A parallel algorithmic version of the Local Lemma

Noga Alon *
Department of Mathematics
Raymond and Beverly Sackler Faculty of Exact Sciences
Tel Aviv University, Tel Aviv, Israel

Abstract

The Lovász Local Lemma is a tool that enables one to show that certain events hold with positive, though very small probability. It often yields existence proofs of results without supplying any efficient way of solving the corresponding algorithmic problems. J. Beck has recently found a method for converting some of these existence proofs into efficient algorithmic procedures, at the cost of loosing a little in the estimates. His method does not seem to be parallelizable. Here we modify his technique in several points and achieve an algorithmic version that can be parallelized, thus obtaining deterministic NC^1 algorithms for various interesting algorithmic search problems. Two representing examples are the problem of finding a directed even simple cycle in a digraph in which the maximum outdegree is at most a small exponent of the minimum indegree and the problem of exhibiting a satisfying assignment for a Boolean CNF Formula in which the maximum number of occurences of each literal is bounded by a small exponent of the minimum number of literals in a disjunction.

1 Introduction

In a typical application of the probabilistic method we try to prove the existence of a combinatorial structure (or a substructure of a given structure) with certain prescribed properties. To do so, we show that a randomly chosen element from an appropriately defined sample space satisfies all the required properties with positive probability. In most applications, this probability is not only positive, but is actually high and frequently tends to 1 as the parameters of the problem tend to infinity. In such cases, the proof usually supplies an efficient randomized algorithm for producing a structure of the desired type, and in many cases this algorithm can be derandomized and converted into an efficient deterministic one. By efficient we mean here, as usual, an algorithm whose running time -(or expected running time, in case we consider randomized algorithms)- is polynomial in the length of the input.

There are, however, certain examples, where one can prove the existence of the required combinatorial structure by probabilistic arguments that deal with rare events; events that hold with positive probability which is exponentially small in the size of the input. Such proofs usually yield neither randomized nor deterministic efficient procedures for the corresponding algorithmic problems.

A class of examples demonstrating this phenomenon is the class of results proved by applying the Lovász Local Lemma, proved in [14]. The exact statement of this lemma (for the symmetric case) is the following.

Lemma 1.1 *Let $A_1, \ldots, A_n$ be events in an arbitrary probability space. Suppose that the probability of each of the n events is at most p, and suppose that each event A_i is mutually independent of all but at most b of the other events A_j. If $ep(b+1) < 1$ then with positive probability none of the events A_i holds.*

Many applications of this lemma can be found in [14], [21], [2], [4], [7], [8], [9], [10], [11]. For several years there has been no known method of converting the proofs of any of these examples into

*Research supported in part by a United States Israel BSF Grant and by a Bergmann Memorial Grant

an efficient algorithm, and the problem of finding such algorithms is mentioned, among other places, in [21]. Very recently J. Beck [12] found such a method that works for many of these examples with a little loss in the constants. He demonstrated his elegant method by considering the problem of hypergraph 2-coloring. A hypergraph $H = (V, E)$ is 2-*colorable* if there is a two coloring of V so that no edge in E is monochromatic. The following result is due to Erdös and Lovász. Its derivation from the Local Lemma is very simple [14].

Theorem 1.2 *If $e(d+1) < 2^{n-1}$ then any hypergraph H in which every edge has at least $n \geq 2$ vertices and no edge intersects more than d other edges is 2-colorable.*

Recall that a hypergraph is n-*uniform* if each of its edges contains precisely n vertices. Observe that there is no loss of generality in assuming that H is n-uniform in Theorem 1.2, since otherwise we can simply replace each edge of H by a subset of size n of it. Let H be an n-uniform hypergraph with N edges satisfying the assumptions in Theorem 1.2, and suppose, for simplicity, that n, d are fixed. Can we find a proper two coloring of H (i.e., a vertex coloring in which no edge is monochromatic) efficiently? Beck showed that indeed we can, in case d is somewhat smaller, say $d < O(2^{n/48})$. In this case there is a randomized as well as a deterministic algorithm whose running time is polynomial in N for finding a proper two coloring. Beck's method does not seem to provide a parallel efficient algorithm (i.e., an algorithm that runs in polylogarithmic time using a polynomial number of processors). Here we modify Beck's elegant idea and combine it with some additional ones to obtain a version which is parallelizable. The main ingredients in our algorithm that differ from those of the one in [12] are the recoloring phase (together with the definition of bad edges) and its analysis, and the application of the derandomization techniques dealing with (almost) d-wise independent events in small sample spaces.

We first describe the randomized version of our algorithm and then comment briefly on its derandomization and parallelization. Let us denote, as usual, the binary entropy function by $H(x) = -x \log_2 x - (1-x)\log_2(1-x)$.

Theorem 1.3 *Suppose $n \geq 2, d$ are fixed and suppose that for some $\alpha > 0$*

$$4ed^3 < 2^{n(1-H(\alpha))} \tag{1}$$

and

$$2e(d+1) < 2^{\alpha n}. \tag{2}$$

Then there is a randomized algorithm that finds a proper 2-coloring of any given n-uniform hypergraph H with N edges in which no edge intersects more than d others in expected running time $N^{O(1)}$. This algorithm can be derandomized and parallelized, providing a deterministic algorithm that finds a proper coloring in time $O(\log N)$ using $N^{O(1)}$ processors.

We note that for large n, any $d \leq 2^{n/8}$ satisfies the above (by taking an appropriate $\alpha > 1/8$). We also note that by assuming that d is smaller, say that $d < 2^{n/500}$, the expected running time can be reduced to almost linear in N. In addition, if indeed $d < 2^{n/500}$ the assumption that n, d are fixed can be omitted.

Our technique is useful for exhibiting NC^1 algorithms to several additional algorithmic search-problems. These include the problem of finding an even directed simple cycle in a digraph with some (weak) regularity properties; a problem which is related to the study of sign-solvability of systems of linear equations ([22]). Another example (first considered by Beck in [12]) is that of finding a satisfying assignment of a conjunction of disjunctions, in which no literal appears more times than a small exponent in the length of the shortest disjunction. Several problems of obtaining a covering of the edges of a given graph by a small number of subgraphs satisfying various conditions can also be solved in NC^1 by similar techniques. Some of these problems are suggested by the analysis of radio networks. More details appear in Section 3. We note that although the technique seems rather general there are some known applications of the Local Lemma for which it does not seem to provide efficient solutions for the corresponding algorithmic problems. An example of such an application appears in [10].

The rest of this paper is organized as follows. In Section 2 we describe the basic algorithm, including an essentially full proof of Theorem 1.3, and mention some of its extensions. In Section 3 we describe briefly some of the additional algorithmic search problems that can be solved in NC^1 using similar methods.

2 The basic algorithm

Let H be as in Theorem 1.3. Here is the randomized algorithm for finding a proper 2-coloring of H. In the First Pass we color all the vertices of H, randomly and independently by two colors, where each point is colored either red or blue with equal probability. Call an edge *bad* if at most αn of its points are red or at most αn of its points are blue. The probability of a fixed edge to be bad is clearly at most $2\sum_{i\leq \alpha n}\binom{n}{i}/2^n \leq 2\cdot 2^{(H(\alpha)-1)n}$. Put $p = 2\cdot 2^{(H(\alpha)-1)n}$ and let B denote the set of all bad edges.

Let G be the dependency graph for the problem, i.e., the graph whose vertices are the edges of H in which two are adjacent iff they intersect. Observe that if S is an independent set in G then the probability that $S \subset B$ is at most $p^{|S|}$, since these $|S|$ events are mutually independent. Let us call a set of vertices C of G a $1,2$-tree if C is the set of vertices of a connected subgraph in the square of G. I.e., C is a $1,2$-tree if the $A_i \in C$ are such that drawing an arc between $A_i, A_j \in C$ if their distance in G is either 1 or 2 the resulting graph is connected.

Lemma 2.1 *The probability that every $1,2$-tree in G all of whose vertices belong to B has size at most $d\log(2N)$ is at least $1/2$.*

Proof Call $T \subseteq G$ a $2,3$-tree if the $A_i \in T$ are such that all their mutual distances in G are at least two and so that, drawing an arc between $A_i, A_j \in T$ if their distance in G is either 2 or 3 the resulting graph is connected. We first bound the number of $2,3$-trees of size u in G. Consider the graph on the set of vertices of G in which two vertices are adjacent if their distance in G is either 2 or 3. Every $2,3$-tree on a set T of vertices of G must contain a tree on T in this new graph. The new graph has maximum degree smaller than $D = d^3$. It is well known (see [18]) that an infinite D-regular rooted tree contains precisely $\frac{1}{(D-1)u+1}\binom{Du}{u}$ rooted subtrees of size u, and this easily implies that the number of trees of size u containing one specific given vertex in any graph with maximum degree at most D does not exceed this number, which is smaller than $(eD)^u$.

For any particular $2,3$-tree T of size u we know that $\Pr[T \subseteq B] \leq p^u$. Hence the expected number of $2,3$-trees $T \subseteq B$ of size u is at most $N(eDp)^u$. As $eDp < 0.5$ by (1), if $u = \log(2N)$ this term is at most $1/2$. Thus with probability at least $1/2$ there is no $2,3$-tree of size bigger than $\log(2N)$ all of whose vertices are in B. We actually want to bound the size of any $1,2$-tree C of G all of whose vertices lie in B. A maximal $2,3$-tree T in such a C must have the property that every $A_i \in C$ is a neighbor (in G) of an $A_j \in T$. (This is because otherwise there is an $A_l \in C$ which is not a neighbor of any set in T and yet it is of distance 2 or 3 from some $A_j \in T$. Such an A_l can be added to T, contradicting its maximality). There are less than d neighbors A_i of any given A_j so that $\log(2N) \geq |T| \geq |C|/d$ and so

$$|C| \leq d\log(2N)$$

completing the proof of the lemma. □

Let us call the First Pass *successful* if there is no $1,2$-tree of size greater than $d\log(2N)$ all of whose vertices lie in B. By the last lemma the probability the First Pass is successful is at least $1/2$. In case it is not, we simply repeat the entire procedure. In expected *linear* time the First Pass is successful.

We can now fix the coloring by recoloring, in the Second Pass, the vertices of H that belong to the bad edges. Let us call an edge *dangerous* if it contains at least αn vertices that belong to bad edges. (Thus, in particular, bad edges are also dangerous). Observe that if an edge is not dangerous then it will not become monochromatic after the recoloring. This is because less than αn of its points will change color, and it has at least αn points of each color before the recoloring. Thus we only have to worry about the dangerous edges. However, if we recolor all the vertices in bad edges randomly and independently than we recolor at least αn vertices in each dangerous edge, and hence the probability it becomes monochromatic does not exceed $2^{-\alpha n}$. Since each dangerous edge intersects at most d others it follows from (2) and from Theorem 1.2 (or directly from the Lovász Local Lemma) that there exists a recoloring in which no edge is monochromatic.

The crucial point is that the recoloring of the points in the edges of each maximal $1,2$-tree C of bad edges can be done separately. This is because there is no dangerous edge that intersects edges from two distinct such maximal $1,2$-trees, and hence it suffices to recolor the points in the edges of each such C in a way that only makes sure

that no dangerous edge intersecting an edge in C becomes monochromatic. Since each of the $1,2$-trees C as above has only $O(\log N)$ vertices that have to be recolored, we can find the required recoloring by *exhaustive enumeration!* Examining all possible two-colorings in each such C only takes time $O(2^{O(\log N)}) = N^{O(1)}$ and hence doing it for all the above C-s can be done in polynomial time.

This completes the description of the randomized algorithm with expected polynomial running time. In case $d < 2^{cn}$ for a smaller c we can make another pass similar to the first one in each $1,2$-tree separately, get new $1,2$-trees of size $O(\log logN)$ and complete as before obtaining an expected running time which is nearly linear- $O(N(\log N)^{O(1)})$. We omit the detailed computation.

The randomized algorithm above is trivially parallelizable and can be implemented on a standard $EREW$-PRAM in time $O(\log N)$ using $N^{O(1)}$ parallel processors. (See [17] for the basic definitions of an $EREW$-PRAM and the complexity classes NC and NC^1. Here we extend the definition of NC^1 and consider it as a class containing search problems as well as decision problems.) Moreover, the algorithm can be derandomized maintaining the running time (with some increase in the number of processors), showing that the problem can be solved in NC^1. To see this observe that the recoloring step is deterministic even in the version described above, so the only problem is the derandomization of the First Pass. This can be done by applying the techniques developed in [20] and improved and simplified in [6]. (It is also possible to apply the methods of [19] and [13], but this will not supply NC^1-algorithms, although it would yield algorithms with a smaller number of processors). The basic idea is that for every constant c there is a constant $b = b(c)$ such that for every m there are explicit sample spaces of size at most m^b in which one can embed m uniform $\{0,1\}$ random variables in which every set of $c \log m$ of the variables is nearly independent. The details of these constructions appear in the above mentioned papers. Let us only mention here that one possibility of performing the first pass described above deterministically is the following.

Let $q \geq N^b$ be a prime, where b is a constant depending only on n. Let χ be the quadratic character defined on the elements of the finite field $GF(q)$, i.e., $\chi(x) = 1$ if x is a quadratic residue modulo q and $\chi(x) = -1$ otherwise. Define a family of q two-colorings of the set of vertices of H as follows. For each $i \in GF(q)$, the color of the j-th point in the i-th coloring is blue if $\chi(i-j) = 1$ and is red if $\chi(i-j) = -1$. Using the results in [6] (based on the ideas in [20] (see also [3])), it can be shown that at least one of the q colorings defined above will produce a successful First Pass. All these (deterministically defined) colorings can be checked in parallel, completing the proof of Theorem 1.3. □

It is worth noting that one can give a different variant of the first pass in the basic algorithm, which is also parallelizable, but since it yields an NC^2 and not an NC^1 algorithm we do not describe it here.

The argument in the proof of Theorem 1.3 can be extended to the case where n and d are not fixed, provided d is somewhat smaller as a function of n, say, $d \leq 2^{n/500}$. Moreover, we can also obtain a coloring in which each edge contains *many* points of each of the two colors. By iterating such a coloring procedure one can obtain a coloring with more than 2 colors in which every edge contains points in each color. The detailed description of these algorithms will be given in the full version of the paper. Here we merely state their existence in the following two statements.

Theorem 2.2 *For any $n \geq 2, d$ satisfying $d \leq 2^{n/500}$ there is an NC^1 algorithm whose input is a hypergraph $H = (V,E)$ in which each edge has at least n vertices and each edge intersects at most d other edges, and whose output is a vertex two-coloring of H such that each edge has at least $n/1000$ points of each color. The algorithm runs in time $O(\log N)$ on an EREW-PRAM with $N^{O(1)}$ parallel processors, where N is the number of edges of H.*

We note that the constants 500 and 1000 can be easily improved, and we make no attempt to optimize them here.

Theorem 2.3 *For every fixed k there is an $\epsilon = \epsilon(k) > 0$ such that the following holds. There is an NC^1 algorithm whose input is a hypergraph $H = (V,E)$ in which each edge has at least $n \geq 2$ vertices and each edge intersects at most d others, where $d \leq 2^{\epsilon n}$, and whose output is a k-coloring of the vertices of H so that each color occures in each edge. The algorithm runs in time $O(\log N)$ on an EREW-PRAM with $N^{O(1)}$ parallel processors, where N is the number of edges of H.*

3 Additional algorithmic applications

In this section we describe several additional algorithmic problems for which the technique above yields efficient deterministic parallel algorithms. Some of these can be obtained directly from the results on hypergraph coloring and some require additional ideas based on the same basic approach. We do not know of any efficient (deterministic or randomized, sequential or parallel) algorithm for any of these problems which does not use the methods of [12] or these of the present paper.

3.1 Even cycles in directed graphs

The problem of deciding whether a given directed graph contains a (simple) even (directed) cycle arises in various contexts, including the study of the problem of sign-solvability of a system of linear equations, (see [22]). There is no known polynomial time algorithm for answering this decision problem, but it is also not known to be NP-complete. (It easily follows from the results of [15] that the similar decison problem of deciding whether a given directed graph contains an even cycle through a given edge is NP-complete). There are several results that supply a sufficient condition for a digraph to contain such a cycle. Friedland ([16]) showed that for every $r \geq 7$, any r-regular digraph contains an even cycle. His proof supplies no polynomial time agorithm for finding an even cycle in such a digraph. In [9] it is shown that one can assume much less than exact regularity. In fact, any digraph in which the maximum indegree is not much bigger than the minimum outdegree contains an even cycle. The proof applies the Local Lemma. The results here enable us to convert this proof into an efficient NC^1 algorithm, as shown in the next theorem. (Here and from now on we make no attempt to optimize the constants.)

Theorem 3.1 *There is an NC^1 algorithm whose input is a digraph $D = (V, E)$ in which each outdegree is at least $\delta > 0$ and each indegree is at most Δ, where*

$$\Delta \leq \frac{2^{(\delta+1)/500}}{\delta+1},$$

and whose output is an even directed cycle in D. The algorithm runs in time $O(\log N)$ on an EREW-PRAM with $N^{O(1)}$ parallel processors, where N is the number of edges of D.

Proof We may assume, by deleting edges if necessary, that every outdegree is precisely δ. For each vertex v of D, let $N^+(v)$ denote the set of all vertices u of D such that there is a directed edge from v to u. Let H be the hypergraph whose vertex set $V(H)$ is the set V of all vertices of D, and whose edges are all the sets $v \cup N^+(v)$, where $v \in V$. Every edge of H has precisely $\delta + 1$ vertices. We claim that every edge intersects at most $2^{(\delta+1)/500}$ other edges. To see this observe that for any vertex v of D and for any $w \in N^+(v)$ there are at most $\Delta - 1$ vertices u other then v such that $w \in N^+(u)$. There are at most Δ vertices u such that $v \in N^+(u)$ and there are at most δ additional vertices u with $u \in N^+(v)$. Altogether, there are at most $(\Delta-1)\delta+\Delta+\delta = \Delta(\delta+1)$ vertices u with $(u \cup N^+(u)) \cap (v \cup N^+(v)) \neq \emptyset$, and this number is, by assumption, at most $2^{(\delta+1)/500}$, as claimed.

It thus follows from Theorem 2.2 that we can find a proper two coloring $f : V \mapsto \{0, 1\}$ of H in time $O(\log N)$ using $N^{O(1)}$ parallel processors. For each vertex $v \in V$, let $u = u(v)$ be a vertex so that $f(u) \neq f(v)$ and (v, u) is a directed edge of D. (Such a u exists and can be found quickly for all v since f is not a constant on $v \cup N^+(v)$). Let G denote the subgraph of D consisting of all the (directed) edges $(v, u(v))$. Observe that each directed cycle in G is even, since the values of f on the vertices along this cycle alternate. Moreover, every outdegree in G is precisely 1. It is easy to see that one can find a directed cycle in G in time $O(\log N)$ using $N^{O(1)}$ processors, completing the proof. □

The assertion of the last theorem can be generalized by applying Theorem 2.3 rather then Theorem 2.2 for deriving an algorithmic version of the general result in [9], dealing with cycles of length 0 modulo k. This gives the following result, whose detailed proof (which is analogous to that of the previous theorem) is omitted.

Theorem 3.2 *For every fixed k there is an $\epsilon = \epsilon(k) > 0$ such that the following holds. There is an NC^1 algorithm whose input is a digraph $D = (V, E)$ in which each outdegree is at least $\delta > 0$ and each indegree is at most Δ, where*

$$\Delta \leq 2^{\epsilon\delta},$$

and whose output is a (simple) directed cycle in D whose length is divisible by k. The algorithm runs in time $O(\log N)$ on an EREW-PRAM with $N^{O(1)}$ parallel processors, where N is the number of edges of D.

3.2 Star arboricity

A *star forest* is a forest whose connected components are stars. The *star arboricity* of a graph G, denoted by $st(G)$, is the minimum number of star forests whose union covers all edges of G. For an integer $\Delta \geq 1$ define $st(\Delta) = Max\{st(G)\}$, where the maximum is taken over all simple graphs with maximum degree Δ.

The star arboricity of graphs was introduced by Akiyama and Kano [1] and has been studied in various papers. The asymptotic behaviour of $st(\Delta)$ is determined in [11], improving a previous estimate from [2]. In particular it is known that for every $\gamma > 0$ there is a $\Delta = \Delta_0(\gamma)$ such that for every $\Delta \geq \Delta_0$, $st(\Delta) \leq (\frac{1}{2} + \gamma)\Delta$.

Given a small fixed $\gamma > 0$ and a graph G with maximum degree Δ, where $\Delta > \Delta_0(\gamma)$ is fixed, can we find efficiently a family of at most $(\frac{1}{2}+\gamma)\Delta$ star-forests that cover all edges of G? As shown in the next subsection this question is naturally suggested by the study of a certain communication model. The next result, shows that the desired star-forests can be found efficiently and in parallel. The proof here requires some additional ideas besides the basic approach presented in Section 2 and will appear in the full version of the paper.

Theorem 3.3 *For every fixed $\gamma > 0$ there is a $\Delta_0 = \Delta_0(\gamma)$ such that the following holds. There exists an algorithm whose input is a (simple) graph G with a fixed maximum degree $\Delta > \Delta_0$, and whose output is a family of at most $(\frac{1}{2} + \gamma)\Delta$ star-forests whose union covers all edges of G. If N is the number of vertices of the graph, then the algorithm runs on an EREW-PRAM in time $O(\log N)$ using $N^{O(1)}$ processors.*

3.3 Radio networks

The study of star arboricity is naturally suggested by the analysis of certain communication networks. A *radio network* is a synchronous network of processors that communicate by transmitting messages to their neighbors. A processor P can receive at most one message in one step. Let us mention here two possible models.

Type I : P receives a message from its neighbor Q in a given step if P is silent, Q transmits and P chooses to receive from Q in this step.

Type II: P receives a message from its neighbor Q if P is silent, and Q is the *only* neighbor of P that transmits in this step.

Suppose, now, that the model is the Type I model and the network is represented by an undirected graph $G = (V, E)$ whose vertices are the processors and two are adjacent if they can transmit to each other. Suppose, further, that we need to transmit once along every edge (in one of the two possible directions), say, in order to check that there is indeed a connection between each adjacent pair. It is easy to see that the minimum number of steps in which we can finish all the required transmissions is precisely $st(G)$, since the set of edges corresponding to the transmissions performed in a single step forms a star forest. Theorem 3.3 thus supplies an upper bound for the minimum number of required steps. Observe that if G is a Δ-regular graph on N vertices then clearly at least $\Delta/2$ steps are necessary, simply because each star-forest has at most $N - 1$ edges. Therefore, the $(\frac{1}{2} + \gamma)\Delta$ estimate is almost tight in many cases, and the required communication pattern can be found in this case efficiently in parallel. We note that obtaining a covering with Δ star-forests is trivial, so the saving here is only by a factor of 2 (which could still be significant if the communication time is our main concern). As shown below, in the analogous problem corresponding to type II networks we improve the trivial bound by more than a constant factor.

The model of type II networks is much more popular, and has been considered in many papers (see, e.g., [4] and its many references). A basic parameter of a network represented by a directed graph D in this model is its *hitting number* $h(D)$ defined as follows. Given a digraph $D = (V, E)$ and a subset A of V, we say that A *hits* a directed edge $(x, y) \in E$ if $x \in A$, $y \notin A$ and the only vertex in $N^-(y) \cap A$ is x, where $N^-(y)$ is the set of all vertices v such that there is a directed edge from v to y. A family of subsets of V *hits* D if for every directed edge e of D there is a set in the family that hits e. Let $h(D)$ denote the minimum cardinality of a family that hits D. One can easily see that this is precisely the minimum number of steps in which it

is possible to complete a transmission along every edge of D. This parameter is of central interest in the study of the task referred to as a Single Round Simulation in [5].

For simplicity let us restrict our attention to symmetric graphs, i.e., (x, y) is an edge of D iff (y, x) is an edge. In [4], [5] it is shown that there exist two positive constants c_1 and c_2 such that
(i) For every (symmetric) digraph D with maximum degree Δ, $h(G) \leq c_1 \Delta \log \Delta$.
(ii) For every Δ there is a symmetric digraph with maximum degree Δ such that $h(D) > c_2 \Delta \log \Delta$.

The proof of the upper bound stated in (i) applies the Local Lemma and improves a trivial bound of Δ^2. Given a symmetric digraph D with maximum degree Δ, can we find efficiently and in parallel a family of cardinality $O(\Delta \log \Delta)$ that hits it? This problem is suggested naturally by the discussion in [4], [5]. The technique described here can be used to answer this as follows.

Theorem 3.4 *There exists a constant $c > 0$ such that for every fixed Δ there exists an algorithm whose input is a symmetric digraph $D = (V, E)$ with maximum degree Δ, and whose output is a family F of subsets of V that hits D, where $|F| \leq c\Delta \log \Delta$. Such a family can be computed in time $O(\log N)$ on an EREW-PRAM with $N^{O(1)}$ processors, where N is the number of vertices of D.*

The detailed proof will appear in the full version of the paper.

3.4 Satisfying assignments for restricted CNF formulae

The proof of the following Theorem is almost identical to the proof of the results dealing with hypergraph 2-coloring, and is omitted. An efficient sequential algorithm to this problem was first given in [12].

Theorem 3.5 *For any $n \geq 2, d$ satisfying $d \leq 2^{n/100}$ there is an NC^1 algorithm whose input is a Conjunction of Disjunctions in which each Disjunction contains at least n literals and each literal appears in at most d disjunctions and whose output is a satisfying assignment for this formula. The algorithm runs in time $O(\log N)$ on an EREW-PRAM with $N^{O(1)}$ parallel processors, where N is the length of the formula.*

3.5 Acyclic edge-coloring

We have seen several algorithmic problems that can be solved efficiently in parallel using the technique described in the first sections. In his original paper, Beck [12] describes several other results obtained by the Local Lemma, for which he obtains efficient (sequential) algorithms. There are several additional examples, including the main results in [7] and in [8] that can be converted into efficient algorithms by similar methods. It is not yet clear if this technique suffices for converting all the known applications of the Local Lemma into efficient algorithms. An example which looks difficult is finding an algorithmic version for the following theorem proved in [10]:

Theorem 3.6 *The edges of any graph with maximum degree Δ can be colored by 64Δ colors so that there are no two adjacent edges having the same color and there is no 2-colored cycle.*

Acknowledgement I would like to thank József Beck for sending me an early draft of his paper and Joel Spencer for helpful comments.

References

[1] J. Akiyama and M. Kano, *Path factors of a graph*, in: Graph Theory and its Applications (Wiley and Sons, New York, 1984).

[2] I. Algor and N. Alon, *The star arboricity of graphs*, Discrete Mathematics 75 (1989), 11-22.

[3] N. Alon, L. Babai and A. Itai, *A fast and simple randomized parallel algorithm for the maximal independent set problem*, J. of Algorithms 7 (1986), 567-583.

[4] N. Alon, A. Bar-Noy, N. Linial and D. Peleg, *On the complexity of radio communication*, Proc. 21^{th} ACM STOC, Seattle, Washington, ACM Press (1989), 274-285.

[5] N. Alon, A. Bar-Noy, N. Linial and D. Peleg, *Single round simulation on radio networks*, to appear in J. of Algorithms.

[6] N. Alon, O. Goldreich, J. Hastad and R. Peralta , Simple constructions of almost k-wise independent random variables, Proc. 31^{st} FOCS, St. Louis, Missouri, IEEE (1990), 544-553.

[7] N. Alon, *The linear arboricity of graphs*, Israel J. Math. 62(1988), 311-325.

[8] N. Alon, *The strong chromatic number of a graph*, to appear in Random Structures and Algorithms.

[9] N. Alon and N. Linial, *Cycles of length* 0 *modulo k in directed graphs*, J. Combinatorial Theory, Ser. B 47(1989), 114-119.

[10] N. Alon, C. McDiarmid and B. Reed, *Acyclic coloring of graphs*, to appear in Random Structures and Algorithms.

[11] N. Alon, C. McDiarmid and B. Reed, *Star Arboricity*, to appear in Combinatorica.

[12] J. Beck, *An algorithmic approach to the Lovász Local Lemma I*, to appear in Random Structures and Algorithms.

[13] B. Berger and J. Rompel, *Simulating* $log^c n$*-wise independence in NC*, Proc. 30^{th} IEEE FOCS (1989), 2-7.

[14] P. Erdös and L. Lovász, *Problems and results on* 3*-chromatic hypergraphs and some related questions*, in:"Infinite and Finite Sets" (A. Hajnal et. al. eds), Colloq. Math. Soc. J. Bolyai 11, North Holland, Amsterdam, 1975, pp. 609-627.

[15] S. Fortune, J.E.Hopcroft and J.Wyllie, *The directed subgraph homeomorphism problem*, Theoret. Comp. Sci. 10(1980),111-120.

[16] S. Friedland, *Any* 7*-regular digraph contains an even cycle*, J. Combinatorial Theory Ser. B, 46 (1989), 249-252.

[17] R. M. Karp and V. Ramachandran, *A survey of parallel algorithms for shared memory machines*, in: Handbook of Theoretical Computer Science (J. Van Leeuwen Ed.), Chapter 17, MIT Press (1990), 871-941.

[18] D. Knuth, *The Art of Computer Programming*, Vol. I, Addison Wesley, London, 1969, p. 396 (Exercise 11).

[19] R. Motwani, J. Naor and M. Naor, *The probabilistic method yields deterministic parallel algorithms*, Proc. 30^{th} IEEE FOCS (1989), 8-13.

[20] J. Naor and M. Naor, *Small-bias probability spaces: efficient constructions and applications*, Proc. 22^{nd} annual ACM STOC, ACM Press (1990), 213-223.

[21] J. Spencer, *Ten Lectures on the Probabilistic Method*, SIAM, Philadelphia, 1987.

[22] C. Thomassen, *Sign-nonsingular matrices and even cycles in directed graphs*, Linear Algebra Appl. 41 (1986), 27-42.