

A Generalization of Spira’s Theorem and Circuits with Small Segregators or Separators

Anna Gál* and Jing-Tang Jang**

Dept. of Computer Science, University of Texas at Austin,
Austin, TX 78712-1188, USA
{panni, keith}@cs.utexas.edu

Abstract. Spira [28] showed that any Boolean formula of size s can be simulated in depth $O(\log s)$. We generalize Spira’s theorem and show that any Boolean *circuit* of size s with segregators of size $f(s)$ can be simulated in depth $O(f(s) \log s)$. If the segregator size is at least s^ε for some constant $\varepsilon > 0$, then we can obtain a simulation of depth $O(f(s))$. This improves and generalizes a simulation of polynomial-size Boolean circuits of constant treewidth k in depth $O(k^2 \log n)$ by Jansen and Sarma [17]. Since the existence of small balanced separators in a directed acyclic graph implies that the graph also has small segregators, our results also apply to circuits with small separators. Our results imply that the class of languages computed by non-uniform families of polynomial-size circuits that have constant size segregators equals non-uniform NC^1 .

Considering space bounded Turing machines to generate the circuits, for $f(s) \log^2 s$ -space uniform families of Boolean circuits our small-depth simulations are also $f(s) \log^2 s$ -space uniform. As a corollary, we show that the Boolean Circuit Value problem for circuits with constant size segregators (or separators) is in deterministic $SPACE(\log^2 n)$. Our results also imply that the Planar Circuit Value problem, which is known to be P -Complete [16], can be solved in deterministic $SPACE(\sqrt{n} \log n)$.

Key words: Boolean circuits, circuit size, circuit depth, Spira’s theorem, Turing machines, space complexity

1 Introduction

Spira [28] proved the following theorem.

Theorem A [28] *Let F be any Boolean formula of size s . Then F can be simulated by an equivalent formula of depth $O(\log s)$.*

There are several results improving or extending Spira’s theorem. Bonet and Buss [3] improved the constants in the depth bounds and the size of the simulation for Boolean formulas, Wegener [30] proved the statement for monotone

* Supported in part by NSF Grant CCF-1018060

** Supported in part by MCD fellowship from Dept. of Computer Science, University of Texas at Austin, and NSF Grant CCF-1018060

Boolean formulas, and Brent [5], Bshouty et. al. [6] extended it for arithmetic formulas. All these results study formulas, i.e. tree-like circuits with fan-out 1.

Valiant, Skyum, Berkowitz and Rackoff [29] showed that arithmetic circuits of size s and degree d can be simulated in size $O((sd)^{O(1)})$ and $O(\log s \log d)$ depth. This implies that polynomial-size and polynomial-degree arithmetic circuits can be simulated in NC^2 . However, very little is known for size vs. depth for general Boolean circuits. The strongest results so far for general Boolean circuits by Paterson and Valiant [23], and Dymond and Tompa [12] give a simulation of arbitrary Boolean circuits of size s in depth $O(s/\log s)$.

In this paper, we generalize Spira's technique to circuits with small segregators or small separators. Informally, the separator of a graph is a subset of the nodes whose removal yields two subgraphs of comparable sizes. (See the following section for a formal definition.) Graphs with small separators include trees, planar graphs [20], graphs with bounded genus [15], graphs with excluded minors [1], as well as graphs with bounded treewidth [25].

Segregators are a relaxed version of separators of directed acyclic graphs. Paul et al. [24], and Santhanam [26] used segregators to study the computation graph of Turing machines. Directed acyclic graphs with small separators also have small segregators, but the reverse may not necessarily hold.

Jansen and Sarma [17] studied the question of simulating Boolean circuits with bounded treewidth by small-depth circuits. They showed that polynomial-size circuits with constant treewidth k can be simulated in depth $O(k^2 \log n)$, and thus the class of languages with non-uniform polynomial-size bounded treewidth circuits equals non-uniform NC^1 .

We extend this result to arbitrary circuits with small segregators and show that any Boolean circuit of size s with segregators (or separators) of size $f(s)$ can be simulated in depth $O(f(s) \log s)$. For circuits with segregators of size k , thus also for graphs with treewidth k , this gives a simulation in depth $k \log s$, improving the bound in [17]. If the segregator size is at least s^ϵ for some constant $\epsilon > 0$, then we can obtain a simulation of depth $O(f(s))$. Our results imply that the class of languages computed by non-uniform families of polynomial-size circuits that have constant-size segregators equals non-uniform NC^1 .

In [14] we observed that the two-person pebble game of Dymond and Tompa can be used to simulate circuits with small separator size in small depth, giving essentially the same dependence of the depth on the separator size as in the current paper. The approach in [14] based on the two person pebble game can also be extended to graphs with small segregators. However, the simulation based on the two person pebble game is non-uniform, and it seems that the resulting circuits cannot be produced efficiently using this approach. Jansen and Sarma's [17] simulation of bounded treewidth circuits is also non-uniform.

For circuits with constant-size segregators or separators, the simulating circuits we obtain in this paper can be generated in space $O(\log^2 s)$. We also note that our simulation works for any circuit, and if the circuit has a segregator of size $f(s)$, we obtain a simulating circuit of depth at most $O(f(s) \log s)$, the value $f(s)$ does not have to be provided in advance. In contrast, the simulation

in [17] assumes that the treewidth k is known in advance, and a tree decomposition is available along with the description of the circuit to be simulated. It would be desirable to generate the simulating circuits even more efficiently with respect to space or circuit depth, especially in the case of polynomial-size circuits with constant-size segregators or separators, since in that case, as in the case of formulas in Spira’s theorem, the resulting circuits are NC^1 circuits. Note however, that even in the case of formulas (tree-like circuits) Spira’s theorem is non-uniform. It is not known if the restructuring procedure for formulas in Spira’s theorem producing the simulating $O(\log s)$ depth circuits can be directly implemented in less than $O(\log^2 s)$ space, or less than $O(\log^2 s)$ depth [8, 9].

The question of finding a uniform version of Spira’s theorem has direct relevance for the complexity of the Boolean Formula Value problem. While a logspace uniform version of Spira’s restructuring algorithm is still not known, it was proved (by a different approach), that for Boolean formulas presented as parenthesized expressions the Boolean Formula Value problem is in $SPACE(\log n)$ [21], and in $DLOGTIME$ -uniform NC^1 [8, 9].

Our generalization of Spira’s theorem allows us to bound the space complexity of the Circuit Value Problem (CVP) for circuits with small separators and segregators. Ladner [18] showed that the Circuit Value Problem is P-complete. The space complexity of the CVP is not known to be $o(n/\log n)$ for general Boolean circuits. It is a straightforward consequence of Borodin’s theorem [4] (see Theorem C) that the CVP for logspace uniform depth d circuits is in $SPACE(d)$ for $d \geq \log n$. It is also easy to see that the CVP for small-width circuits can be solved in small space. Barrington, Lu, Miltersen and Skyum [2] showed that the Monotone Planar Circuit Value Problem is in $LOGDCFL$, and thus in $SPACE(\log^2 n)$. See [10, 19] for recent results on variants of the Monotone Planar Circuit Value Problem. As far as we know, the only other variant that was previously shown to be computable in small (polylog) space is the Boolean Formula Value Problem, that is the Circuit Value Problem for tree-like circuits [8, 9, 21]. We show that the Boolean Circuit Value Problem for circuits with constant-size segregators (or separators) is in deterministic $SPACE(\log^2 n)$. Our results also imply that the Planar Circuit Value problem, which is known to be P-Complete [16], can be solved in deterministic $SPACE(\sqrt{n} \log n)$.

2 Preliminaries

2.1 Space Bounded Turing Machines

For the space complexity of Turing machines, we follow the convention of considering Turing machines with a separate *read-only* input tape, and additional work tapes. If the machine has to produce an output string (instead of just accepting or rejecting its input), then we also assume a separate *write-only* output tape. The space used by a Turing machine on a given input is defined as the number of work tape cells visited during the computation over all work tapes. The input tape and the output tape do not contribute to the space bound of the computation. This allows us to consider computations with sublinear space.

$SPACE(s(n))$ denotes the class of languages decidable by deterministic Turing machines with a separate read-only input tape and a separate write-only output tape using $O(s(n))$ space on the work tapes.

In the following, whenever we talk about space bounds of Turing Machines, it is assumed that the input tape is read-only, the output tape is write-only and the space bound refers to the space used on the work tapes. See Papadimitriou [22] for more details on space bounded Turing machines.

2.2 The Circuit Model

A Boolean circuit is a labeled directed acyclic graph (DAG), where every node is labeled by either a variable from $\{x_1, \dots, x_n\}$, or an operation from $\{\wedge, \vee, \neg\}$. The inputs of a Boolean circuit are the nodes with in-degree (fan-in) zero, and the outputs of a Boolean circuit are the nodes with out-degree (fan-out) zero. We refer to the nodes (including the inputs) as *gates*. A formula (or tree-like circuit) is a circuit whose fan-out is one for every gate except the output. The size of a Boolean circuit is the number of its gates. We will consider Boolean circuits with gates of fan-in at most 2 from the basis $\{\wedge, \vee, \neg\}$. The *depth of a gate g* is the length of the longest path from any input to g . The *depth of a circuit C* is the depth of the output gate. See [31] for more on Boolean circuits.

Definition 1. *A family of Boolean circuits $\{C_n\}$ is called $h(n)$ -space uniform, if there exists a deterministic Turing machine M that on input 1^n , outputs the standard description of C_n using space $O(h(n))$ for all n . In particular, $\{C_n\}$ is logspace uniform if $h(n) = \log n$.*

2.3 Separators and Segregators

Informally, a node separator of a graph G is a set of nodes whose removal yields two disjoint subgraphs of G . In this paper we only consider *balanced* separators, that yield subgraphs that are comparable in size. In the next definition each of the two subDAGs could consist of several weakly connected components.

Definition 2. *A separator of size k of a DAG $G = (V, E)$ is a set of k nodes $S \subseteq V$ such that $G \setminus S$ is not weakly connected (i.e. the underlying undirected graph is not connected); and the removal of S partitions $G \setminus S$ into two subDAGs, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, such that $|V_i| \leq \frac{2}{3}|V|$ for $i = 1, 2$, and there are no edges either from G_1 to G_2 , or from G_2 to G_1 in $G \setminus S$.*

Segregators are a relaxation of separators in directed acyclic graphs [24, 26].

Definition 3. *A segregator of size k of a DAG $G = (V, E)$ is a set of k nodes $S \subseteq V$ such that every node in $G \setminus S$ has at most $\frac{2}{3}|V|$ predecessors in $G \setminus S$.*

The following lemma follows directly from the definitions.

Lemma 1. *Any DAG with a separator of size k has a segregator of size k .*

Notice that the reverse is not true in general, since a node in a DAG may have much smaller number of predecessors than the size of the component that contains the node in the underlying undirected graph.

3 Boolean Circuits with Small Segregators or Separators

Definition 4. We say that a Boolean circuit C has separators of size $f()$ if the underlying DAG of every subcircuit of C with s gates has a separator of size at most $f(s)$.

We say that a Boolean circuit C has segregators of size $f()$ if the underlying DAG of every subcircuit of C with s gates has a segregator of size at most $f(s)$.

The above definition is reasonable, since we typically consider classes of circuits based on properties of their underlying DAGs that are closed with respect to subDAGs, for example planar circuits, circuits with small treewidth, etc.

We talk about constant-size separators (resp. segregators), if the size of the separator (resp. segregator) is bounded by a fixed constant that does not depend on the size of the circuit.

By Lemma 1, if the circuit has separators of size $f()$, then it must also have segregators of size $f()$. Therefore in the following we will focus on circuits with small segregators. We prove the following generalization of Spira's theorem.

Theorem 2. Any Boolean circuit of size s with segregators of size $f()$ can be simulated in depth $O(f(s))$ if $f(s) = \Omega(s^\varepsilon)$ for some constant $\varepsilon > 0$, and in depth $O(f(s) \log s)$ otherwise.

Proof. The construction is defined recursively. Let $U = \{u_1, \dots, u_p\}$ be the segregator of C with size $p \leq f(s)$. Let C_1, \dots, C_p be the subcircuits of C corresponding to the nodes of the segregator, that is the node u_j is the output of the subcircuit C_j , for $j = 1, \dots, p$. Let g_j be the Boolean function computed by C_j . Let v be the output node of the circuit C , and let \hat{C} be the circuit with output node v , obtained from C by replacing the nodes in U by new variables y_1, \dots, y_p . Thus, if the original circuit C has n variables, then \hat{C} may have up to $p + n$ variables. It is possible that \hat{C} has less than $p + n$ variables, if some of the original inputs get disconnected from the output v after removing the nodes of the segregator from the circuit.

We enumerate all Boolean vectors $c \in \{0, 1\}^p$. Let $c_i = \langle c_{i,1}, c_{i,2}, \dots, c_{i,p} \rangle$ be the i th Boolean vector of length p , for $i = 1, \dots, 2^p$, according to some fixed ordering. Let \hat{C}_i be the circuit obtained from \hat{C} by fixing the values of the variables y_1, \dots, y_p to the bits $c_{i,1}, \dots, c_{i,p}$, respectively. Let h_i be the Boolean function computed by the circuit \hat{C}_i .

Then, the Boolean function computed by the circuit C can be represented using the following expression:

$$\bigvee_{i=1}^{2^p} \left(h_i \wedge \bigwedge_{j=1}^p ((g_j \wedge c_{i,j}) \vee (\neg g_j \wedge \neg c_{i,j})) \right) \quad (1)$$

Next we will represent the functions h_i for $i = 1, \dots, 2^p$ and g_j for $j = 1, \dots, p$. We could proceed with a straightforward recursion, if we could claim

that each subcircuit C_1, \dots, C_p and each circuit \hat{C}_i for $i = 1, \dots, 2^p$ has size at most $2s/3$. In fact, we do know that every subDAG of the underlying DAG of C with the nodes of U removed has size at most $2s/3$. However, the output node of the subcircuit C_j is u_j , and u_j is a member of the segregator U . Note that the underlying DAGs of the circuits \hat{C}_i are identical (they only differ from each other in the substituted constants), and their output node v is the output node of the “original” circuit C . The node v may or may not participate in the segregator. If the node v participates in the segregator, then the functions h_i are constants and the recursion stops.

We can compute the function g_j (computed at gate u_j) by an additional gate if we compute the functions computed at the two children of the gate u_j . If none of the children participates in the segregator, then we know that their subcircuits must have size at most $2s/3$. However, it is possible that children of segregator nodes are also included in the segregator. Let S_j be the set of nodes in the segregator, that are predecessors of u_j , such that there is a path from each of them to u_j that consists only of segregator nodes. We also include u_j in S_j . That is, S_j forms a subcircuit with output u_j that consists of segregator nodes. Let B_j be the “boundary” of S_j formed by nodes that are not in the segregator, that is, B_j contains the children of the nodes in S_j that are not included in the segregator. Then we can compute the function g_j from the functions computed at the nodes in B_j (these can be computed by subcircuits of size at most $2s/3$) with an additional set of gates corresponding to the segregator nodes in S_j . Since $|S_j| \leq p$, this takes additional depth at most p .

To summarize, we can compute the functions h_i and g_j , by first computing in parallel the functions corresponding to all subcircuits after removing the nodes of the segregator. We know that each such subcircuit has size at most $2s/3$, and we can use our construction recursively on these smaller size circuits. Then we finish computing every function h_i and g_j we need, by adding the gates corresponding to the nodes participating in the segregator. This will take at most an additional $p \leq f(s)$ depth. Then we can compute the function computed by C by expression (1). This takes at most an additional $p + \lceil \log(p+1) \rceil + 3 = O(f(s))$ depth. Thus, in each iteration, we increase the depth by at most $O(f(s))$. Since the size is reduced by a constant factor in each iteration, we are done after $O(\log s)$ steps. More precisely, the depth of the final circuit is $O\left(\sum_{i=0}^{\lceil \log_{3/2} s \rceil} f\left((2/3)^i s\right)\right)$. Thus the depth of the final circuit is $O(f(s))$ if $f(s) = s^\epsilon$ for some constant $\epsilon > 0$, or $O(f(s) \log s)$ otherwise. \square

Theorem 3. *The class of languages decided by non-uniform families of polynomial-size circuits with constant-size segregators equals non-uniform NC^1 .*

Proof. Immediately follows from Theorem 2. \square

Robertson and Seymour [25] showed that if a graph has treewidth k , then the graph also has separator size $O(k)$. Together with Lemma 1 and Theorem 3, a polynomial-size circuit with treewidth k can be simulated in depth $O(k \log n)$. This improves a result in [17], which showed that Boolean circuits of size $n^{O(1)}$

and treewidth k can be simulated in non-uniform depth $O(k^2 \log n)$. We refer interested readers to [11] and [13] for more background on treewidth.

4 Finding minimum size segregators in small space

4.1 Segregators of directed acyclic graphs

In this section, we give a space-efficient algorithm to find a minimum size segregator in arbitrary directed acyclic graphs.

We will use the following space-efficient algorithm for reachability in directed graphs by Savitch [27], to count the number of predecessors of a given node.

Theorem B [27] *Given a directed graph G on s nodes and two nodes $u, v \in G$, there exists a deterministic Turing machine that decides if there is a path from u to v in G using space $O(\log^2 s)$.*

Lemma 4. *Let G be a DAG with s nodes. There exists a deterministic Turing machine M such that, on input G , if G has a segregator of size $f(s)$, then M outputs a segregator of G of size at most $f(s)$ using space $O(f(s) \log s + \log^2 s)$.*

Proof. We first define a Turing machine M_1 that takes G and a node $v \in G$ as input, and computes the number of predecessors of v in G , i.e. the number of nodes u such that there exists a directed path from u to v in G . In the beginning M_1 initializes a counter to 1. Then M_1 uses Theorem B to check, one-by-one, for each node $u \in G \setminus \{v\}$ if there is a directed path from u to v in G . For each node $u \in G \setminus \{v\}$ such that v is reachable from u , the counter is incremented. The space used to check the reachability of v from u is reused when checking for reachability from the next node in $G \setminus \{v\}$. Thus M_1 uses $O(\log^2 s)$ space and computes the size of the subDAG with v as the root.

We now define M in Lemma 4 as follows. First M enumerates integers k such that $1 \leq k \leq s$ in increasing order. For a fixed k , M enumerates subsets W of size k of the nodes in G in lexicographic order. For a given W , for every node $u \in G \setminus W$, let $G(u)$ denote the set of predecessors of u in $G \setminus W$. That is, $G(u)$ is the subDAG in $G \setminus W$ with u as its root. M uses M_1 to compute $|G(u)|$. If there exists one node $u \in G \setminus W$ such that $|G(u)| > \frac{2}{3}s$, then M continues to the next W , or the next k if every W of the current size has been already checked. Also, every time before continuing to the next W or the next k , M clears unnecessary information from the work tape.

We now argue that M will find a segregator of the smallest size. Observe that the set of nodes of G is a segregator of size s , so M is guaranteed to find a segregator. Since we try every k in increasing order, and we check for every subset W of size k whether or not it is a segregator, it is guaranteed that we will find a segregator of the smallest possible size in G .

We now argue that M only uses $O(f(s) \log s + \log^2 s)$ space. The description of G can be read using a counter of size $O(\log s)$. The enumeration and the storing of W both take $O(k \log s) = O(f(s) \log s)$ space. The computation of $|G(u)|$

takes $O(\log^2 s)$ space since M_1 uses $O(\log^2)$ space. Thus the space complexity to find a segregator of smallest size is $O(f(s) \log s + \log^2 s)$. \square

Note that in the proof for Lemma 4, the input of M consists of only the description of the graph. M does not know the value of $f(s)$ in advance. Also, by Lemma 1, for graphs with separators of size k , the algorithm in Lemma 4 will also find a segregator of size at most k .

4.2 Segregators of uniform circuits

Intuitively, Lemma 4 seems to apply directly to circuits since circuits are also DAGs. However, the input of the Turing machine that has to generate the circuit C_n for a uniform family of circuits, is the unary representation of n (1^n), so the graph of the circuit C_n is not available directly. Since we want to generate the segregator using small space, we cannot store the description of C_n on the work tapes. As it is standard in such situations, we will generate the description of C_n as needed for the machine in the proof of Lemma 4, but never store the complete description. We then have the following lemma.

Lemma 5. *Let \mathcal{C} be a $h(n)$ -space uniform family of circuits. Let $C_n \in \mathcal{C}$ be the Boolean circuit in the family with n inputs, and assume that C_n has size $s = s(n)$ and a segregator of size $f(s)$. Then there exists a deterministic Turing machine \tilde{M} that on input 1^n , outputs a segregator of C_n of size at most $f(s)$ using space $O(h(n) + f(s) \log s + \log^2 s)$.*

As in the case for directed graphs, for circuits with separators of size $f(s)$, the algorithm in Lemma 5 will also find a segregator of size at most $f(s)$.

5 Generating the simulating circuits in small space

Let v be any node and Z be any set of nodes in the underlying graph of a circuit C_n . We denote by $C_{v,Z}$ the circuit obtained from the subcircuit C_v of C_n with output v by replacing every node in Z that participates in C_v by a new input variable.

Lemma 6. *Let \mathcal{C} be a $h(n)$ -space uniform family of circuits. Let $C_n \in \mathcal{C}$ be the circuit with n inputs in the family, and assume that C_n has size $s = s(n)$. Let v be any node and Z be any set of nodes in the underlying graph of C_n . Then there exists a Turing machine M_2 such that on input 1^n , v and Z , M_2 outputs the standard description of the circuit $C_{v,Z}$. Furthermore, M_2 runs in space $O(h(n) + \log^2 s)$.*

Note that if $Z = \emptyset$, or if Z does not contain any predecessors of v then $C_{v,Z}$ is simply the subcircuit C_v . Similarly to the circuit \hat{C} in the proof of Theorem 2, if the size of Z is r , and C_v depends on n' input variables, then $C_{v,Z}$ may have up to $n' + r$ variables. If $v \in Z$, then $C_{v,Z}$ is simply a new variable. The proof of this lemma is standard, and we leave it for the full version.

Lemma 7. *Let \mathcal{C} be a $h(n)$ -space uniform family of circuits. Let $C_n \in \mathcal{C}$ be the circuit with n inputs in the family, and assume that C_n has size $s = s(n)$. Let v be any node and Z be any set of nodes in the underlying graph of C_n . Also assume that C_n has segregators of size $f(\cdot)$. Then there exists a Turing machine M_3 such that on input 1^n , v and Z , M_3 outputs a minimum size segregator of $C_{v,Z}$ using space $O(h(n) + f(s) \log s + \log^2 s)$.*

Proof. Let M_2 be the Turing machine in Lemma 6 that generates the description of $C_{v,Z}$ in space $O(h(n) + \log^2 s)$. Let M be the Turing machine in the statement of Lemma 4, that takes a directed graph G as input, and outputs a minimum size segregator of G . The machine M_3 will simulate M on the underlying directed graph of $C_{v,Z}$. However, as before, the full description of the graph will never be stored. Instead, whenever M_3 needs some information about the graph, it lets M_2 run, (without recording its output), until the required information is generated. The size of the subcircuit $C_{v,Z}$ is $s' \leq s$. Since C_n has segregators of size $f(\cdot)$, we know that $C_{v,Z}$ has a segregator of size $f(s')$. Recall that M always finds a minimum size segregator, thus it will find a segregator of size $f(s') \leq f(s)$. Since M runs in space $O(f(s) \log s + \log^2 s)$, the total space used will be $O(h(n) + f(s) \log s + \log^2 s)$. \square

Now we are ready to prove a uniform version of Theorem 2.

Theorem 8. *Let \mathcal{C} be an $h(n)$ -space uniform family of Boolean circuits. Let $C_n \in \mathcal{C}$ be the Boolean circuit on n inputs with size $s = s(n)$. Suppose that C_n has segregators of size $f(\cdot)$. Let $g(s) = f(s)$ if $f(s) = \Omega(s^c)$ for some constant $c > 0$ and $f(s) \log s$ otherwise. Then \mathcal{C} can be simulated by a $O(h(n) + g(s) \log s)$ -space uniform family of Boolean circuits of depth $O(g(s))$.*

Proof. We show that the construction in the proof of Theorem 2 can be generated by a machine M^* within the appropriate space bounds. M^* on input 1^n will output the description of the depth $O(g(s))$ circuit simulating the circuit $C_n \in \mathcal{C}$.

M^* generates the simulating circuit essentially as described in the proof of Theorem 2. In each step of the recursion, M^* has to do the following:

1. Find a segregator S of the current subcircuit, and store the list of nodes of S in workspace.
2. Find and store the list of nodes that participate in $B = \cup_{j=1}^{|S|} B_j$. Note that a given node may belong to B_j for more than one j , but $|\cup_{j=1}^{|S|} B_j| \leq 2|S|$, since B_j contains only children of segregator nodes. Thus, if $|S| = p$, it takes $O(p \log s)$ space to store the list of nodes in B . We can generate this list using \hat{M}_1 , where \hat{M}_1 is the Turing machine that on input 1^n generates the description of C_n using space $O(h(n))$. We will run \hat{M}_1 several times, reusing space, and never store the full description of the circuit, as discussed before. For finding the set B_j , we have to find the set S_j and store it until we are finished generating B_j . For each j this takes $O(p \log s)$ workspace. We reuse this space when we move on to the next j . For each node of B_j that we find, we check if we have already added it to the list, so the full list B takes at most $O(p \log s)$ workspace to store.

3. Output the description of the part of the circuit that corresponds to the current subcircuit. This is based on the expression (1), and the sets B_j and S_j . We produce the description of the part of the circuit to compute g_j , while we have B_j and S_j stored in memory. We reuse space when we move on to the next j . Recall that the output is not part of the space bound. (We do keep S and the full list B until the end of processing the subcircuit, and maybe longer as we see below.)

The recursion will continue to process the subcircuits \hat{C}_i (functions h_i) defined in the proof of Theorem 2, and the subcircuits of the nodes in B . Recall that each of these subcircuits has size at most $2/3$ of the last subcircuit. The recursion stops when a subcircuit is either constant or an input variable. We need a counter of size p to enumerate the Boolean vectors substituted, and to enumerate the functions h_i , for $i = 1, \dots, 2^p$.

We reuse space as we proceed to the next recursive step. However, to be able to proceed with the recursion, we need to retain some information about the segregators S , the sets B and list of values substituted for segregator nodes from previous recursive steps to be able to generate and process the current subcircuits. We process the subcircuits similarly to a depth first search in the recursion tree, starting with the subcircuits corresponding to the set B and leaving the subcircuit for the functions h_i for last. Recall that there is only one subcircuit to consider for the functions h_i , they just differ in the values of constants substituted.

We keep S , B and list of values substituted for nodes in S from previous steps along the current path in the recursion tree. Since there are $\log s$ stages of the recursion, at any point we keep at most $\log s$ segregators with their corresponding set B and list of values. This takes $O(\sum_{i=1}^{\log s} f(s/2^i) \log s) = O(g(s) \log s)$ space.

At the first iteration, we simply use the machine \hat{M} from Lemma 5 to find a segregator. Now we describe how to find a segregator of the current subcircuit during the recursion. To find a segregator for the subcircuits with outputs in the sets B described above, we use M_3 with input 1^n , u where u is the output of the subcircuit, and $Z = \emptyset$. (For processing the subcircuits corresponding to nodes in the sets B we do not need to worry about the segregators that we stored from previous levels of the recursion.) For the subcircuits \hat{C}_i (functions h_i) we use M_3 with input 1^n , v , where v is the output node of the subcircuits \hat{C}_i (recall that they have the same output node, they only differ in the constants substituted), and Z where Z is the union of all the segregators currently stored.

In each step of the recursion, M_3 finds the current segregator in at most $h(n) + O(\log^2 s + f(s) \log s)$ space by Lemma 7. Note that after each invocation of Lemma 7, its workspace can be reused.

Thus on input 1^n , the space used to construct the new circuit is at most $O(h(n) + \log^2 s + g(s) \log s) = O(h(n) + g(s) \log s)$ since $g(s) = \Omega(\log s)$. \square

6 Circuit Value Problem

The Boolean Circuit Value problem is defined as follows: given the standard description of a circuit C and an assignment x to the variables of C as the input, compute the value of the output of the circuit C evaluated on the assignment x . As an application of Theorem 8, we obtain a bound on the space complexity of the problem for Boolean circuits with small segregators (or separators). We need the following theorem of Borodin [4].

Theorem C [4] *Any language decided by a $h(n)$ -space uniform circuit family of depth $h(n) \geq \log n$, can be decided by a Turing machine in space $O(h(n))$.*

Theorem 9. *The Boolean Circuit Value problem for circuits that have size s and segregators (or separators) of size $f(s)$ is in deterministic $SPACE(f(s) \log s)$ if $f(s) = \Omega(s^\varepsilon)$ for some constant $\varepsilon > 0$, and $SPACE(f(s) \log^2 s)$ otherwise.*

Proof. Let $g(s) = f(s)$ if $f(s) = \Omega(s^\varepsilon)$ for some constant $\varepsilon > 0$, and $g(s) = f(s) \log s$ otherwise. Since the description of C is given in the input, by the proof of Theorem 8, using $O(g(s) \log s)$ space, we can generate a circuit C' of depth $O(g(s))$ that simulates C . Then we can evaluate C' in the given assignment using the argument of Theorem C using space $O(g(s))$.

Theorem 9 immediately implies the following theorem.

Theorem 10. *The Boolean Circuit Value problem for circuits with constant-size segregators (or separators) is in deterministic $SPACE(\log^2 n)$.*

Lipton and Tarjan [20] gave the following “planar separator theorem”.

Theorem D [20] *Any planar graph of size s has a separator of size $O(\sqrt{s})$.*

We use this to obtain our result about the space complexity of the Circuit Value Problem for planar graphs.

Theorem 11. *The Planar Circuit Value Problem can be decided in deterministic $SPACE(\sqrt{n} \log n)$.*

Proof. Immediately follows from Theorem D and Theorem 9.

References

1. N. Alon, P. Seymour and R. Thomas: A Separator Theorem for Graphs with an Excluded Minor and its Applications. *Proceedings of STOC*, pp. 293–299 (1990).
2. D. Barrington, C. Lu, P. B. Miltersen and S. Skyum: On monotone planar circuits. *Proceedings of IEEE Conference on Computational Complexity*, pp. 24–33, (1999).
3. M. Bonnet and S. R. Buss: Size-depth tradeoffs for Boolean formulae. *Information Processing Letters*, 49(3), pp. 151–155 (1994).
4. A. Borodin: On Relating Time and Space to Size and Depth. *SIAM Journal on Computing*, 6(4), pp. 733–744 (1977).

5. R. P. Brent: The Parallel Evaluation of General Arithmetic Expressions. *Journal of the ACM*, 21(2), pp. 201–206 (1974).
6. N. Bshouty, R. Cleve and W. Eberly: Size-Depth Tradeoffs for Algebraic Formulas. *SIAM Journal on Computing*, 24(4), pp. 682–705 (1995).
7. P. Bürgisser, M. Clausen, and M. Amin Shokrollahi: Algebraic complexity theory. Springer-Verlag, Berlin-Heidelberg-New York (1997).
8. S. R. Buss: The Boolean formula value problem is in ALOGTIME. *Proceedings of STOC*, pp. 123–131 (1987).
9. S. Buss, S. Cook, A. Gupta, and V. Ramachandran: An Optimal Parallel Algorithm for Formula Evaluation. *SIAM Journal on Computing*, 21(4), pp. 755–780 (1992).
10. T. Chakraborty and S. Datta: One-Input-Face MPCVP Is Hard for L, but in LogDCFL. *Proceedings of FSTTCS 2006, LNCS 4337*, pp. 57–68 (2006).
11. R. G. Downey and M. R. Fellows: Parameterized Complexity. Springer (1999).
12. P. Dymond and M. Tompa: Speedups of Deterministic Machines by Synchronous Parallel Machines. *J. Comp. and Sys. Sci.*, 30(2), pp. 149–161 (1985).
13. J. Flum and M. Grohe: Parameterized Complexity Theory. Springer (2006).
14. A. Gál and J. Jang: The size and depth of layered Boolean circuits. *Proceedings of LATIN 2010, LNCS 6034*, pp. 372–383 (2010).
15. J.R. Gilbert, J.P. Hutchinson, and R.E. Tarjan: A Separator Theorem for Graphs of Bounded Genus. *Journal of Algorithms*, 5(3), pp. 391–407 (1984).
16. L. Goldschlager: The monotone and planar circuit value problem is complete for P. *SIGACT News*, pp. 25–27 (1977).
17. M. Jansen and J. Sarma M. N.: Balancing Bounded Treewidth Circuits. *5th International Computer Science Symposium in Russia*, pp. 228–239 (2010).
18. R. E. Ladner: The circuit value problem is log-space complete for P. *SIGACT News*, 6(2), pp. 18–20 (1975).
19. N. Limaye, M. Mahajan, and J. Sarma: Upper bounds for monotone planar circuit value and variants. *Computational Complexity*, 18, pp. 377–412 (2009).
20. R. Lipton and R.E. Tarjan: A Separator Theorem for Planar Graphs. *SIAM J. Appl. Math.*, 36, pp. 177–189 (1979).
21. N. A. Lynch: Log space recognition and translation of parenthesis languages. *J. Assoc. Comput. Mach.*, 24, pp. 583–590 (1977).
22. C. H. Papadimitriou: Computational complexity. Addison-Wesley (1994).
23. M. S. Paterson and L. G. Valiant: Circuit Size is Nonlinear in Depth. *Theoretical Computer Science*, 2(3), pp. 397–400 (1976).
24. W. J. Paul, N. Pippenger, E. Szemerédi, and W. T. Trotter: On determinism versus non-determinism and related problems. *Proceedings of FOCS*, pp. 429–438 (1983).
25. N. Robertson and P. D. Seymour: Graph Minors II. Algorithmic aspects of tree width. *Journal of Algorithms*, 7, pp. 309–322 (1986).
26. R. Santhanam: On separators, segregators and time versus space. *Proceedings of the Sixteenth Annual Conference on Computational Complexity*, pp. 286–294 (2000).
27. W. J. Savitch: Relationships Between Nondeterministic and Deterministic Tape Complexities. *Journal of Computer and System Sciences*, 4(2), pp. 177–192 (1970).
28. P. M. Spira: On time-hardware complexity tradeoffs for Boolean functions. *Proc. 4th Hawaii Symp. on System Sciences*, pp. 525–527 (1971).
29. L. Valiant, S. Skyum, S. Berkowitz, and C. Rackoff: Fast Parallel Computation of Polynomials Using Few Processors. *SIAM J. Comp.*, 12(4), pp. 641–644 (1983).
30. Ingo Wegener: Relating monotone formula size and monotone depth of Boolean functions. *Information Processing Letters*, 16(1), pp. 41–42 (1983).
31. Ingo Wegener: The Complexity of Boolean Functions. (1987).