By the previous lemma, the number of tardy color classes in Stage t is $O(n/v_t^5)$. On the other hand, a color class that is active but not tardy in Stage t by definition is of size at least v_t^4, so that there are at most n/v_t^4 such color classes. Consequently, the total number of color classes active at the beginning of Stage t is $O(n/v_t^4)$, which implies that the total number of processors and memory cells allocated in Stage t is $O(n/v_t^3)$. We may therefore conclude from Lemma 7.3 that each stage can be executed in constant time using $O(n)$ processors and $O(n)$ space. ∎

What remains is to show that the algorithm satisfies conditions (A) and (B) of Lemma 9.1. An attempt to carry out the proof in complete analogy with the proofs of Lemmas 8.8–8.10 is thwarted by the fact that the incomplete allocation of resources introduces complex dependencies between color classes (which, formerly, were independent), so that the simple-minded applications of martingale arguments in the earlier proofs cannot be justified in the present setting. In order to cope with this situation, we have to identify and analyze random experiments that *are* truly independent. We could choose to focus on the elementary random choices made by the processors, i.e., the random numbers generated, but the ("long-term") effect of a change in one of these quantities is complicated and difficult to handle. Fortunately, the graduated conditional scatterings carried out within a single stage are independent, provided that we fix all random choices made by the algorithm in earlier stages and in

the incomplete allocation in the stage under consideration. Put differently, we consider a random experiment comprising only the graduated conditional scatterings in Stage t, for some $t \in \{1, \ldots, T\}$. As noted previously, if we can show that a given property of this random experiment holds with probability at least p for any fixed setting of the earlier random choices, we can conclude that the property holds with probability at least p even in the actual experiment consisting of the whole execution of the algorithm, where earlier random choices in fact are not fixed. When using this principle, we will indicate the restricted probability space under consideration by stating that the analysis is carried out "for any fixed past of the GCS in Stage t". In the following two proofs, let $\delta_{i,t} = 1$ if $\widehat{b}_i$ becomes defined in Stage t and $\delta_{i,t} = 0$ otherwise, for $i = 1, \ldots, n$ and $t = 1, \ldots, T$.

Lemma 9.5: With high probability, $\sum_{i=1}^{n} \widehat{b}_i = O(n)$.

Proof: Let $i \in \{1, \ldots, n\}$ and $t \in \{1, \ldots, T\}$ and assume first that $b_i > 0$. Translated to the present setting, the proof of Lemma 8.8 showed that for any fixed past of the GCS in Stage t and for all $a \geq 1$, $\Pr(\widehat{b}_i > ab_i \mid \delta_{i,t} = 1) \leq (K/a)^6$. Since the proof of Lemma 8.9 depends only on the integer random variable $\widehat{b}_i$ satisfying the relation $\Pr(\widehat{b}_i > ab_i) \leq (K/a)^6$, for all $a \geq 1$, the same proof now yields that $E(\widehat{b}_i \mid \delta_{i,t} = 1) = O(b_i + 1)$, still for any fixed past of the GCS in Stage t. The same obviously holds if $b_i = 0$, since in that case $\widehat{b}_i = K$ (provided that $\widehat{b}_i$ becomes defined in some stage).

Considering again the entire execution as the random experiment, we will show that with high probability, $\sum_{i=1}^{n} \widehat{b}_i \cdot \delta_{i,t} = O(\sum_{i=1}^{n}(b_i + 1) \cdot \delta_{i,t} + n^{3/4})$, for $t = 1, \ldots, T$. Summing this over all values of t yields $\sum_{i=1}^{n} \widehat{b}_i = O(\sum_{i=1}^{n}(b_i + 1) + n^{3/4} \log^* n) = O(n)$, as desired.

We show the relation $\sum_{i=1}^{n} \widehat{b}_i \cdot \delta_{i,t} = O(\sum_{i=1}^{n}(b_i + 1) \cdot \delta_{i,t} + n^{3/4})$ by demonstrating that it holds with high probability for any fixed past of the GCS in Stage t. Given such a fixed past, n random variables $Z_1, \ldots, Z_n$ will be independent, provided that Z_i is defined exclusively in terms of the GCS of $\mathcal{B}_i$ in Stage t (or of the fact that $\mathcal{B}_i$ does not carry out any GCS in Stage t), for $i = 1, \ldots, n$. For a certain constant $C > 0$, we will use this with $Z_i = (C(b_i + 1) - \widehat{b}_i) \cdot \delta_{i,t}$, for $i = 1, \ldots, n$. Choose C sufficiently large to make $E(Z_i \mid \delta_{i,t} = 1) \geq 0$, for $i = 1, \ldots, n$, which is possible by what was shown above. Since clearly $E(Z_i \mid \delta_{i,t} = 0) = 0$, we have $E(Z_i) \geq 0$, for $i = 1, \ldots, n$. But then $E(S) \geq n^{3/4}$, where $S = \sum_{i=1}^{n} Z_i + n^{3/4}$, and S changes by $O(n^{1/8})$ in response to a change in a single Z_i. An application of Lemma 2.2(b) therefore shows that with high probability, $S \geq E(S)/2$. In particular, with high probability $S \geq 0$ or, what is the same, $\sum_{i=1}^{n} \widehat{b}_i \cdot \delta_{i,t} \leq C \sum_{i=1}^{n}(b_i + 1) \cdot \delta_{i,t} + n^{3/4}$. ∎

Lemma 9.6: For all $a \geq 1$, with high probability $\sum_{b_i > a\widehat{b}_i}(b_i + \widehat{b}_i) = O(n/2^a + n^{3/4})$.

Proof: Since $\widehat{b}_i < b_i$ whenever $b_i > a\widehat{b}_i$ and since $\widehat{b}_i > 0$ for all $i \in \{1, \ldots, n\}$ with high probability (Lemma 9.2), it suffices to show that with high probability, $\sum_{b_i > a\widehat{b}_i > 0} b_i = O(n/2^a + n^{3/4})$.

Given an assertion P, let $\nu(P) = 1$ if P is true, and $\nu(P) = 0$ if P is false. A martingale argument shows that for any $t \in \{1, \ldots, T\}$ and for any fixed past of the GCS in Stage t, with high probability

$$\sum_{i=1}^{n} b_i \cdot \nu(b_i > a\widehat{b}_i \text{ and } \delta_{i,t} = 1) = O\left(\sum_{i=1}^{n} b_i \cdot \Pr(b_i > a\widehat{b}_i \text{ and } \delta_{i,t} = 1) + n^{5/7}\right). \quad (5)$$

We are interested in the double sum obtained by summing (5) for $t = 1, \ldots, T$. At this point recall that the proof of Lemma 8.10 showed that $\sum_{t=1}^{T} \Pr(b_i > a\widehat{b}_i$ and $\delta_{i,t} = 1) \leq 2^{-a}$ for all $i \in \{1, \ldots, n\}$ by partitioning the index set $\{1, \ldots, T\}$ of the summation into three sets I_1, I_2 and I_3, corresponding to the three cases "b_i is too large", "b_i is of suitable size" and "b_i is too small", respectively (the partition depends on i). In the first two cases the earlier analysis of a stage can be seen to hold for any fixed past of the GCS in that stage — in fact, the analysis formally was carried out in the restricted probability space. Summing (5) for $t = 1, \ldots, T$, but including only the terms with $t \in I_1 \cup I_2$, therefore yields

$$\sum_{i=1}^{n} \sum_{t \in I_1 \cup I_2} b_i \cdot \nu(b_i > a\widehat{b}_i \text{ and } \delta_{i,t} = 1) = O\left(\sum_{i=1}^{n} b_i \cdot 2^{-a} + n^{5/7} \log^* n\right) = O(n/2^a + n^{3/4}).$$

Concerning the third case, "b_i is too small", recall that $b_i > a\widehat{b}_i$ and $\delta_{i,t} = 1$ for some $i \in \{1, \ldots, n\}$ and $t \in I_3$ only if $\mathcal{B}_i$ is tardy in Stage t and $v_t \geq 2^a$. Therefore $\sum_{i=1}^{n} \sum_{t \in I_3} b_i \cdot \nu(b_i > a\widehat{b}_i$ and $\delta_{i,t} = 1) \leq \sum_{v_t \geq 2^a} \sum_{i=1}^{n} b_i \cdot \nu(\mathcal{B}_i$ is tardy in Stage $t)$, where the outer summation is over those $t \in \{1, \ldots, T\}$ with $v_t \geq 2^a$. For $t = 1, \ldots, T$, a color class which is tardy in Stage t is of size at most v_t^4, so that Lemma 9.3 implies that with high probability the total size of the color classes that are tardy in Stage t is $O(n/v_t)$. But then the above double sum is $\sum_{v_t \geq 2^a} O(n/v_t) = O(n/2^a)$. Altogether, we have shown that with high probability, $\sum_{b_i > a\widehat{b}_i > 0} b_i = \sum_{i=1}^{n} \sum_{t=1}^{T} b_i \cdot \nu(b_i > a\widehat{b}_i$ and $\delta_{i,t} = 1) = O(n/2^a + n^{3/4})$. This ends the proofs of Lemmas 9.6 and 9.1. ∎

The next major goal is to devise an n-processor semisorting algorithm. The basic idea is simply to use techniques similar to those of Section 4 to compact the elements of each color class into an array of suitable size. However, this approach meets with major complications. Multiplicities must be estimated using the algorithm of Lemma 9.1, which means that the estimates obtained are not very reliable. Arrays must be allocated as described in Section 7, and colors cannot be handled independently, as far as the placement in arrays is concerned, since the failure probability for small color classes

cannot be ignored. Instead it is necessary to monitor the progress of the colors throughout the process, pushing more resources towards colors that are not keeping pace with the rest.

Lemma 9.7: There is a constant $\epsilon > 0$ such that for all given $n \in \mathbb{N}$, n-color semisorting problems of size n can be solved on a TOLERANT PRAM using $O(\log^* n)$ time, n processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: In order to let the basic idea stand out clearly, we first describe and analyze a simplified algorithm that ignores a number of complications, and afterwards motivate the various bells and whistles that have to be added to the algorithm to actually make it work.

In the following the word "element" will be used exclusively to denote elements of color classes. The simplified algorithm begins by computing a sequence $v_1, \ldots, v_T$ similar to the one used in the proof of Lemma 8.3 (the exact requirements will be specified below), and then executes the following:

```
Let all elements be active;
for t := 1 to T do
  for each color class B_i pardo
    begin
      Allocate an array A_{i,t} of size 6b_i/v_t^3 to B_i;
      v_t-compact the active elements in B_i to A_{i,t},
        deactivating every lucky element in B_i;
    end;
```

The algorithm hence consists of T stages, each of which attempts to place the elements of each color class in an array of suitable size. The elements that are successfully placed in the array become inactive and do not participate in subsequent stages. Conceptually, if the algorithm succeeds in deactivating all elements, the elements of each color class $\mathcal{B}_i$ afterwards are stored in T arrays $A_{i,1}, \ldots, A_{i,T}$. In actual fact, there is no need to preserve $A_{i,t}$ beyond the end of Stage t, for $i = 1, \ldots, n$ and $t = 1, \ldots, T$, and the algorithm can reclaim the space allocated in each stage for reuse in the following stage. Instead, in Stage t, the algorithm remembers the total size $\sum_{l=1}^{t-1} |A_{i,l}|$ of the arrays allocated to $\mathcal{B}_i$ in earlier stages, for $i = 1, \ldots, n$ and $t = 1, \ldots, T$, and each element of $\mathcal{B}_i$ deactivated in Stage t adds this offset to its position in $A_{i,t}$ and stores the resulting absolute position. Informally, this has the effect of gluing the array $A_{i,t}$ onto the right end of an array already containing $A_{i,1}, \ldots, A_{i,t-1}$. After the last stage, it is therefore an easy matter to use the algorithm of Theorem 7.1 to allocate a single array A_i of size $|A_{i,1}| + \cdots + |A_{i,T}|$ to $\mathcal{B}_i$ and to place the elements of $\mathcal{B}_i$ in A_i (informally, to move

the elements of $\mathcal{B}_i$ from $A_{i,1}, \ldots, A_{i,T}$ to A_i). Provided that $\sum_{i=1}^{n} |A_i| = O(n)$, this produces a solution to the semisorting problem. In other words, if all elements are deactivated, the correctness of the algorithm will be guaranteed if we can show that $\sum_{i=1}^{n} \sum_{t=1}^{T} |A_{i,t}| = O(n)$. In the idealized algorithm above, this condition is satisfied, since $\sum_{i=1}^{n} \sum_{t=1}^{T} |A_{i,t}| = \sum_{t=1}^{T} \sum_{i=1}^{n} (6b_i/v_t^3) \leq \sum_{t=1}^{T} (6n/v_t^3) = O(n)$.

In order to analyze the rate with which elements are deactivated, fix $i \in \{1, \ldots, n\}$, let $t \in \{1, \ldots, T\}$ and assume by way of induction that the number of active elements in $\mathcal{B}_i$ has decreased to at most b_i/v_t^4 before Stage t. Then $\mathcal{B}_i$ is well-supplied in the v_t-compaction in Stage t, i.e. (Lemma 6.1(a)), the probability that a fixed element in $\mathcal{B}_i$ remains active is at most 2^{-v_t}. Therefore the expected number of elements in $\mathcal{B}_i$ that remain active is at most $b_i \cdot 2^{-v_t}$, which, for a suitable choice of the sequence $v_1, \ldots, v_T$, is significantly smaller than the b_i/v_{t+1}^4 required for the induction.

We now proceed to discuss the problems with the above algorithm. One such problem is that color classes may be too small to exhibit a "reliable" behavior, in a statistical sense. E.g., in the analysis in the preceding paragraph, even though the expected number of elements in $\mathcal{B}_i$ that remain active is significantly below b_i/v_{t+1}^4, the probability that their actual number exceeds b_i/v_{t+1}^4 may not be negligible (cf. Lemma 6.1(b), which yields little unless z is much larger than v). We counter this problem by treating small color classes specially; in particular, the space allocated to small color classes is larger, relative to their sizes, than for other color classes.

More significantly, we do not know the multiplicities $b_1, \ldots, b_n$, so we have to resort to estimates $\widehat{b}_1, \ldots, \widehat{b}_n$. One consequence of this is that we do not really know whether a color class is small; the color classes that are treated specially, as mentioned in the previous paragraph, are hence those whose estimates let them "appear" small.

Another difficulty is posed by the allocation of space to color classes. Similarly as in the algorithm of Lemma 9.1, we have to resort to the incomplete allocation of Lemma 7.3 (the color classes themselves will be requesting elements in the sense of Lemma 7.3, but our convention in this section is to use the word "element" only for elements of color classes). As a consequence of the incomplete allocation, in each stage certain color classes will be unlucky (recall that this means that they do not receive the resources that they requested), so that they cannot participate in the v_t-compaction; this adds another complication to the analysis. Furthermore, in order for the resource requirements of the incomplete allocation to remain $O(n)$ in spite of the increase in v_t over the stages, it is necessary to ensure that the number of requests per stage decreases over the course of the execution. We therefore initially let all nonempty color classes be active, declare a color class to become inactive when it loses its last active element, and allocate space only to active color classes. Since the number of active color classes cannot exceed the number of active elements, the number of active color classes will decrease as required, provided

that the number of active elements does so.

Finally, although for most color classes the incomplete compaction in a particular stage will succeed in deactivating most elements in the color class, for some color classes almost all elements may remain active; in particular, this surely happens for color classes that are unlucky in the stage under consideration, and it is likely to happen for those whose sizes were heavily underestimated. Given the algorithm as described so far, the problem will be aggravated over successive stages, since the available space decreases. In such cases we need to resort to an "emergency escape", which will be to compact into an array of size $6v_t^2\widehat{b}_i$, rather than $6\widehat{b}_i/v_t^3$. Note that we certainly cannot use $6v_t^2\widehat{b}_i$ space always, since the total space requirements would be superlinear. In order to decide when to apply the emergency escape, we begin each stage with a "test scattering" for each active color class, except those that appear small. The test scattering is simply a conditional scattering by the active elements in the color class, with parameters chosen to allow color classes that are in need of the emergency escape to be roughly distinguished from those that are not; the actual bit observed is whether or not the fullness of the scattering equals 1.

Before describing the algorithm proper, we define the sequence $v_1, \ldots, v_T$ and establish a lemma that parallels Lemma 8.4. Similarly as in Section 8, let $f(z) = z - 44\log z$, for $z > 0$. Then $f(2) < 0$, $f(2^9) > 0$, $f(z) \to \infty$ for $z \to \infty$, and f' has only one zero, so that $z > 44\log z \geq 22\lceil \log z \rceil$ for all $z \geq 2^9$. The algorithm below hence outputs a finite sequence, which we take to be $v_T, \ldots, v_1$.

$z := \lceil n^{1/88} \rceil$;
repeat
  write(z);
  $z := 22\lceil \log z \rceil$;
until $z < 2^{15}$;

Clearly $v_1 < 2^{2^{15}} = O(1)$. We will assume without loss of generality that $T \geq 2$ (otherwise n is bounded by a constant). Then $2^{15} \leq v_1 < v_2 < \cdots < v_T = \lceil n^{1/88} \rceil$. For a sufficiently large constant z_0, clearly $22\lceil \log(22\lceil \log z \rceil) \rceil \leq 44\log(44\log z) = 44(\log(44) + \log\log z) \leq \log z$ for all $z \geq z_0$. Hence $T \leq 2\log^* n + O(1)$.

Lemma 9.8:

(a) For $t = 1, \ldots, T-1$, $v_{t+1}^{22} \leq 2^{v_t} \leq v_{t+1}^{44}$;

(b) For $t = 1, \ldots, T-1$, $v_t^{22} \leq v_{t+1}$;

(c) For $t = 1, \ldots, T$, $v_t \geq 2^t$.

Proof:

(a) $\log(v_{t+1}^{22}) = 22 \log v_{t+1} \leq 22\lceil \log v_{t+1} \rceil = v_t \leq 44 \log v_{t+1} = \log(v_{t+1}^{44})$, from which the relation follows by exponentiation.

(b) We noted above that $2^z \geq z^{44}$ for $z \geq 2^9$. Since $v_t \geq 44 \cdot 2^9$, part (a) therefore implies that $v_{t+1} \geq 2^{v_t/44} \geq (v_t/44)^{44} \geq v_t^{22}$.

(c) The claim follows trivially from part (b) by induction. ∎

The algorithm begins by computing a profile $\widetilde{b}_1, \ldots, \widetilde{b}_n$ for $x_1, \ldots, x_n$ with the properties described in Lemma 9.1. By Theorem 8.2, we can assume that $\widetilde{b}_i$ is a linear overestimate for b_i for all $i \in I$, where I is a known subset of $\{1, \ldots, n\}$ with the property that $i \in I$ for all colors i with $b_i \geq n^{1/88}$. Since we can clearly remove all colors i with $\widetilde{b}_i < n^{1/88}$ from I without affecting the property of I just mentioned, we can also assume that $|I| = O(n^{87/88})$. We want to apply (the nonoptimal part of) the algorithm of Theorem 6.2 to place the elements of $\mathcal{B}_i$ in an array of size $O(b_i)$ (with $\widetilde{b}_i$ serving as the limit for $\mathcal{B}_i$), for all $i \in I$. Recall that the algorithm of Theorem 6.2 can cope with up to $\Theta(n^{1-\delta})$ colors, for arbitrary fixed $\delta > 0$, as long as we provide a means of performing the necessary space allocation. Since Theorem 7.1 is now available, the latter condition no longer is a problem. Using the algorithms of Theorem 4.6 and Lemma 8.1 to replace the colors in I by colors in a range of size $O(n^{87/88})$, we can therefore apply the algorithm of Theorem 6.2 as stated. This preprocessing serves to let us assume without loss of generality that $b_i \leq n^{1/88}$, for $i = 1, \ldots, n$. We want to work with estimates that are at least 1 and at most $n^{1/88}$ and therefore take $\widehat{b}_i = \min\{\max\{\widetilde{b}_i, 1\}, \lfloor n^{1/88} \rfloor\}$, for $i = 1, \ldots, n$. It is easy to see that properties (A) and (B) of Lemma 9.1 continue to hold for the modified profile $\widehat{b}_1, \ldots, \widehat{b}_n$. Additionally, $b_i \leq v_T \widehat{b}_i$ and $\widehat{b}_i \leq n^{1/88}$, for $i = 1, \ldots, n$. The algorithm now proceeds as follows.

```
(1)  Let all elements and all nonempty color classes be active;
(2)  for t := 1 to T do (* Stage t *)
(3)    for each active color class B_i pardo
(4)      begin
(5)        if b̂_i > v_t^17 (* not apparently-small *) then
(6)          begin (* test scattering *)
(7)            v_t-allocate v_t memory cells and v_t processors to B_i;
(8)            if B_i was unlucky in the allocation in line (7)
(9)            then goto line (21);
(10)           Let the elements in B_i carry out a conditional
                 scattering S_i with probability v_t^7/b̂_i and of width v_t;
(11)           Huge_{i,t} := (S_i has fullness 1);
(12)         end;
```

(13) **if** $\widehat{b}_i \leq v_t^{17}$ (* apparently-small *) **then** $Size_{i,t} := 6v_t^{19}$

(14) **else if** $Huge_{i,t}$ **then** (* emergency escape *) $Size_{i,t} := 6v_t^2\widehat{b}_i$

(15) **else** (* normal case *) $Size_{i,t} := \lceil 6\widehat{b}_i/v_t^3 \rceil$;

(16) v_t-allocate an array $A_{i,t}$ of size $Size_{i,t}$ to $\mathcal{B}_i$;

(17) **if** $\mathcal{B}_i$ was lucky in the allocation in line (16) **then**

(18) v_t-compact the active elements in $\mathcal{B}_i$ to $A_{i,t}$,
deactivating every lucky element in $\mathcal{B}_i$;

(19) **if** no element in $\mathcal{B}_i$ remains active

(20) **then** make $\mathcal{B}_i$ inactive;

(21) **end**;

For $t = 1, \ldots, T$, let *Stage* t be the tth execution of lines (3)–(21) and say that a color class is active in Stage t if it is active at the beginning of Stage t and that it is unlucky in Stage t if it is unlucky in the incomplete allocation in either line (7) or line (16) in Stage t. A color class that is unlucky in some stage drops out of that stage and rejoins the computation in the beginning of the next stage, if any; this is realized via a goto instruction in line (9) and a conditional instruction in line (17).

The goal of the analysis is to show that with high probability, the algorithm deactivates all elements using $O(n)$ processors and $O(n)$ space and allocating arrays $A_{i,t}$ of total size $O(n)$. Since a stage can be executed in constant time, the algorithm is then correct and its resource requirements are as claimed in Lemma 9.7. A key property established below is that the number of active elements (and hence of active color classes) decreases rapidly over the execution of the algorithm. More precisely, we will show (Lemma 9.14) that with high probability the number of elements active at the beginning of Stage t, for $t = 1, \ldots, T$, is $O(n/v_t^{22})$. The proof of this key property consists of two main parts. We first identify certain favorable conditions that may apply to a color class in a stage, show that these conditions together imply that the color class is well-supplied, in the sense of Lemma 6.1, in the v_t-compaction in the given stage (Lemma 9.9), and note that for well-supplied color classes the rate of deactivation is essentially as in the idealized analysis of the simplified algorithm earlier in this section (Lemma 9.12). We then show that only very few color classes lack the favorable conditions (Lemma 9.13).

For $i = 1, \ldots, n$ and $t = 1, \ldots, T$, denote by $N_{i,t}$ the number of active elements in $\mathcal{B}_i$ at the start of Stage t. $\mathcal{B}_i$ is said to be *apparently-small* in Stage t if $\widehat{b}_i \leq v_t^{17}$, and to be *apparently-huge* in Stage t if a test scattering for $\mathcal{B}_i$ is carried out in Stage t and achieves a fullness of 1 (i.e., if the value *true* is assigned to $Huge_{i,t}$). $\mathcal{B}_i$ is *well-estimated* in Stage t if $b_i \leq v_t\widehat{b}_i$ (i.e., the size of $\mathcal{B}_i$ may have been underestimated, but at most by a factor of v_t), and $\mathcal{B}_i$ is *well-supplied* in Stage t if it is active in Stage t, lucky in the incomplete allocations in lines (7) and (16) in Stage t, and well-supplied in the incomplete compaction in line (18) in Stage t. Recalling the definition of "well-supplied"

in Section 6, we observe that $\mathcal{B}_i$ is well-supplied in Stage t if and only if an array $A_{i,t}$ of size at least $6v_tN_{i,t}$ is allocated to $\mathcal{B}_i$ in Stage t. A sufficient set of conditions for this to happen is formulated in the following lemma.

Lemma 9.9: For $i = 1, \ldots, n$ and $t = 1, \ldots, T$, $\mathcal{B}_i$ is well-supplied in Stage t if it is active in Stage t and each of the following conditions holds:

(1) $\mathcal{B}_i$ is well-estimated in Stage t;

(2) $\mathcal{B}_i$ is lucky (in both v_t-allocation steps) in Stage t;

(3) $\mathcal{B}_i$ is apparently-small or apparently-huge in Stage t, or $N_{i,t} \leq \widehat{b}_i/v_t^4$.

Proof: Let $i \in \{1, \ldots, n\}$ and $t \in \{1, \ldots, T\}$ and assume that $\mathcal{B}_i$ is active in Stage t and that conditions (1)–(3) hold. In particular, an array $A_{i,t}$ of size $Size_{i,t}$ is allocated to $\mathcal{B}_i$ (condition (2)). If $\mathcal{B}_i$ is apparently-small in Stage t, then $b_i \leq v_t\widehat{b}_i \leq v_t^{18}$ (condition (1)) and $Size_{i,t} = 6v_t^{19}$. Otherwise $Size_{i,t} \geq 6\widehat{b}_i/v_t^3$, and if $N_{i,t} > \widehat{b}_i/v_t^4$, then $Size_{i,t} = 6v_t^2\widehat{b}_i \geq 6v_tb_i$ (conditions (3) and (1)). In all cases $Size_{i,t} \geq 6v_tN_{i,t}$, i.e., $\mathcal{B}_i$ is well-supplied in Stage t. ∎

Define the *density* of a conditional scattering with probability p and of width s carried out by a set of m elements as mp/s.

Lemma 9.10: For $t = 1, \ldots, T$, if a test scattering is executed in Stage t by a color class $\mathcal{B}_i$ with $N_{i,t} > \widehat{b}_i/v_t^4$, then the probability that $\mathcal{B}_i$ does not become apparently-huge in Stage t is at most 2^{-v_t}.

Proof: The density of the test scattering is at least $(\widehat{b}_i/v_t^4) \cdot (v_t^7/\widehat{b}_i) \cdot (1/v_t) = v_t^2$. Hence by Lemma 3.1(c), the probability in question is at most $v_t \cdot 2^{-v_t^2} \leq 2^{-v_t}$. ∎

Lemma 9.11: With high probability, the algorithm deactivates all elements.

Proof: We first show that with high probability, conditions (1)–(3) of Lemma 9.9 are satisfied in Stage T for all active color classes. We already noted that $b_i \leq v_T\widehat{b}_i$, for $i = 1, \ldots, n$, so that every color class is well-estimated in Stage T, i.e., condition (1) is satisfied. By Lemma 7.3(a), the probability that some active color class is unlucky in Stage T is at most $2n \cdot 2^{-v_T}$, i.e., condition (2) is also satisfied with high probability. Condition (3), finally, follows directly from Lemma 9.10.

By what was shown above and Lemma 9.9, with high probability every active color class is well-supplied in Stage T. Lemma 6.1(a) implies that the probability that a fixed active element in an active and well-supplied color class is unlucky in the v_T-compaction in Stage T is at most 2^{-v_T}. Hence with high probability, no element remains active at the end of Stage T. ∎

Lemma 9.12: For $t = 1, \ldots, T-1$, with high probability the number of active elements in well-supplied color classes at the end of Stage t is $O(n/v_{t+1}^{22})$.

Proof: An element whose color class is well-supplied in Stage t remains active at the end of Stage t only if it is unlucky in the v_t-compaction in Stage t. But Lemmas 9.8(a) and 6.1(b) show the number of such elements to be no larger than n/v_{t+1}^{22}, except with probability $2e^{-\zeta}$, where $\zeta \geq (n/v_{t+1}^{22})^2/(2^9 n v_t^3) \geq n/(2^9 v_{t+1}^{45}) = \Omega(n^{1/4})$, i.e., except with negligible probability. ∎

Lemma 9.12 shows that the number of elements in well-supplied color classes decreases as required. In Lemma 9.13 we prove that the elements in color classes that are active but not well-supplied are so few that they can be ignored in this context. Informally, the reason for this is that if an active color class is not well-supplied in Stage t, then either it is unlucky, or its estimate is off by a factor of more than v_t, or the test scattering for the color class does not achieve fullness 1 although its density is at least v_t^2, all of which are unlikely.

Lemma 9.13: Let $t \in \{1, \ldots, T-1\}$ and take $I = \{i : 1 \leq i \leq n$ and $\mathcal{B}_i$ is active but not well-supplied in Stage $t\}$. Then, with high probability, $\sum_{i \in I}(b_i + \widehat{b}_i) = O(n/v_{t+1}^{22})$.

Proof: If a color class $\mathcal{B}_i$ is active but not well-supplied in Stage t, then one of Conditions (1)–(3) of Lemma 9.9 must be violated. Therefore the index sets I', I'' and I''' defined below cover all of I, i.e., $I' \cup I'' \cup I''' = I$; we will show that with high probability the sum $\sum_i (b_i + \widehat{b}_i)$ over each of these index sets is $O(n/v_{t+1}^{22})$.

$$
\begin{aligned}
I' &= \{i \in I : \mathcal{B}_i \text{ is not well-estimated in Stage } t\}, \\
I'' &= \{i \in I : \mathcal{B}_i \text{ is unlucky in Stage } t\}, \quad \text{and} \\
I''' &= \{i \in I \setminus (I' \cup I'') : \mathcal{B}_i \text{ is neither apparently-small} \\
&\qquad \text{nor apparently-huge in Stage } t \text{ and } N_{i,t} > \widehat{b}_i / v_t^4\}.
\end{aligned}
$$

It follows directly from Lemmas 9.1 and 9.8(a) that with high probability $\sum_{i \in I'}(b_i + \widehat{b}_i) = O(n/2^{v_t} + n^{3/4}) = O(n/v_{t+1}^{22})$. In the rest of the proof we consider all random choices made by the algorithm in Stages $1, \ldots, t-1$ to be fixed in an arbitrary manner. Write $I'' = I_1'' \cup I_2''$, where I_1'' and I_2'' are the residue sets of the v_t-allocations in lines (7) and (16), respectively. If we further define $S_1 = \sum_{i \in I_1''}(b_i + \widehat{b}_i)$ and $S_2 = \sum_{i \in I_2''}(b_i + \widehat{b}_i)$ as functions of these residue sets, it is easy to see that both S_1 and S_2 satisfy a Lipschitz condition with constant $O(n^{1/88})$ (recall that $b_i + \widehat{b}_i = O(n^{1/88})$, for $i = 1, \ldots, n$). We also know for each v_t-allocation that a fixed color class is unlucky with probability at most 2^{-v_t} (Lemma 7.3(a)), so that $E(S_1 + S_2) = O(n/2^{v_t})$. By two straightforward applications of Lemma 7.3(c), we thus obtain that with high probability, $\sum_{i \in I''}(b_i + \widehat{b}_i) = S_1 + S_2 = O(n/2^{v_t} + n^{1/88} v_t^3 n^{5/8}) = O(n/v_{t+1}^{22} + n^{3/4}) = O(n/v_{t+1}^{22})$. Finally, if a color class $\mathcal{B}_i$ is

neither apparently-small nor unlucky in Stage t, a test scattering is carried out for $\mathcal{B}_i$ in Stage t. Hence, by Lemma 9.10, $\Pr(i \in I''') \leq 2^{-v_t}$, for $i = 1, \ldots, n$, and therefore $E(\sum_{i \in I'''}(b_i + \widehat{b}_i)) = O(\sum_{i=1}^{n}(b_i + \widehat{b}_i)/2^{v_t}) = O(n/2^{v_t})$. A simple martingale argument now ensures that, with high probability, $\sum_{i \in I'''}(b_i + \widehat{b}_i) = O(n/2^{v_t} + n^{3/4}) = O(n/v_{t+1}^{22})$. ∎

Lemma 9.14: For $t = 1, \ldots, T-1$, with high probability the number of elements (and hence color classes) active at the end of Stage t is $O(n/v_{t+1}^{22})$.

Proof: Immediate from Lemmas 9.12 and 9.13. ∎

We finally show that the total size of the arrays $A_{i,t}$ allocated in Stage t is $O(n/v_t^3)$, for $t = 1, \ldots, T$ (Lemma 9.16), from which will follow not only that the algorithm is correct, but also that it uses $O(n)$ processors and $O(n)$ space. Disregarding the arrays allocated to apparently-huge color classes, this can easily be done using Lemma 9.14. In order to handle the apparently-huge color classes, however, we first have to show the following technical lemma, which says that if a color class $\mathcal{B}_i$ is well-supplied in Stage t, then it is unlikely to contain more than $\max\{v_{t+1}\sqrt{b_i}, b_i/v_{t+1}^8\}$ active elements at the beginning of Stage $t+1$.

Lemma 9.15: Let $t \in \{1, \ldots, T-1\}$ and take $I = \{i : 1 \leq i \leq n,\ \mathcal{B}_i$ is active and well-supplied in Stage t and $N_{i,t+1} > \max\{v_{t+1}\sqrt{b_i}, b_i/v_{t+1}^8\}\}$. Then, with high probability, $\sum_{i \in I}(b_i + \widehat{b}_i) = O(n/v_{t+1}^{22})$.

Proof: Consider all random choices made in the algorithm before the v_t-compaction in line (18) in Stage t to be fixed in an arbitrary way and let $i \in \{1, \ldots, n\}$. Since $b_i/v_{t+1}^8 \geq b_i/2^{v_t}$, Lemmas 6.1(b) and 9.8(b) imply that if $\mathcal{B}_i$ is active and well-supplied in Stage t, then $N_{i,t+1} > \max\{v_{t+1}\sqrt{b_i}, b_i/v_{t+1}^8\}$ with probability at most $2e^{-\zeta}$, where $\zeta = v_{t+1}^2 b_i/(2^9 b_i v_t^3) \geq v_t^{44}/(2^9 v_t^3) \geq 2v_t$. We have thus shown that $\Pr(i \in I) \leq 2 \cdot 2^{-2v_t} \leq 2^{-v_t}$. Similarly as in the proof of Lemma 9.13, let $S = \sum_{i \in I}(b_i + \widehat{b}_i)$ and note that S satisfies a Lipschitz condition with constant $O(n^{1/88})$. Now $E(S) \leq \sum_{i=1}^{n}(b_i + \widehat{b}_i) \cdot 2^{-v_t} = O(n \cdot 2^{-v_t}) = O(n/v_{t+1}^{22})$, and by Lemma 6.1(c), with high probability $S = O(E(S) + n^{1/88} v_t^3 n^{5/8}) = O(n/v_{t+1}^{22} + n^{3/4}) = O(n/v_{t+1}^{22})$. ∎

Lemma 9.16: For $t = 1, \ldots, T$, the total size of the arrays $A_{i,t}$ allocated in Stage t is $O(n/v_t^3)$.

Proof: The claim is obvious for $t = 1$ since $v_1 = O(1)$, so fix $t \in \{2, \ldots, T\}$. By Lemma 9.14, with high probability the total number of arrays allocated in Stage t is $O(n/v_t^{22})$. Hence the total size of the arrays allocated in Stage t to color classes that are

not apparently-huge is

$$O\Big((n/v_t^{22})\cdot(v_t^{19}+1)+\sum_{i=1}^{n}(\widehat{b}_i/v_t^3)\Big)=O(n/v_t^3),$$

as desired.

What remains is to bound the total size of the arrays allocated to apparently-huge color classes. Let $I=\{i:1\le i\le n$ and $\mathcal{B}_i$ is apparently-huge in Stage $t\}$. It suffices to show that $\sum_{i\in I}\widehat{b}_i=O(n/v_t^5)$ with high probability, since then the total size of the arrays allocated to apparently-huge color classes in Stage t is $O(\sum_{i\in I}v_t^2\widehat{b}_i)=O(n/v_t^3)$. To this end we partition I into three subsets:

$$\begin{aligned}
I' &= \{i\in I:\mathcal{B}_i \text{ is not well-supplied in Stage } t-1 \text{ or}\\
&\qquad \mathcal{B}_i \text{ is not well-estimated in Stage } t\},\\
I'' &= \{i\in I\backslash I': N_{i,t}>\widehat{b}_i/v_t^7\}, \quad \text{and}\\
I''' &= \{i\in I\backslash I': N_{i,t}\le\widehat{b}_i/v_t^7\}.
\end{aligned}$$

By Lemmas 9.13 and 9.1, $\sum_{i\in I'}\widehat{b}_i=O(n/v_t^{22}+n/2^{v_t}+n^{3/4})=O(n/v_t^5)$ with high probability. Suppose next that $i\in I''$. Then $N_{i,t}>\widehat{b}_i/v_t^7$ and $\mathcal{B}_i$ is well-estimated in Stage t, i.e., $\widehat{b}_i\ge b_i/v_t$. Also, since $\mathcal{B}_i$ is apparently-huge in Stage t, it cannot be apparently-small in Stage t, so $\widehat{b}_i>v_t^{17}$. It follows that $N_{i,t}>b_i/v_t^8$ and also that $N_{i,t}>\widehat{b}_i^{1/2}v_t^{17/2}/v_t^7\ge\sqrt{b_i/v_t}\cdot v_t^{3/2}=v_t\sqrt{b_i}$. But then, by Lemma 9.15, $\sum_{i\in I''}\widehat{b}_i=O(n/v_t^5)$ with high probability.

As concerns I''', finally, we use the fact that a color class $\mathcal{B}_i$ with $N_{i,t}\le\widehat{b}_i/v_t^7$ is very unlikely to become apparently-huge in Stage t. Specifically, according to Lemma 3.1(d), the probability of this event is at most

$$\left(\frac{\widehat{b}_i}{v_t^7}\cdot\frac{v_t^7}{\widehat{b}_i}\cdot e\cdot\frac{1}{v_t}\right)^{v_t}\le 2^{-v_t},$$

and a simple martingale argument shows that $\sum_{i\in I'''}\widehat{b}_i=O(n/v_t^5+n^{1/88}n^{3/4})=O(n/v_t^5)$ with high probability. ∎

Lemma 9.17: With high probability, the algorithm is correct and uses $O(n)$ processors and $O(n)$ space.

Proof: We have already argued that the correctness of the algorithm follows from Lemmas 9.11, 9.16 and 9.8(c). It uses $O(n)$ processors and $O(n)$ space, plus the resources needed for the incomplete allocations in lines (7) and (16), which are $O(n)$ by Lemmas 9.14, 9.16 and 7.3, and the resources needed for the incomplete compaction in line (18), which are $O(n)$ by Lemmas 9.14 and 6.1. This ends the proofs of Lemmas 9.17 and 9.7. ∎

Corollary 9.18: There is a constant $\epsilon > 0$ such that for all given $n \in \mathbb{N}$, the following problem can be solved on a TOLERANT PRAM using $O(\log^* n)$ time, n processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas): Given n integers $x_1, \ldots, x_n$ in the range $0..n$, compute n nonnegative integers $\widehat{b}_1, \ldots, \widehat{b}_n$ such that
(A) $\sum_{i=1}^{n} \widehat{b}_i = O(n)$;
(B) For $i = 1, \ldots, n$, $\widehat{b}_i \geq b_i$,
where $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, n$.

Proof: After semisorting the input elements into an array A of size $O(n)$ using the algorithm of Lemma 9.7, we can use the algorithm of Lemma 2.8(a) to store them in a linked list in the order in which they occur in A. This makes it easy to compute the first and the last element in A of each nonempty color class, which identifies nonoverlapping subarrays $A_1, \ldots, A_n$ of A such that A_i contains all elements of $\mathcal{B}_i$, for $i = 1, \ldots, n$. All that remains is to take $\widehat{b}_i = |A_i|$, for $i = 1, \ldots, n$. ∎

The final goal in this section is to take the step from the nonoptimal algorithm of Lemma 9.7 to an optimal semisorting algorithm. We first have to show two lemmas that gradually increase our ability to compute good profiles with optimal speedup.

Lemma 9.19: There is a constant $\epsilon > 0$ such that for all given $n \in \mathbb{N}$ with $\tau \geq \log^* n$, the following problem can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo): Given n integers $x_1, \ldots, x_n$ in the range $0..n$, compute n nonnegative integers $\widehat{b}_1, \ldots, \widehat{b}_n$ such that
(A) $\sum_{i=1}^{n} \widehat{b}_i = O(n)$;
(B) $\sum_{b_i > \sqrt{\tau \widehat{b}_i}} b_i = O(n/\tau)$,
where $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, n$.

Proof: Assume that $\tau \leq n^{1/4}$ and that $b_i \leq n^{1/8}$, for $i = 1, \ldots, n$. Obtain $\widehat{b}_1, \ldots, \widehat{b}_n$ exactly as in the proof of Lemma 8.11, except that in Step 3 Corollary 9.18 is used instead of Lemma 8.3. Since the use of Lemma 8.3 was the only reason for the lower bound of $(\log^* n)^2$ on τ, this gives an algorithm with the stated resource requirements. We now demonstrate the correctness of the algorithm. Observe first that since the profile of Corollary 9.18 satisfies conditions (A) and (B) in the definition of a coarse-profile, Lemmas 8.12–8.14 continue to hold for the modified algorithm. Lemma 8.12 directly implies condition (A) of Lemma 9.19, and Lemma 8.14 states that $\Pr(b_i > \sqrt{\tau \widehat{b}_i}) \leq 2^{-\sqrt{\tau}}$, for $i = 1, \ldots, n$. We would like to conclude from this that $\sum_{b_i > \sqrt{\tau \widehat{b}_i}} b_i = O(n/\tau)$ with high probability, i.e., that condition (B) of Lemma 9.19 is satisfied as well. This would follow easily from a martingale argument if $\widehat{b}_1, \ldots, \widehat{b}_n$ were independent, but this is not necessarily the case. At this point, however, we can use the fact that the algorithm

of Corollary 9.18 overestimates all multiplicities. Assume that we were to execute the algorithm of Lemma 8.11 as above, but with $\widehat{b}_i^{\mathcal{Z}} = b_i^{\mathcal{Z}}$, for $i = 1, \ldots, n$, and let $\overline{b}_1, \ldots, \overline{b}_n$ be the resulting profile (corresponding to $\widehat{b}_1, \ldots, \widehat{b}_n$). In this "ideal" setting $\widehat{b}_1^{\mathcal{Z}}, \ldots, \widehat{b}_n^{\mathcal{Z}}$ are independent, since they are fixed, which makes the above martingale argument valid. Hence $\sum_{b_i > \sqrt{\tau}\overline{b}_i} b_i = O(n/\tau)$. The profile $\widehat{b}_1^{\mathcal{Z}}, \ldots, \widehat{b}_n^{\mathcal{Z}}$ actually used by our algorithm satisfies $\widehat{b}_i^{\mathcal{Z}} \geq b_i^{\mathcal{Z}}$, for $i = 1, \ldots, n$. Given the workings of the algorithm of Lemma 8.11, this can easily be seen to imply that $\widehat{b}_i \geq \overline{b}_i$, for $i = 1, \ldots, n$ (if a subroutine returns larger estimates, the final estimates become no smaller). But then

$$\sum_{b_i > \sqrt{\tau}\widehat{b}_i} b_i \leq \sum_{b_i > \sqrt{\tau}\overline{b}_i} b_i = O(n/\tau). \quad \blacksquare$$

Lemma 9.20: There is a constant $\epsilon > 0$ such that for all given $n \in \mathbb{N}$ with $\tau \geq \log^* n$, the following problem can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo): Given n integers $x_1, \ldots, x_n$ in the range $0..n$, compute n nonnegative integers $\widetilde{b}_1, \ldots, \widetilde{b}_n$ such that

(A) $\sum_{i=1}^n \widetilde{b}_i = O(n)$;

(B) $\sum_{b_i > \widetilde{b}_i} b_i = O(n/\tau)$,

where $b_i = |\{j : 1 \leq j \leq n \text{ and } x_j = i\}|$, for $i = 1, \ldots, n$.

Proof: Assume that $\tau \leq n^{1/4}$ and that $b_i \leq n^{1/8}$, for $i = 1, \ldots, n$. First use the algorithm of Lemma 9.19 to compute a profile $\widehat{b}_1, \ldots, \widehat{b}_n$ for $x_1, \ldots, x_n$ with $\sum_{i=1}^n \widehat{b}_i = O(n)$ such that $\sum_{b_i > \sqrt{\tau}\widehat{b}_i} b_i = O(n/\tau)$. Then apply two-pass scattering in time with phase count τ to the primary input $X_1, \ldots, X_n$ and the profile input $\widehat{b}_1, \ldots, \widehat{b}_n$ and let $\mathcal{X}'$ and $\mathcal{X}''$ be the resulting sets of successful and unsuccessful elements, respectively. It follows almost exactly as in the proof of Theorem 8.15 that $|\mathcal{X}''| = O(n/\tau)$ with high probability. We can hence use the algorithms of Theorem 4.6 and Lemma 8.1 to store $\mathcal{X}''$ in an array of size $O(n/\tau)$ and to replace the values of elements in $\mathcal{X}''$ by values in a range of size $O(n/\tau)$, after which we can use the algorithm of Corollary 9.18 to compute a profile $\widehat{b}_1'', \ldots, \widehat{b}_n''$ such that $\sum_{i=1}^n \widehat{b}_i'' = O(n/\tau)$, but $\widehat{b}_i'' \geq |\mathcal{B}_i \cap \mathcal{X}''|$, for $i = 1, \ldots, n$ (take $\widehat{b}_i'' = 0$ for each $i \in \{1, \ldots, n\}$ with $\mathcal{B}_i \cap \mathcal{X}'' = \emptyset$).

Now draw a random sample $\mathcal{Y}$ from $\mathcal{X}'$ (not from $\mathcal{X}$) by including each element of $\mathcal{X}'$ in $\mathcal{Y}$ with probability $1/\tau$ and independently of other elements. By Chernoff bound (a), with high probability $|\mathcal{Y}| = O(n/\tau)$. Exactly as described for $\mathcal{X}''$ above, we can compute a profile $\widehat{b}_1^{\mathcal{Y}}, \ldots, \widehat{b}_n^{\mathcal{Y}}$ such that $\sum_{i=1}^n \widehat{b}_i^{\mathcal{Y}} = O(n/\tau)$, but $\widehat{b}_i^{\mathcal{Y}} \geq |\mathcal{B}_i \cap \mathcal{Y}|$, for $i = 1, \ldots, n$. For $i = 1, \ldots, n$, let $b_i' = |\mathcal{B}_i \cap \mathcal{X}'|$, $b_i'' = |\mathcal{B}_i \cap \mathcal{X}''|$ and $b_i^{\mathcal{Y}} = |\mathcal{B}_i \cap \mathcal{Y}|$. For $i = 1, \ldots, n$, if $\widehat{b}_i \leq \tau^2$, then take $\widetilde{b}_i = |\mathcal{B}_i \cap \mathcal{X}'| + \widehat{b}_i'' \geq b_i' + b_i'' = b_i$; otherwise take $\widetilde{b}_i = \max\{2\tau\widehat{b}_i^{\mathcal{Y}}, \tau^2\} + \widehat{b}_i''$.

Property (A) is satisfied, since $\sum_{i=1}^{n} \widetilde{b}_i \leq \sum_{i=1}^{n} (|\mathcal{B}_i \cap \mathcal{X}'| + 2\tau \widehat{b}_i^{\mathcal{Y}} + \widehat{b}_i'') + n = O(|\mathcal{X}'| + 2\tau(n/\tau) + |\mathcal{X}''| + n) = O(n)$. As for property (B), fix $i \in \{1, \ldots, n\}$ and note first that we cannot have $b_i > \widetilde{b}_i$ unless $b_i' > \max\{2\tau \widehat{b}_i^{\mathcal{Y}}, \tau^2\}$. But $b_i' > 2\tau \widehat{b}_i^{\mathcal{Y}}$ implies $b_i^{\mathcal{Y}} < b_i'/(2\tau)$, which under the condition $b_i' > \tau^2$ happens with probability at most $e^{-b_i'/(8\tau)} \leq e^{-\tau/8}$, by Chernoff bound (b). The desired result now follows by a martingale argument. ∎

Theorem 9.21: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$ with $\tau \geq \log^* n$, n-color semisorting problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Observe first that it suffices to partition the input into two subsets and to semisort these into arrays Q' and Q'' of size $O(n)$ each. For then, as in the proof of Corollary 9.18, we can divide Q' into nonoverlapping subarrays $A_1', \ldots, A_n'$ and Q'' into nonoverlapping subarrays $A_1'', \ldots, A_n''$ such that each element of $\mathcal{B}_i$ is stored either in A_i' or in A_i'', for $i = 1, \ldots, n$, after which we can use the algorithm of Theorem 7.1 to allocate an array A_i of size $|A_i'| + |A_i''|$ to $\mathcal{B}_i$ from a base array of size $\sum_{i=1}^{n} (|A_i'| + |A_i''|) = O(n)$ and store all elements of $\mathcal{B}_i$ in A_i, for $i = 1, \ldots, n$.

By this observation, it suffices to semisort the n input elements with a "waste" of $O(n/\tau)$ elements, i.e., with $O(n/\tau)$ elements not placed in the output array. This is because the elements that could not be placed are sufficiently few to be semisorted by the algorithm of Lemma 9.7 (following a compaction according to Theorem 4.6), after which we are in the situation described above.

As usual, assume that $\tau \leq n^{1/4}$ and that $b_i \leq n^{1/8}$, for $i = 1, \ldots, n$. The algorithm begins by computing a profile $\widetilde{b}_1, \ldots, \widetilde{b}_n$ for $x_1, \ldots, x_n$ with the properties described in Lemma 9.20, after which it applies one-pass scattering in time with phase count τ to the primary input $X_1, \ldots, X_n$ and the profile input $\widetilde{b}_1, \ldots, \widetilde{b}_n$. Similarly as in the proof of Theorem 8.15, call $\mathcal{B}_i$ *well-estimated* if $b_i \leq \widetilde{b}_i$, and call each element of $\mathcal{B}_i$ *good* if $\mathcal{B}_i$ is well-estimated, and *bad* otherwise, for $i = 1, \ldots, n$. Our first source of "waste" are the bad elements; by property (B) of Lemma 9.20, their number is $O(n/\tau)$ with high probability. A second source of "waste" are the colliding good elements. Since a good element collides with probability at most $1/\tau$, a martingale argument shows that the number of colliding good elements is also $O(n/\tau + n^{3/4}) = O(n/\tau)$ with high probability. A third and last source of "waste" will be good elements that cannot be placed in the output array although they did not collide. We now describe a procedure that uses the output of the scattering in time to semisort most of the noncolliding good elements.

Recall that scattering in time with profile input $\widetilde{b}_1, \ldots, \widetilde{b}_n$, as described in Section 8, uses arrays $A_1, \ldots, A_n$ of list headers and counters allocated within a base array of size $O(n)$, where $|A_i| = \widetilde{b}_i$, for $i = 1, \ldots, n$. For $i = 1, \ldots, n$, divide A_i into $\lceil \widetilde{b}_i/\tau \rceil$ *segments*,

each of size at most τ, and say that an element is stored in a segment if it belongs to a list whose header is stored in (a cell in) the segment. Further take $b'_i = |\mathcal{B}_i \cap \mathcal{X}'|$, where $\mathcal{X}'$ is the set of noncolliding elements.

Now associate a *target array* with each segment as follows: For $i = 1, \ldots, n$, if $\widetilde{b}_i \leq \tau$, then the target array of the (single) segment of A_i is of size $\min\{b'_i, 2\tau\}$ (as argued in Section 8, this quantity is readily available). If $\widetilde{b}_i > \tau$, on the other hand, the target array of each segment of A_i is of size 2τ, and the target arrays of all segments of A_i form a contiguous block of memory cells — this is easy to ensure, since they are all of the same size. Note that the total size of the target arrays is $O(n)$, so that they can be allocated according to Theorem 7.1 from a base array of size $O(n)$, which will be the output array of the semisorting.

We finally associate with each segment the task of placing $\min\{m, s\}$ elements stored in the segment in its target array, where m is the number of elements stored in the segment and s is the size of its target array, and execute all the tasks using operation allocation, as described in Section 7; if we take the length of a task to be the sum of the size of its associated segment and the size of the corresponding target array, the necessary prerequisites are easily seen to be satisfied (since every task is of length at most 3τ, it suffices to show how to process a task in linear sequential time, which is straightforward).

We want to show that with high probability, the number of elements not placed in the corresponding target arrays in the above computation is $O(n/\tau)$. To this end note that the choice of a list number in the scattering in time implicitly is a choice of a segment, and that the elements stored in a segment can be placed in the corresponding target array if their number is no larger than the size of the target array, i.e., if their number is at most 2τ. The expected number of elements of a well-estimated color class $\mathcal{B}_i$ choosing a particular segment is at most τ (since $b_i \leq \widetilde{b}_i$, the number of lists associated with $\mathcal{B}_i$ at least equals the number of elements in $\mathcal{B}_i$). Hence by Chernoff bound (a), the probability that a fixed element of a well-estimated color class $\mathcal{B}_i$ finds itself in a segment containing $2\tau + 1$ or more elements of $\mathcal{B}_i$ is at most $e^{-\tau/3}$. A martingale argument now shows that with high probability, the number of noncolliding good elements that cannot be placed in the appropriate target arrays is $O(n/\tau)$. ∎

10 Applications of Semisorting

This section describes a few relatively straightforward applications of Theorem 9.21. A number of less immediate applications were mentioned in the introduction.

Our first goal is to extend the semisorting result to strong semisorting. Recall that whereas usual semisorting places the elements of each color class in a subarray of a base array, strong semisorting additionally requires the size of the subarray of each color

class to be proportional to the size of the color class, a property that is often useful in applications.

Going from usual semisorting to strong semisorting obviously is a matter of compacting each color class into linear space. Treating color classes independently, we can use the algorithm of Theorem 5.3 to choose a suitable size for the destination array of each color and carry out the actual compaction using the algorithms of Section 4. Since color classes may be small, however, their sizes may be overestimated (as well as underestimated) by the algorithm of Theorem 5.3; as a result, although the compaction of a color class succeeds, it may fail in the sense that the destination array is too large. Since this is an infrequent event, we have enough resources to retry each unsuccessful compaction many times, which achieves a high reliability. An indispensable prerequisite for this, however, is the ability to tell whether a particular compaction was indeed into linear space. We therefore need a *certified approximate counting* algorithm that with high probability estimates the number of ones among n bits correctly, up to a constant factor, and that explicitly reports failure if it is unable to do so, i.e., a Las Vegas algorithm for approximate counting (Lemma 10.3). Our idea for obtaining such an algorithm is simple: Compacting the ones in the input into an array A furnishes a proof that their number b is at most $|A|$. On the other hand, subsequently compacting the free cells in A into an array Q proves that $|A| - b \leq |Q|$, which yields a lower bound on b.

We already know how to compact b elements into an array of size cb, where c is a constant. It turns out, however, that for the above scheme to work we cannot allow c to be arbitrarily large; in fact, we must demand that $c < 2$. We therefore briefly depart from our usual philosophy of ignoring constant factors to show that the relevant result in Section 4 (Theorem 4.6) actually holds for any constant $c > 1$ (Lemma 10.1). Observations similar to Lemmas 10.1 and 10.2 were made independently and first reported by Goodrich (1991).

Lemma 10.1: For every fixed $\mu > 0$ there is a constant $\epsilon > 0$ such that for all given $n, d, \tau \in \mathbb{N}$ with $\tau \geq \log^* n$, complete compaction problems of size n and with parameters $d \underset{(1+\mu)d}{\longrightarrow} 0$ can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Without loss of generality we can assume that $\mu d \geq 12$, since otherwise the number of active input elements is bounded by a constant, that $d \leq n$ and that μ is rational and at most 1. It suffices to describe a basic algorithm with a failure probability of $2^{-d^{\Omega(1)}}$, since for $d \leq \sqrt{n}$ the active elements can be compacted into an array of size $O(\sqrt{n})$ using the algorithm of Corollary 4.3, after which the basic algorithm can be applied independently $\Theta(\sqrt{n})$ times. It also suffices, for a certain constant $K \in \mathbb{N}$, to place all except $\mu d/K$ elements in an array of size $s = \lceil (1 + \mu/2)d \rceil$, since, provided

that K is sufficiently large, the algorithm of Theorem 4.6 can then be used to place the remaining elements in an array of size $\lceil \mu d/3 \rceil$, which for $\mu d \geq 12$ is at most $(1+\mu)d - s$.

We do this using repeated 1-scattering over a fixed array A of size s. Initially let all elements be active, and then carry out a number of stages. In each stage the remaining active elements are 1-scattered over A, whereby colliding elements as well as elements that hit an element placed in a previous stage remain active, while the other elements are placed in A and become inactive.

Assume that some stage starts with more than $\mu d/K$ active elements. It is easy to see that a fixed element collides or hits an element placed in a previous stage with probability at most $d/s \leq 1/(1+\mu/2)$, so that the expected number of elements deactivated in the stage is at least $(\mu d/K)(1 - 1/(1+\mu/2)) = \mu^2 d/(2K(1+\mu/2)) \geq (\mu^2/(4K)) \cdot d$. By Lemma 2.2(b), with high probability the stage under consideration deactivates at least $(\mu^2/(8K)) \cdot d$ elements. We may conclude that with high probability, $\lceil 8K/\mu^2 \rceil$ stages suffice to reduce the number of active elements to at most $\mu d/K$, as desired. ∎

Lemma 10.2: For every fixed $\mu > 0$ there is a constant $\epsilon > 0$ such that for all given $n, \tau \in I\!N$ and $\tau \geq \log^* n$, the following problem can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo): Given n bits $x_1, \ldots, x_n$, compute a nonnegative integer $\widehat{b}$ such that $b \leq \widehat{b} \leq (1+\mu)b$, where $b = |\{j : 1 \leq j \leq n \text{ and } x_j = 1\}|$.

Proof: Define an input element to be *active* if its value is 1, and assume without loss of generality that μ is rational and at most 1. Take $\theta = \mu/3$ and begin by using the algorithm of Theorem 5.3 to compute an integer $\widetilde{b}$ such that with high probability, $b/K \leq \widetilde{b} \leq b$, for some constant $K \geq 1$. If $\theta\widetilde{b} \leq 1$, solve the problem in a trivial manner. Otherwise repeatedly use the algorithm of Lemma 10.1 with $\mu = \theta$ to attempt to compact the active elements with limit $d = \widetilde{b}, \widetilde{b} + \lfloor \theta\widetilde{b} \rfloor, \widetilde{b} + 2\lfloor \theta\widetilde{b} \rfloor, \ldots$, stopping after the first successful complete compaction, and return as $\widehat{b}$ the quantity $\lfloor (1+\theta)d \rfloor$, where d is the limit of the last (successful) attempt.

The size of the destination array of the successful compaction is at most $(1+\theta)d$, i.e., the relation $b \leq \widehat{b}$ is satisfied. On the other hand, the compaction will succeed with high probability for any limit which is at least b. Provided that indeed $\widetilde{b} \leq b$, the first limit with this property in the above series is at most $b + \theta\widetilde{b} \leq (1+\theta)b$, so that with high probability, $\widehat{b} \leq (1+\theta)^2 b \leq (1+3\theta)b = (1+\mu)b$. It is easy to see from this that provided that indeed $\widetilde{b} = \Omega(b)$, with high probability the algorithm of Lemma 10.1 is applied only a constant number of times, i.e., the running time is $O(\tau)$. ∎

Informally, the "true" output of the Las Vegas algorithm below for approximate counting is the integer $\widehat{b}$. $y = 0$ indicates the correctness of the output, whereas $y = 1$ signifies that the execution failed.

Lemma 10.3: For every fixed $\mu > 0$ there is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$ with $\tau \geq \log^* n$, the following problem can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space: Given n bits $x_1, \ldots, x_n$, compute a nonnegative integer $\widehat{b}$ and a bit y such that

(a) If $y = 0$, then $b \leq \widehat{b} \leq (1+\mu)b$, where $b = \sum_{j=1}^{n} x_j$;

(b) $\Pr(y = 1) \leq 2^{-n^\epsilon}$.

Proof: Assume that μ is rational, choose $\theta < 1$ to make $(1-4\theta)(1+\mu) = 1$ and begin by applying the algorithm of Lemma 10.2 to obtain a nonnegative integer $\widetilde{b}$ such that with high probability, $b \leq \widetilde{b} \leq b/(1-\theta)$. Taking $\widehat{b} = \lfloor (1+\theta)\widetilde{b} \rfloor$, we now verify the two inequalities $b \leq \widehat{b} \leq (1+\mu)b$ and set $y = 1$ if the verification fails. Assume that $\widetilde{b} \geq 1$, since for $\widetilde{b} = 0$ the verification can be done trivially according to Lemma 2.7.

Again define an input element to be active if its value is 1. Let A be an array of size $\widehat{b}$ and use the algorithm of Lemma 10.1 with $\mu = \theta$ and $d = \widetilde{b}$ to attempt to place the active elements in A. If this succeeds, it clearly proves that $b \leq \widehat{b}$. On the other hand, since $\widetilde{b} \geq b$ with high probability, the compaction succeeds with high probability.

Assuming that the compaction into A succeeds, we next use the algorithm of Lemma 10.1 with $\mu = 1$ and $d = \lfloor 2\theta\widehat{b} \rfloor$ to attempt to place the free cells in A in an array Q of size $\lfloor 4\theta\widehat{b} \rfloor$. More precisely, this entails deriving from A a bit sequence $x'_1, \ldots, x'_{|A|+n}$ such that $x'_j = 1$ if and only if the jth cell of A contains no input element, for $j = 1, \ldots, |A|$, and $x'_j = 0$ for $j = |A|+1, \ldots, |A|+n$, and then using $x'_1, \ldots, x'_{|A|+n}$ as input to the algorithm of Lemma 10.1 ($x'_{|A|+1}, \ldots, x'_{|A|+n}$ are added only to ensure that the algorithm works correctly with high probability). Take $y = 0$ if and only if both compactions according to Lemma 10.1 succeed.

The following happens with high probability: $b \geq (1-\theta)\widetilde{b}$, so the number $\widehat{b} - b$ of free cells in A is at most $(1+\theta)\widetilde{b} - (1-\theta)\widetilde{b} = 2\theta\widetilde{b} \leq 2\theta\widehat{b}$, and the compaction into Q succeeds. If it does, this is proof that the number of free cells in A is at most $|Q|$, and hence that $b \geq \widehat{b} - \lfloor 4\theta\widehat{b} \rfloor \geq (1-4\theta)\widehat{b}$, from which follows that $\widehat{b} \leq (1+\mu)b$. ∎

Theorem 10.4: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$ with $\tau \geq \log^* n$, n-color strong semisorting problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Assume that $\tau \leq n^{1/4}$. We begin by semisorting the input according to Theorem 9.21. As in the proof of Corollary 9.18, we can view this as providing us with n disjoint subarrays $A_1, \ldots, A_n$ of a base array A of size $O(n)$ such that the elements of $\mathcal{B}_i$ are placed in A_i, for $i = 1, \ldots, n$. Our goal is to move the elements in $\mathcal{B}_i$ from A_i to a subarray of A_i of size $O(b_i)$, for $i = 1, \ldots, n$, which provides a solution to the strong semisorting problem.

We process $A_1, \ldots, A_n$ using operation allocation, as described in Section 7. The sequential processing of an array is simply exact compaction by means of prefix summation. The parallel processing of A_i is as follows, for $i = 1, \ldots, n$: Apply the algorithm of Lemma 10.3 to A_i with $\mu = 1$ to obtain a pair $(\widehat{b}_i, y_i)$, where $\widehat{b}_i$ is an estimate of b_i and y_i is an indication of the validity of $\widehat{b}_i$ (if $y_i = 0$, then $b_i \leq \widehat{b}_i \leq 2b_i$). Subsequently apply the algorithm of Lemma 10.1 with $\mu = 1$ to attempt to place $\mathcal{B}_i$ in a subarray A_i' of A_i of size at most $2\widehat{b}_i$ (if $2\widehat{b}_i \geq |A_i|$, simply take $A_i' = A_i$). If either $y_i = 1$ or the compaction of $\mathcal{B}_i$ into A_i' fails, we will say that the processing of A_i fails. Take $y_i' = 1$ if the processing of A_i fails, and $y_i' = 0$ otherwise.

By Lemmas 10.1 and 10.3, the processing of A_i fails with probability at most $2 \cdot 2^{-|A_i|^\delta}$, for some fixed $\delta > 0$ and for $i = 1, \ldots, n$. In particular, with high probability the processing of an array of size $n^{1/8}$ or more does not fail. As another consequence, $E(\sum_{i=1}^n y_i' \cdot 2^{|A_i|^\delta}) = O(n)$. Furthermore, by a martingale argument, $\sum_{i=1}^n y_i' q_i = O(n)$ with high probability, where $q_i = \min\{2^{\lceil |A_i|^\delta \rceil}, \lceil n^{1/4} \rceil\}$, for $i = 1, \ldots, n$. But this means that if the processing of A_i fails, for some $i \in \{1, \ldots, n\}$, then we can expend $\theta(q_i)$ operations in a second attempt to process A_i. We again use operation allocation, now with a new collection of tasks. Since $|A_i| \leq n^{1/8}$ with high probability, we can clearly compact $\mathcal{B}_i$ exactly in $O(q_i)$ sequential time. Furthermore, if $q_i = 2^{\lceil |A_i|^\delta \rceil}$, we can use prefix summation (Lemma 2.4) to compact $\mathcal{B}_i$ in $O(\tau)$ time using $\lceil q_i/\tau \rceil$ processors, while if instead $q_i = \lceil n^{1/4} \rceil$, $\lceil q_i/\tau \rceil$ processors suffice to carry out $\Theta(n^{1/8})$ independent attempts to process A_i in $O(\tau)$ time as above, at least one of which will succeed with high probability. ∎

Corollary 10.5: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in \mathbb{N}$ with $\tau \geq \log^* n$, n-color fine-profiling problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Immediate from Theorem 10.4 and Lemma 2.8(a) (see the proof of Corollary 9.18). ∎

A second application of semisorting is to integer chain-sorting. Recall that the chain-sorting problem is to store given keys in sorted order in a linked list. In the formal definition below, the linked list is represented by a circular structure and a pointer to the last list element.

Definition: For all $n, m \in \mathbb{N}$, the *m-color chain-sorting* problem of size n is the following: Given n integers $x_1, \ldots, x_n$ in the range $1..m$, compute a cyclic permutation $\pi_1, \ldots, \pi_n$ of $1, \ldots, n$ and an integer $q \in \{1, \ldots, n\}$ such that for all $j \in \{1, \ldots, n\} \backslash \{q\}$, we have $x_{\pi_j} \geq x_j$.

Theorem 10.6: There is a constant $\epsilon > 0$ such that for all given $n, \tau \in I\!N$ with $\tau \geq \log^* n$, n-color chain-sorting problems of size n can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Begin by semisorting the input elements into an array of size $O(n)$ according to Theorem 9.21. Using the algorithm of Lemma 2.8(a), it is then easy to construct a linked list containing precisely the elements of $\mathcal{B}_i$ in the order in which they occur in A, for $i = 1, \ldots, n$. The remaining problem is to concatenate these lists in the right order. This can be done by applying the algorithm of Lemma 2.8(a) a second time, now to an n-bit input whose ith bit is 1 if and only if $\mathcal{B}_i \neq \emptyset$, for $i = 1, \ldots, n$. ∎

A claim similar to Theorem 10.6 above was made in (Gil *et al.*, 1991). It seems unlikely, however, that any algorithm based on the outline given in (Gil *et al.*, 1991) can be made to run in linear space.

An important application of Theorem 10.6 is to (standard) integer sorting. Let us restrict attention to the problem of sorting n integers in the range $1..n$ on a CRCW PRAM. Rajasekaran and Reif (1989) describe a randomized algorithm with optimal speedup for this problem that uses $O(\log n)$ time and $O(n/\log n)$ processors with high probability. Bhatt *et al.* (1991) give a deterministic algorithm that works in $O(\log n/\log\log n)$ time using $O(n(\log\log n)^2/\log n)$ processors. We show how to combine the time bound of (Bhatt *et al.*, 1991) with the time-processor product of (Rajasekaran and Reif, 1989), thus achieving at the same time optimal speed and optimal speedup. Similar results were found independently by Matias and Vishkin (1991) and Raman (1991); note, however, that the algorithms of these authors (which are quite similar) are inherently much less reliable than the algorithm given here — the failure probability is $\Omega(2^{-(\log n)^\alpha})$, for some fixed α, to be contrasted with our failure probability of $2^{-n^{\Omega(1)}}$.

Our algorithm makes use of a subroutine for monotonic list ranking with optimal speedup. The *monotonic list ranking* problem of size n is, given a linked list of n elements such that which of two given list elements precedes the other can be determined in constant time by a single processor, to mark each element of the list with its position within the list.

Lemma 10.7 (Bhatt *et al.*, 1991): For all given integers $n \geq 4$ and $\tau \geq \log n/\log\log n$, monotonic list ranking problems of size n can be solved on a (deterministic) TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space.

Theorem 10.8: There is a constant $\epsilon > 0$ such that for all given integers $n \geq 4$ and $\tau \geq \log n/\log\log n$, n integers in the range $1..n$ can be sorted on a TOLERANT PRAM using $O(\tau)$ time, $\lceil n/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: Chain-sort the input elements using the algorithm of Theorem 10.6 and compute the position of each element within the resulting list using the algorithm of Lemma 10.7. In order to determine the relative order of two elements with the same value, compare their positions in the semisorted array output by the algorithm of Theorem 9.21. ∎

11 Nonoptimal Algorithms

This section investigates the effect for the problems considered of allowing slightly superlinear processor and space bounds. In some cases, we also have to generalize the problems by introducing a so-called *slack parameter* (this notion already appeared in Lemmas 6.1 and 7.3). We begin by showing that compaction with slack can be done in constant time.

Although, technically, the results stated in this section allow k and τ to vary independently as functions of n, it is probably most useful to imagine that $\tau = k$ is constant. Our informal discussion makes this assumption.

Theorem 11.1: There is a constant $\epsilon > 0$ such that for all given $n, d, k, \tau \in \mathbb{N}$ with $\tau \geq k$, complete compaction problems of size n and with parameters $d \underset{O(s)}{\longrightarrow} 0$, where $s = d\lfloor \log^{(k)} n \rfloor$, can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil kn/\tau \rceil$ processors and $O(n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: We can assume that $(2\lceil \log^{(k+1)} d \rceil)^3 \leq \log^{(k)} d$, since otherwise $k = \Omega(\log^* d)$ and we can apply the algorithm of Theorem 4.6, and that $s \leq n$. Then apply the algorithm of Lemma 4.4 $O(k)$ times. The number of operations needed is $O(kn)$, which translates into $\lceil kn/\tau \rceil$ processors, for any $\tau \geq k$. Omitting the size σ of the destination array from the notation $d_1 \underset{\sigma}{\longrightarrow} d_2$, we can express the process symbolically as follows:

$$d \leq \frac{d\lfloor \log^{(k)} d \rfloor}{(2\lceil \log^{(k+1)} d \rceil)^3} \longrightarrow \frac{d\lfloor \log^{(k)} d \rfloor}{2^{6\lceil \log^{(k+1)} d \rceil}}$$

$$\leq \frac{d}{\lceil \log^{(k)} d \rceil^3} \longrightarrow \frac{d}{\lceil \log^{(k-1)} d \rceil^3} \longrightarrow \cdots \longrightarrow \frac{d}{\lceil \log d \rceil^3} \longrightarrow \frac{d}{d^3}.$$

The last step in the above sequence reduces the number of active elements below 1, i.e., to zero. The destination array used in the first step is of size $O(d \log^{(k)} d) = O(s)$, and the sizes of the destination arrays used in the remaining steps sum to $O(d)$. Hence all active elements can indeed be placed in an array of size $O(s)$. ∎

We now extend Theorem 11.1 to the case of several colors. In contrast with the algorithm of Theorem 11.1, the generalized algorithm of Theorem 11.2 needs superlinear space.

Definition: For all $n, m \in \mathbb{N}$, $d_1, \ldots, d_m \geq 0$ and $\lambda \geq 1$, the complete colored compaction problem of size n and with limits $d_1, \ldots, d_m$ and slack λ is, given n integers $x_1, \ldots, x_n$ in the range $0..m$ such that $|\{j : 1 \leq j \leq n \text{ and } x_j = i\}| \leq d_i$, for $i = 1, \ldots, m$, to compute a complete placement for $x_1, \ldots, x_n$ with bounds $\lambda d_1, \ldots, \lambda d_m$.

Theorem 11.2: There is a constant $\epsilon > 0$ such that for all given $n, m, k, \tau, d_1, \ldots, d_m \in \mathbb{N}$ with $m = (\log n)^{O(1)}$ and $\tau \geq k$, complete colored compaction problems of size n with limits $d_1, \ldots, d_m$ and with slack $O(\log^{(k)} n)$ can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil kn/\tau \rceil$ processors and $O((n + \sum_{i=1}^{m} d_i) \log^{(k)} n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: The idea of the proof is to apply the nonoptimal part of the algorithm of Theorem 6.2 (with $\delta = 1/2$) in a situation in which some of the elements have already been deactivated. Define $\mathcal{B}_i$ to be *large* if $b_i > n^{1/4}$, for $i = 1, \ldots, m$, and recall that the algorithm of Theorem 6.2 essentially applies the algorithm of Lemma 6.1 $\log^* n$ times to reduce the fraction of active elements in each large color class below a certain threshold, after which the compaction is finished using negligible resources. In the present setting, where we are allowed $O(\log^{(k)} n)$ slack, we can speed up the deactivation by first 1-scattering the elements of $\mathcal{B}_i$ over an array of size $2d_i\lceil \log^{(k)} n \rceil$, for $i = 1, \ldots, m$. Lemma 3.6(b) shows that the number of elements in a fixed large color class $\mathcal{B}_i$ that collide in the 1-scattering is at most $\max\{b_i/\log^{(k)} n, b_i^{3/4}\}$, except with probability at most $e^{-\zeta}$, where $\zeta = (b_i^{3/4})^2/(32b_i) = \Omega(b_i^{1/2}) = \Omega(n^{1/8})$, so that with high probability the fraction of active elements left in any large color class is $O(1/\log^{(k)} n)$. It is now easy to see that all but the last $O(k)$ applications of the algorithm of Lemma 6.1 in the algorithm of Theorem 6.2 can be omitted. Since all subroutines used can be made to run in $O(\tau/k)$ time using $\lceil kn/\tau \rceil$ processors, we can therefore deactivate all elements in $O(\tau)$ time. ∎

Armed with Theorem 11.2, we can easily use the reductions of interval allocation and interval marking to colored compaction given in Section 7 to derive similar results for these problems.

Definition: For all $n \in \mathbb{N}$ and $\lambda \geq 1$, the interval allocation problem of size n and with slack λ is, given n nonnegative integers $x_1, \ldots, x_n$, to compute a complete interval placement for $x_1, \ldots, x_n$ with slack λ.

Theorem 11.3: There is a constant $\epsilon > 0$ such that for all given $n, k, \tau \in \mathbb{N}$ with $\tau \geq k$, interval allocation problems of size n and with slack $O(\log^{(k)} n)$ can be solved on a TOLERANT PRAM using $O(\tau)$ time, $\lceil kn/\tau \rceil$ processors and $O(n \log^{(k)} n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: As the proof of Theorem 7.1, except that Theorem 11.2 is used instead of Theorem 6.2, and that Corollary 10.5 is used instead of Theorem 5.3 (the latter substitution serves only to obtain a Las Vegas algorithm). ∎

In (Bast *et al.*, 1992), Theorem 11.3 is used to prove a related result: For any $k \in \mathbb{N}$, usual interval allocation problems (i.e., with constant slack) of size n can be solved in $O(k)$ time using $O(n \log^{(k)} n)$ processors and $O(n \log^{(k)} n)$ space with high probability.

Definition: For all $n \in \mathbb{N}$ and $\lambda \geq 1$, the interval marking problem of size n and with slack λ is the following: Given n nonnegative integers $x_1, \ldots, x_n$, compute nonnegative integers $s, z_1, \ldots, z_s$ such that

(1) For all integers i, j, k with $1 \leq i \leq j \leq k \leq s$, if $z_i = z_k \neq 0$, then $z_j = z_i$;

(2) For $i = 1, \ldots, n$, $|\{j : 1 \leq j \leq s \text{ and } z_j = i\}| = x_i$;

(3) $s = O(\lambda \sum_{j=1}^{n} x_j)$.

Theorem 11.4: There is a constant $\epsilon > 0$ such that for all given $n, k \in \mathbb{N}$, interval marking problems of size n and with slack $O(\log^{(k)} n)$ can be solved on a TOLERANT PRAM using $O(k)$ time, $O(n + W \log^{(k)} n)$ processors and $O((n + W) \log^{(k)} n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas), where W is the sum of the input numbers.

Proof: As the proof of Theorem 7.2, using Theorem 11.3 instead of Theorem 7.1. ∎

Our goal in the remainder of this section is to derive constant-time algorithms for coarse-profiling and for a variant of semisorting. Recall from Sections 8 and 9 that the n-processor algorithms for these problems proceed in a number of stages, each except the last of which essentially performs bootstrapping for the following stage, while only the last stage actually solves the entire problem. The key observation for the present section is that if we simply omit the bootstrapping of the first stages, the algorithms still operate in a well-defined way, but with a certain increase in their resource requirements and, possibly, a certain degradation in the quality of their output. It even turns out that if we start with Stage t_0, for some t_0, the bootstrapping effect of Stage t_0 is not affected by the absence of Stages $1, \ldots, t_0 - 1$, so that our original analysis applies without modification to all stages following Stage t_0. It therefore suffices to reanalyze Stage t_0 with respect to its resource requirements and its effect on the output.

Theorem 11.5: There is a constant $\epsilon > 0$ such that for all given $n, k \in \mathbb{N}$, n-color coarse-profiling problems of size n can be solved on a TOLERANT PRAM using $O(k^2)$ time, $O(n \log^{(k)} n)$ processors and $O(n \log^{(k)} n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo).

Proof: Execute only Stages $t_0, \ldots, T$ of the algorithm of Lemma 8.3, where $t_0 = \max\{1, T - 2(k+1)\}$, and note that $v_{t_0} = O(\log^{(k+1)} n)$ (because $v_{t-2} \leq \log v_t$ for

t larger than some constant). The number of processors and memory cells allocated in Stage t_0 clearly is $O(nv_{t_0})$. By the proof of Lemma 8.6, the probability that a fixed color class is tardy in Stage t is still bounded by v_t^{-6}, for $t = t_0 + 1, \ldots, T$. Therefore the number of processors and memory cells allocated in each of Stages $t_0 + 1, \ldots, T$ remains $O(n)$ (Lemma 8.7). The amount of resources consumed is hence as desired, and a running time of $O(k^2)$ is achieved by allocating the resources needed in each of the $O(k)$ stages using the algorithms of Theorems 11.3 and 11.4.

Intuitively, the quality of an estimate is essentially independent of the stage in which it becomes defined. Correspondingly, it is not difficult to see that Lemmas 8.8 and 8.9 hold without change; hence the linear-sum condition is satisfied with high probability. To show property (B) of a coarse-profile, fix $a \geq 1$ and $i \in \{1, \ldots, n\}$ with $b_i > 0$. In the present setting in which Stages 1 to $t_0 - 1$ are omitted, the proof of Theorem 8.10 yields that $\sum_{t=t_0+1}^{T} \Pr(b_i > a\widehat{b}_i \text{ and } D_t) \leq 2^{-2a}$, where D_t denotes the event that $\widehat{b}_i$ becomes defined in Stage t, for $t = 1, \ldots, T$. Note that t_0 is excluded from the sum; this is necessary because small color classes are certain to be tardy in Stage t_0, thus invalidating the analysis in Case 3 of the proof of Lemma 8.10. By the analysis of Case 2, however, $\Pr(b_i > a\widehat{b}_i \text{ and } D_{t_0}) \leq 2^{-8a}$, allowing us to conclude that $\Pr(b_i > a\widehat{b}_i) \leq 2^{-2a} + 2^{-8a} \leq 2^{-a}$. The output of the modified algorithm therefore indeed is a coarse-profile. ∎

Theorem 11.5 represents the best that we can do on the TOLERANT PRAM; in particular, the number of processors needed is superlinear. On the ARBITRARY PRAM, on the other hand, we obtain a constant-time algorithm with optimal speedup with the sole drawback of superlinear space requirements.

Theorem 11.6: There is a constant $\epsilon > 0$ such that for all given $n, k \in \mathbb{N}$, n-color coarse-profiling problems of size n can be solved on an ARBITRARY PRAM using $O(k^2)$ time, n processors and $O(n \log^{(k)} n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Monte Carlo).

Proof: Consider the algorithm of Theorem 11.5 and note that the processors allocated by the algorithm serve exclusively to determine the 1-rows of a number of graduated conditional scatterings according to Lemma 3.4(a). By Lemma 3.4(b), on the ARBITRARY PRAM the same can be achieved with just one processor for each element scattered, plus one additional processor per GCS. Since these processors can be allocated in a trivial manner from a pool of $O(n)$ processors and the allocation of space can be done using the algorithm of Theorem 11.3, the claim follows. ∎

Definition: For all $n, m \in \mathbb{N}$ and $\lambda \geq 1$, the *m-color semisorting* problem of size n and with slack λ is the following: Given n integers $x_1, \ldots, x_n$ in the range $0 \,..\, m$, compute n

nonnegative integers $y_1, \ldots, y_n$ such that

(1) For $1 \leq i < j \leq n$, if $x_i \neq 0$, then $y_i \neq y_j$;

(2) For all $i, j, k \in \{1, \ldots, n\}$, if $y_i < y_j < y_k$ and $x_i = x_k$, then $x_j = x_i$;

(3) $\max\{y_j : 1 \leq j \leq n\} = O(\lambda n)$.

Theorem 11.7: There is a constant $\epsilon > 0$ such that for all given $n, k \in \mathbb{N}$, n-color semisorting problems of size n and with slack $O(\log^{(k)} n)$ can be solved on a TOLERANT PRAM using $O(k)$ time, $O(n \log^{(k)} n)$ processors and $O(n \log^{(k)} n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: As in the proof of Lemma 9.7, we first compute a profile with properties (A) and (B) of Lemma 9.1. The procedure is similar to that of the proof of Theorem 11.5: With $t_0 = \max\{1, T - 2(k+1)\}$, we execute only Stages $t_0, \ldots, T$ of the algorithm of Lemma 9.1. By the proof of Lemma 9.3, the number of tardy color classes in Stage t is still $O(n/v_t^4)$, for $t = t_0 + 1, \ldots, T$, so that Stages $t_0 + 1, \ldots, T$ can be executed using $O(n)$ processors and $O(n)$ space (Lemma 9.4). The amount of resources allocated in Stage t_0 is $O(nv_{t_0})$; hence, by Lemma 7.3, Stage t_0 can be executed using $O(nv_{t_0}^4) = O(n \log^{(k)} n)$ processors and memory cells. Furthermore, Lemma 9.5 still holds, so the linear-sum condition is satisfied. Using an argument in the proof of Theorem 11.5 showing that $\sum_{t=t_0}^{T} \Pr(b_i > a\widehat{b}_i \text{ and } D_t) \leq 2^{-a}$, the proof of Lemma 9.6 carries over to the modified algorithm, and condition (B) of Lemma 9.1 is satisfied as well.

Now execute Stages $t_0, \ldots, T$ of the algorithm of Lemma 9.7, still with $t_0 = \max\{1, T - 2(k+1)\}$. Although the series $v_1, \ldots, v_T$ is now different, it is again the case that $v_{t_0} = O(\log^{(k+1)} n)$. Recall that each stage deactivates elements by placing them in suitably-sized arrays, one for each color. In every stage, the size of the array used for a particular color is chosen on the basis of a test scattering for that color, which roughly estimates the number of remaining active elements of that color. As an important consequence of this "self-correcting" mechanism, we were able to analyze the deactivation capability of a stage without relying on the deactivation carried out in earlier stages (if earlier stages perform poorly, the resource requirements of the stage at hand go up, but it will still reduce the number of remaining active elements to the required level). Therefore Lemma 9.11 remains true (the last stage always deactivates all remaining active elements), and with the additional restriction $t \geq t_0$, the same holds for Lemmas 9.12–9.15. We must show that the algorithm is correct and bound its resource requirements. As in Section 9, this essentially boils down to bounding the total size of the arrays $A_{i,t}$ allocated in Stage t, for $t = t_0, \ldots, T$. For $t \geq t_0 + 1$, this quantity can be seen to be $O(n/v_t^3)$ (Lemma 9.16); the reason is that in the analysis of a particular stage, Lemmas 9.13–9.15 can be applied to the previous stage. As regards Stage t_0 itself, it is easy to see from lines (7) and (13)–(15) in the algorithm that the total size of the

arrays allocated in Stage t_0 is $O(nv_{t_0}^{19})$. Since $v_{t_0}^{22} = O(\log^{(k)} n)$, by the choice of t_0, it now follows essentially as in the proof of Lemma 9.17 that with high probability, the algorithm uses $O(n \log^{(k)} n)$ processors and $O(n \log^{(k)} n)$ space and solves the semisorting problem with slack $O(n \log^{(k)} n)$. ∎

Corollary 11.8: There is a constant $\epsilon > 0$ such that for all given $n, k \in \mathbb{N}$, n-color chain-sorting problems of size n can be solved on a TOLERANT PRAM using $O(k)$ time, $O(n \log^{(k)} n)$ processors and $O(n \log^{(k)} n)$ space with probability at least $1 - 2^{-n^\epsilon}$ (Las Vegas).

Proof: As the proof of Theorem 10.6, using Theorem 11.7 instead of Theorem 9.21 and part (b) of Lemma 2.8 instead of part (a). ∎

Acknowledgment: We are grateful to the anonymous referees for their meticulous reading and for many useful suggestions. We also thank Peter Miltersen for pointing us to the papers by Stockmeyer and Ajtai and Ben-Or, and Prabhakar Ragde for simplifying the proof of Lemma 3.5.

References

AJTAI, M., AND BEN-OR, M. (1984), A Theorem on Probabilistic Constant Depth Computations, *in* Proc. 16th Annual ACM Symposium on Theory of Computing, pp. 471–474.

ALON, N., AND MEGIDDO, N. (1990), Parallel Linear Programming in Fixed Dimension Almost Surely in Constant Time, *in* Proc. 31st Annual Symposium on Foundations of Computer Science, pp. 574–582.

BAST, H., DIETZFELBINGER, M., AND HAGERUP, T. (1992), A Perfect Parallel Dictionary, *in* Proc. 17th International Symposium on Mathematical Foundations of Computer Science, Springer Lecture Notes in Computer Science, Vol. 629, pp. 133–141.

BAST, H., AND HAGERUP, T. (1991), Fast and Reliable Parallel Hashing, manuscript. A preliminary version appears *in* Proc. 3rd Annual ACM Symposium on Parallel Algorithms and Architectures, pp. 50–61.

BEAME, P., AND HASTAD, J. (1989), Optimal Bounds for Decision Problems on the CRCW PRAM, *J. ACM* **36**, pp. 643–670.

BERKMAN, O., AND VISHKIN, U. (1989), Recursive *-Tree Parallel Data-Structure, *in* Proc. 30th Annual Symposium on Foundations of Computer Science, pp. 196–202.

BHATT, P. C. P., DIKS, K., HAGERUP, T., PRASAD, V. C., RADZIK, T., AND SAXENA, S. (1991), Improved Deterministic Parallel Integer Sorting, *Inform. and Comput.* **94**, pp. 29–47.

BOLLOBÁS, B. (1987), Martingales, Isoperimetric Inequalities and Random Graphs, *in* Colloq. Math. Soc. J. Bolyai 52, pp. 113–139.

CHLEBUS, B. S., DIKS, K., HAGERUP, T., AND RADZIK, T. (1988), Efficient Simulations between Concurrent-Read Concurrent-Write PRAM Models, *in* Proc. 13th Symposium on Mathematical Foundations of Computer Science, Springer Lecture Notes in Computer Science, Vol. 324, pp. 231–239.

CHLEBUS, B. S., DIKS, K., HAGERUP, T., AND RADZIK, T. (1989), New Simulations between CRCW PRAMs, *in* Proc. 7th International Conference on Fundamentals of Computation Theory, Springer Lecture Notes in Computer Science, Vol. 380, pp. 95–104.

COLE, R., AND VISHKIN, U. (1989), Faster Optimal Parallel Prefix Sums and List Ranking, *Inform. and Comput.* **81**, pp. 334–352.

FICH, F. E., RAGDE, P., AND WIGDERSON, A. (1988a), Simulations Among Concurrent-Write PRAMs, *Algorithmica* **3**, pp. 43–51.

FICH, F. E., RAGDE, P., AND WIGDERSON, A. (1988b), Relations Between Concurrent-Write Models of Parallel Computation, *SIAM J. Comput.* **17**, pp. 606–627.

GIL, J. (1990), Lower Bounds and Algorithms for Hashing and Parallel Processing, Ph. D. Thesis, The Hebrew University, Jerusalem.

GIL, J. (1991), Fast Load Balancing on a PRAM, *in* Proc. 3rd IEEE Symposium on Parallel and Distributed Processing, pp. 10–17.

GIL, J., MATIAS, Y., AND VISHKIN, U. (1991), Towards a Theory of Nearly Constant Time Parallel Algorithms, *in* Proc. 32nd Annual Symposium on Foundations of Computer Science, pp. 698–710.

GOODRICH, M. T. (1991), Using Approximation Algorithms to Design Parallel Algorithms that May Ignore Processor Allocation, *in* Proc. 32nd Annual Symposium on Foundations of Computer Science, pp. 711–722.

GROLMUSZ, V. (1991), Large Parallel Machines Can Be Extremely Slow for Small Problems, *Algorithmica* **6**, pp. 479–489.

GROLMUSZ, V., AND RAGDE, P. (1987), Incomparability in Parallel Computation, *in* Proc. 28th Annual Symposium on Foundations of Computer Science, pp. 89–98. Journal version: *Disc. Appl. Math.* **29** (1990), pp. 63–78.

HAGERUP, T. (1992a), Fast and Optimal Simulations between CRCW PRAMs, *in* Proc. 9th Annual Symposium on Theoretical Aspects of Computer Science, Springer Lecture Notes in Computer Science, Vol. 577, pp. 45–56.

HAGERUP, T. (1992b), The Log-Star Revolution, *in* Proc. 9th Annual Symposium on Theoretical Aspects of Computer Science, Springer Lecture Notes in Computer Science, Vol. 577, pp. 259–278.

HAGERUP, T. (1992c), The Parallel Complexity of Integer Prefix Summation, Tech. Rep. no. LSI–92–18–R, Dept. de LSI, Universitat Politècnica de Catalunya, Barcelona, Spain.

HAGERUP, T. (1992d), On a Compaction Theorem of Ragde, *Inform. Process. Lett.* **43**, pp. 335–340.

HAGERUP, T., AND KATAJAINEN, J. (1993), Improved Parallel Bucketing Algorithms for Proximity Problems, *in* Proc. 26th Hawaii International Conference on System Sciences, Vol. II: Software Technology, pp. 318–327.

HAGERUP, T., AND RADZIK, T. (1990), Every Robust CRCW PRAM Can Efficiently Simulate a PRIORITY PRAM, *in* Proc. 2nd Annual ACM Symposium on Parallel Algorithms and Architectures, pp. 117–124.

HAGERUP, T., AND RAMAN, R. (1992), Waste Makes Haste: Tight Bounds for Loose Parallel Sorting, *in* Proc. 33rd Annual Symposium on Foundations of Computer Science (1992), pp. 628–637.

HAGERUP, T., AND RÜB, C. (1990), A Guided Tour of Chernoff Bounds, *Inform. Process. Lett.* **33**, pp. 305–308.

MACKENZIE, P. D. (1992), Load Balancing Requires $\Omega(\log^* n)$ Expected Time, *in* Proc. 3rd Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 94–99.

MATIAS, Y., AND VISHKIN, U. (1991), Converting High Probability into Nearly-Constant Time — with Applications to Parallel Hashing, *in* Proc. 23rd Annual ACM Symposium on Theory of Computing, pp. 307–316.

MCDIARMID, C. (1989), On the Method of Bounded Differences, *in* Surveys in Combinatorics, ed. J. Siemons, London Math. Soc. Lecture Note Series 141, Cambridge University Press, pp. 148–188.

RAGDE, P. (1990), The Parallel Simplicity of Compaction and Chaining, *in* Proc. 17th International Colloquium on Automata, Languages and Programming, Springer Lecture Notes in Computer Science, Vol. 443, pp. 744–751.

RAJASEKARAN, S., AND REIF, J. H. (1989), Optimal and Sublogarithmic Time Randomized Parallel Sorting Algorithms, *SIAM J. Comput.* **18**, pp. 594–607.

RAMAN, R. (1990), The Power of Collision: Randomized Parallel Algorithms for Chaining and Integer Sorting, *in* Proc. 10th Conference on Foundations of Software Technology and Theoretical Computer Science, Springer Lecture Notes in Computer Science, Vol. 472, pp. 161–175.

RAMAN, R. (1991), Optimal Sub-logarithmic Time Integer Sorting on the CRCW PRAM, Tech. Rep. no. 370, Dept. of Computer Science, University of Rochester, Rochester, New York.

REISCHUK, R. (1985), Probabilistic Parallel Algorithms for Sorting and Selection, *SIAM J. Comput.* **14**, pp. 396–409.

SHILOACH, Y., AND VISHKIN, U. (1982), An $O(\log n)$ Parallel Connectivity Algorithm, *J. Alg.* **3**, pp. 57–67.

STOCKMEYER, L. (1983), The Complexity of Approximate Counting, *in* Proc. 15th Annual ACM Symposium on Theory of Computing, pp. 118–126.

VALIANT, L. G. (1990), General Purpose Parallel Architectures, *in* "Handbook of Theoretical Computer Science, Vol. A: Algorithms and Complexity" (J. van Leeuwen, ed.), pp. 943–971, Elsevier, Amsterdam.