

On the Cook-Mertz Tree Evaluation procedure

Oded Goldreich

Department of Computer Science
Weizmann Institute of Science, Rehovot, ISRAEL.

June 30, 2024

Abstract

The input to the Tree Evaluation problem is a binary tree of height h in which each internal vertex is associated with a function mapping pairs of ℓ -bit strings to ℓ -bit strings, and each leaf is assigned an ℓ -bit string. The desired output is the value of the root, where the value of each internal node is defined by applying the corresponding function to the value of its children.

We provide an exposition and a digest of the recent result of Cook and Mertz (*ECCC*, TR23-174), which asserts that Tree Evaluation problem can be solved in space $O((h + \ell) \cdot \log \ell)$. In particular, we point out that the algebraic manipulation (using roots of unity) performed in the original work is merely a special case of univariate polynomial interpolation. Using this observation we provide a more transparent exposition of their main result as well as its low order quantitative improvement (i.e., space complexity $O(\ell + h \cdot \log \ell)$).

Our exposition refers to the “global storage” model rather than to the “catalytic storage” model used by Cook and Mertz, which can be viewed as a special case. We believe that the global storage model is more flexible and intuitive, but our exposition can be easily adapted to the catalytic storage model.

Contents

| | | |
|----------|--|----------|
| 1 | Main text: An alternative exposition of [1, Thm. 10] | 1 |
| 2 | Digest and beyond | 4 |
| 3 | The global storage model (mainly reproduced from [2, Sec. 5.2.4.2]) | 5 |
| 4 | Tedious details | 6 |

1 Main text: An alternative exposition of [1, Thm. 10]

In this section we provide an alternative exposition of Cook and Mertz theorem [1, Thm. 10] that asserts that the Tree Evaluation problem can be solved in space $O(\log n \cdot \log \log n)$, where n denotes the length of the input. In particular, we point out that the algebraic manipulation (relying on roots of unity) performed in [1] is merely a special case of univariate polynomial interpolation. This observation allows for a more transparent presentation of the proof of [1, Thm. 10] as well as a (low order) quantitative improvement [1, Thm. 15] (presented in Section 2).¹

The Tree Evaluation problem ($\text{TrEv}_{h,\ell}$). The input to this computational problem is a rooted binary tree of height h in which internal nodes represent arbitrary gates mapping pairs of ℓ -bit strings to ℓ -bit strings, and each leaf carries an ℓ -bit string. Specifically, nodes in the tree are associated with strings of length at most h such that the nodes $u0$ and $u1$ are the children of the node $u \in U \stackrel{\text{def}}{=} \bigcup_{i=0}^{h-1} \{0,1\}^i$. For every $u \in U$, the internal node u is associated with a gate $f_u : \{0,1\}^{\ell+\ell} \rightarrow \{0,1\}^\ell$, and the leaf $u \in \{0,1\}^h$ is assigned the value $v_u \in \{0,1\}^\ell$. Hence, the input is the description of all $|U| = 2^h - 1$ gates (i.e., all f_u 's) and the values assigned to the 2^h leaves; that is, the length of the input is $(2^h - 1) \cdot (2^{2\ell} \cdot \ell) + 2^h \cdot \ell = \exp(\Theta(h + \ell))$. The desired output is v_λ such that for every $u \in U$ it holds that

$$v_u = f_u(v_{u0}, v_{u1}). \quad (1)$$

(For the history and significance of the Tree Evaluation problem, see [1].)

The straightforward recursive algorithm. Observing that the value at node u is determined by the values at its two children, we compute v_u by first making a recursive call for the value of v_{u0} and then making a recursive call for the value of v_{u1} . Hence, before making the second recursive call, we maintain the value v_{u0} in the local memory of the current execution (which refers to node u). Once we obtain v_{u1} , we compute v_u and output it. The crucial point is that each level of recursion uses a local memory that is different from the memory that is used by other levels. Hence, the space complexity of the algorithm that unravels the recursion is $O(h \cdot \ell)$.

Towards the improved (recursive) algorithm. The first step is conceptual: It consists of abandoning the paradigm of “good programming” under which a recursive call uses a different work space than the execution that calls it. Instead, we shall use the same *global space* for both executions, whereas only a much smaller work space will be allocated to each recursive level as its *local space*. (Such a model, spelled-out in Section 3, was used by us in [2, Sec. 5.2.4.2]; the “catalytic space model” used by [1] is a special case.)

The key question is how to implement the foregoing recursion in this (global storage) model. For starters, suppose that the global memory holds three ℓ -bit strings, denoted x , y and z . Further suppose that we have a procedure that, for any $u \in U$ and $\sigma \in \{0,1\}$, when invoked with $(u\sigma, x, y, z)$ on the global space, returns $(u\sigma, x, y, z \oplus v_{u\sigma})$ on the global space, where $v_{u\sigma}$ is recursively defined as in Eq. (1). Then, when invoked with (u, x, y, z) on the global space, we can return $(u, x, y, z \oplus v_u)$ (such that $v_u = f_u(v_{u0}, v_{u1})$) by proceeding as follows:

¹Referring to the parameters h and ℓ as defined below, the improvement is from the space complexity $O((\ell+h) \cdot \log \ell)$ to space complexity $O(\ell + h \cdot \log \ell)$.

1. Making a recursive call on $(u0, y, z, x)$, we update the global space to $(u0, y, z, x')$, where $x' \stackrel{\text{def}}{=} x \oplus v_{u0}$.
(Note that we re-arranged the parts of the global space so that the variable holding x is updated (to a value denoted x') and the other variables are left intact.)
2. Making a recursive call on $(u1, x', z, y)$, we update the global space to $(u1, x', z, y')$, where $y' \stackrel{\text{def}}{=} y \oplus v_{u1}$.
3. *Miraculously* compute $z' \stackrel{\text{def}}{=} z \oplus f_u(v_{u0}, v_{u1})$ based on $x' = x \oplus v_{u0}$ and $y' = y \oplus v_{u1}$, while preserving the values of x' and y' .
4. Making a recursive call on $(u0, y', z', x')$, we update the global space to $(u0, y', z', x)$.
(Note that $x' \oplus v_{u0}$ equals the original value of x .)
5. Making a recursive call on $(u1, x, z', y')$, we update the global space to $(u1, x, z', y)$.
6. Return (u, x, y, z') .

Indeed, the problem is with the *miraculous step* (i.e., Step 3): We wish to compute $z \oplus f_u(v_{u0}, v_{u1})$, but we don't have v_{u0} and v_{u1} , but rather versions of these values that are masked by the original values of x and y , respectively. There is hope for such a miracle only if we have a few versions of this masking. Suppose, for example, that f_u were a linear (over $\text{GF}(2)$) function and that we have the values of $f_u(x', y')$ and $f_u(x, y)$; then, using $f_u(x', y') \oplus f_u(x, y) = f_u(x' \oplus x, y' \oplus y)$, we can obtain $f_u(x' \oplus x, y' \oplus y) = f_u(v_{u0}, v_{u1})$. This ignores the problem of having to store both $f_u(x', y')$ and $f_u(x, y)$. The last problem can be overcome if we deal with the bits of these ℓ -bit values one at a time; that is, for each $i \in [\ell]$, we first compute the i^{th} of $f_u(x', y')$ and of $f_u(x, y)$, and then obtain the corresponding bit of $f_u(v_{u0}, v_{u1})$.

Multi-linear extensions and interpolation Needless to say, we do not want to assume that the f_u 's are linear. The alternative of using multi-linear extensions (of functions describing the output bits) arises naturally. Indeed, we consider multi-linear extensions of the corresponding functions, where these extensions are in a (prime) field \mathcal{K} that contains at least $2\ell + 2$ elements. Specifically, for every $u \in U$ and $i \in [\ell]$, let $f_{u,i}(x, y)$ equal the i^{th} bit of $f_u(x, y)$. Next, we define $\widehat{f}_{u,i} : \mathcal{K}^\ell \times \mathcal{K}^\ell \rightarrow \mathcal{K}$ to be the multi-linear extension of $f_{u,i} : \{0, 1\}^\ell \times \{0, 1\}^\ell \rightarrow \{0, 1\}$. Now, suppose that we are given the values of $\widehat{f}_{u,i}(j\widehat{x} + v_0, j\widehat{y} + v_1)$ for every $j \in \{1, \dots, 2\ell + 1\} \subset \mathcal{K}$, where $j \cdot (z_1, \dots, z_\ell) = (jz_1, \dots, jz_\ell)$. Using polynomial interpolation (on the degree 2ℓ univariate polynomial (in j) obtained by fixing $u, i, \widehat{x}, \widehat{y}, v_0$ and v_1), we obtain $\widehat{f}_{u,i}(0\widehat{x} + v_0, 0\widehat{y} + v_1)$. Note, however, that a naive implementation of this interpolation involves operating on these $2\ell + 1$ values (after storing them in memory). Fortunately, *the interpolation formula is a linear combination of these $2\ell + 1$ values, and so we need not store these values but can rather operate on them on-the-fly* (while only storing the partial linear combination computed so far).

Actually, as observed in [1], using specific interpolation points allows for a more explicit interpolation that merely sums-up the values (rather than using a more general linear combination). Specifically, these interpolation points are powers of an m^{th} root of unity (in \mathcal{K}), where $m > \ell' \stackrel{\text{def}}{=} 2\ell$

and $m < |\mathcal{K}| = O(\ell)$. Denoting such a root by ω , we observe that for any multi-linear polynomial $p : \mathcal{K}^{\ell'} \rightarrow \mathcal{K}$ it holds that

$$\sum_{j \in [m]} p(\omega^j z_1 + w_1, \dots, \omega^j z_{\ell'} + w_{\ell'}) = m \cdot p(w_1, \dots, w_{\ell'}). \quad (2)$$

(Eq. (2) can be proved by considering each monomial separately.)²

The improved (recursive) algorithm. For sake of simplicity, we first assume that we have oracle access to $F : U \times [\ell] \times \mathcal{K}^{2\ell} \rightarrow \mathcal{K}$ defined by

$$F(u, i, \hat{x}, \hat{y}) \stackrel{\text{def}}{=} \hat{f}_{u,i}(\hat{x}, \hat{y}). \quad (3)$$

The global memory that we use will hold three ℓ -long sequences over \mathcal{K} , denoted \hat{x} , \hat{y} and \hat{z} , as well as a string of length at most h , denoted u . Now, suppose that we have a procedure that, for any $u \in U$ and $\sigma, \tau \in \{0, 1\}$, when invoked with $(u\sigma, \tau, \hat{x}, \hat{y}, \hat{z})$ on the global space, returns $(u\sigma, \hat{x}, \hat{y}, \hat{z} + (-1)^\tau \cdot v_{u\sigma})$ on the global space, where $v_{u\sigma} \in \{0, 1\}^\ell \subset \mathcal{K}^\ell$ is recursively defined as in Eq. (1).³ Then, when invoked with $(u, \tau, \hat{x}, \hat{y}, \hat{z})$ on the global space, we can return $(u, \hat{x}, \hat{y}, \hat{z} + (-1)^\tau v_u)$ such that $v_u = f_u(v_{u0}, v_{u1})$, by proceeding in m iterations.⁴

In iteration $j \in [m]$, for each $i \in [\ell]$, we increment the current value of the i^{th} element of \hat{z} by $(-1)^\tau \cdot \hat{f}_{u,i}(\omega^j \hat{x} + v_{u0}, \omega^j \hat{y} + v_{u1})/m$, while maintaining (u, \hat{x}, \hat{y}) intact.

Recall that, by Eq. (2), $\sum_{j \in [m]} \hat{f}_{u,i}(\omega^j \hat{x} + v_{u0}, \omega^j \hat{y} + v_{u1})/m$ equals $\hat{f}_{u,i}(v_{u0}, v_{u1})$.

The j^{th} iteration proceeds as follows.

1. Making a recursive call on $(u0, 0, \hat{y}, \hat{z}, \omega^j \hat{x})$, we update the global space to $(u0, \hat{y}, \hat{z}, \hat{x}')$, where $\hat{x}' \stackrel{\text{def}}{=} \omega^j \hat{x} + v_{u0}$.
2. Making a recursive call on $(u1, 0, \hat{x}', \hat{z}, \omega^j \hat{y})$, we update the global space to $(u1, \hat{x}', \hat{z}, \hat{y}')$, where $\hat{y}' \stackrel{\text{def}}{=} \omega^j \hat{y} + v_{u1}$.
3. For each $i \in [\ell]$, letting \hat{z}_i denote the i^{th} element of $\hat{z} \in \mathcal{K}^\ell$, compute $\hat{z}_i + (-1)^\tau \cdot F(u, i, \hat{x}', \hat{y}')/m$ by making an oracle call to F , and update the value of \hat{z}_i accordingly. Note that in the i^{th} sub-step only the i^{th} element of the sequence \hat{z} is updated (and that division by m compensates for the factor of m in Eq. (2)).

²For any $I \subseteq [\ell']$, it holds that

$$\begin{aligned} \sum_{j \in [m]} \prod_{i \in I} (\omega^j z_i + w_i) &= \sum_{j \in [m]} \sum_{S \subseteq I} \left(\prod_{i \in S} \omega^j z_i \right) \cdot \left(\prod_{i \in I \setminus S} w_i \right) \\ &= \sum_{S \subseteq I} \sum_{j \in [m]} \omega^{j \cdot |S|} \left(\prod_{i \in S} z_i \right) \cdot \left(\prod_{i \in I \setminus S} w_i \right) \\ &= m \cdot \prod_{i \in I} w_i, \end{aligned}$$

where the last equality uses $\sum_{j \in [m]} \omega^{j \cdot s} = 0$ for $s \in [\ell'] \subseteq [m-1]$ and $\sum_{j \in [m]} \omega^0 = m$.

³The variable/parameter τ allows us to either add or subtract the value $v_{u\sigma}$. In our recursive calls, we shall need both options.

⁴The following description is for the case of $u \in U$. In case $u \in \{0, 1\}^h$, we may just obtain v_u from the input oracle (e.g., augment F such that $F(u) = v_u$).

4. Making a recursive call on $(u0, 1, \hat{y}', \hat{z}, \hat{x}')$, we update the global space to $(u0, \hat{y}', \hat{z}, \omega^j \hat{x})$.
(Note that $\hat{x}' - v_{u0} = \omega^j \hat{x}$.)
5. Making a recursive call on $(u1, 1, \omega^j \hat{x}, \hat{z}, \hat{y}')$, we update the global space to $(u1, \omega^j \hat{x}, \hat{z}, \omega^j \hat{y})$.
6. Re-arrange the global space to contain $(u, \hat{x}, \hat{y}, \hat{z})$, while noting that each \hat{z}_i got incremented by $(-1)^\tau \cdot \hat{f}_{u,i}(\omega^j \hat{x} + v_{u0}, \omega^j \hat{y} + v_{u1})/m$.

Using Eq. (2), we note that (after the m iterations) the value of each \hat{z}_i equals the initial value plus $(-1)^\tau \cdot \hat{f}_{u,i}(v_{u0}, v_{u1}) = (-1)^\tau \cdot v_u$.

The foregoing recursive procedure uses a global space of length $h + O(1) + (3 + o(1)) \cdot \log_2 |\mathcal{K}|^\ell = O(h + \ell \log |\mathcal{K}|)$ and a local space of length $\log_2 m = O(\log \ell)$. (The $o(1) \cdot \log_2 |\mathcal{K}|^\ell$ term accounts for the space complexity of various manipulations (including maintaining the counter $i \in [\ell]$), whereas the local space is used only for recording $j \in [m]$.)

Using a composition lemma akin [2, Lem. 5.10] (reproduced as Lemma 4), it follows that the Tree Evaluation problem (with parameters h and ℓ) can be solved in space $O((h + \ell) \cdot \log \ell)$, when using oracle access to F , which in turn can be evaluated in linear space (i.e., space linear in $h + O(\ell \log |\mathcal{K}|)$).⁵ Using a naive composition (see Section 4 for details), it follows that

Theorem 1 (Cook and Mertz [1, Thm. 10]): *The space complexity of $\text{TrEv}_{h,\ell}$ is $O((h + \ell) \cdot \log \ell)$.*

Recalling that the length of the input to $\text{TrEv}_{h,\ell}$ is exponential in $h + \ell$, it follows that $\text{TrEv}_{h,\ell}$ is solved in space $O(\log n \cdot \log \log n)$, where $n = \exp(\Theta(h + \ell))$.

2 Digest and beyond

As hinted above, we believe that the model of global storage (as outlined in [2, Def. 5.8] and reproduced in Section 3) is more flexible and intuitive than the model of catalytic storage used in [1], which may be viewed as a special case.⁶ Hence, we used the global storage model rather than the catalytic storage model in our exposition.

As hinted above, the interpolation formula given in Eq. (2), which relies on an m^{th} root of unity, is inessential for the proof of Theorem 1. More generally, recalling that the $\hat{f}_{u,i}$'s are polynomials of total degree 2ℓ , we can use univariate polynomial interpolation based on any $2\ell' + 1$ points (on a line that passes through the desired point), while noting that such interpolation can be represented by a linear combination of the polynomial's values (with coefficients that depend on the interpolation points and the desired point). The only advantage of using Eq. (2) is that the interpolation formula is a simple sum (i.e., all coefficients are 1).

Capitalizing on the last paragraph, we can reduce the length of global storage used by the recursive procedure from $O(h + \ell \log \ell)$ to $O(h + \ell)$. This can be done by viewing the f_u 's as functions from $[k]^k \times [k]^k$ to $[k]^k$, where $k^k = 2^\ell$ (i.e., $k = \Theta(\ell / \log \ell)$), and using low degree extensions of the

⁵Recall that computing F calls for computing the corresponding $\hat{f}_{u,i}$, which is a multi-linear extension of $f_{u,i}$. As for computing $\hat{f}_{u,i}$, it requires obtaining all values of $f_{u,i}$ (cf. Footnote 8).

⁶To see that the catalytic is a special case of the global model, we note a correspondence between the catalytic (resp., ordinary) storage of the catalytic model and the global (resp., local) memory of the model defined in Section 3. A “clean computation” (in the catalytic model) that results in adding a desired value to one set of registers while keeping another set of registers intact can be emulated by a corresponding transformation of the global storage.

corresponding $f_{u,i}$'s (i.e., $f_{u,i}(x, y) \in [k]$ is the i^{th} symbol in $f_u(x, y) \in [k]^k$).⁷ Specifically, these extensions are $2k$ -variate polynomials of individual degree $k - 1$ over \mathcal{K} , where \mathcal{K} is a finite field of size $\text{poly}(k)$ that is greater than $m = 2k^2$ (and $[k] \subset \mathcal{K}$).⁸ Thus, $\hat{f}_{u,i} : \mathcal{K}^k \times \mathcal{K}^k \rightarrow \mathcal{K}$ has total degree $2k \cdot (k - 1) < m$, whereas its input length (i.e., $\log_2 |\mathcal{K}^{2k}|$) equals $\log_2(\text{poly}(k)^{2k}) = O(\ell)$. Consequently, we can obtain the value of $\hat{f}_{u,i}(v_0, v_1)$ by univariate polynomial interpolation from the values of $\hat{f}_{u,i}(j\hat{x} + v_0, j\hat{y} + v_1)$ for all $j \in [m]$. Hence, the revised recursive procedure uses a global space of length $O(h + \ell)$ and local space of length $\log_2 m = O(\log \ell)$. Again, using a composition lemma akin [2, Lem. 5.10], it follows that $\text{TrEv}_{h,\ell}$ can be solved in space $O(\ell + h \cdot \log \ell)$, when using oracle access to $F : U \times [\ell] \times \mathcal{K}^{2k} \rightarrow \mathcal{K}$, which in turn can be evaluated in linear space (i.e., space linear in $O(h + \ell)$).⁹ Hence (see Section 4 for details), we obtain

Theorem 2 (Cook and Mertz [1, Thm. 15]): *The space complexity of $\text{TrEv}_{h,\ell}$ is $O(\ell + h \cdot \log \ell)$.*

In particular, for $h = O(\ell / \log \ell)$, the problem can be solved in logarithmic space (because, in this case, the input length is $n = \exp(\Theta(\ell))$, whereas $O(\ell + h \cdot \log \ell) = O(\ell) = O(\log n)$).

Needless to say, the question of whether $\text{TrEv}_{h,\ell}$ can be solved in $O(\ell + h)$ space remains open. The stumbling block for our approach is that we use $O(\log \ell)$ bits of local storage for indexing $\text{poly}(\ell)$ different evaluations of $\hat{f}_{u,i}$.

3 The global storage model (mainly reproduced from [2, Sec. 5.2.4.2])

(This model was introduced in [2, Sec. 5.2.4] in order to facilitate a modular presentation of Reingold's UCONN algorithm [3].)

The aim of this model is to support a composition result that is beneficial in the context of recursive calls. The basic idea is deviating from the paradigm that allocates *separate* input/output and query devices to *each level in the recursion*, and combining all these devices in a single (“global”) device, which will be used by all levels of the recursion. That is, rather than following the “structured programming” methodology of using locally designated space for passing information to the subroutine, we use the “bad programming” methodology of passing information through global variables. (As usual, this notion is formulated by referring to the model of multi-tape Turing machine, but it can be formulated in any other reasonable model of computation.)

Definition 3 (following [2, Def. 5.8]): *A global-tape oracle machine is defined as an oracle machine (cf. [2, Def. 1.11]), except that the input, output and oracle tapes are replaced by a single global-tape. In addition, the machine has a constant number of work tapes, called the local-tapes. The machine obtains its input from the global-tape, writes each query on this very tape, obtains the corresponding*

⁷More generally, we may replace $\{0, 1\}^\ell$ by S^k such that $k = \Theta(\ell / \log \ell)$ and $S \subset \mathcal{K}$ has size $2^{\ell/k}$.

⁸Indeed, for simplicity, we assume that \mathcal{K} is of prime cardinality. In general, for $S \subset \mathcal{K}$, the low degree extension of $f : S^t \rightarrow S$ is given by $\hat{f} : \mathcal{K}^t \rightarrow \mathcal{K}$ such that

$$\hat{f}(x_1, \dots, x_t) = \sum_{a_1, \dots, a_t \in S} \left(\prod_{i \in [t]} \chi_{a_i}(x_i) \right) \cdot f(a_1, \dots, a_t),$$

where $\chi_a(x) \stackrel{\text{def}}{=} \prod_{b \in S \setminus \{a\}} (x - b) / (a - b)$ is a degree $|S| - 1$ univariate polynomial.

⁹We can also assign to F the task of providing the coefficients of the linear combination used in the interpolation. As detailed in Section 4, these coefficients can be computed in $o(\ell)$ space.

answer from this tape, and writes its final output on this tape. (We stress that, as a result of invoking the oracle f , the contents of the global-tape changes from q to $f(q)$.)¹⁰ In addition, the machine can use the global-tape also for its internal computations. The space complexity of such a machine is stated when referring separately to its use of the global-tape and to its use of the local-tapes.

Note that in our presentation of Theorem 1 we also used oracle calls to a function F . This was done for the sake of simplicity, and these oracle calls (unlike the recursive calls) can be modeled by the usual mechanism (of oracle tapes).

Composing global-tape oracle machines. As stated above, global-tape oracle machines are beneficial in the context of *recursive composition*, as indicated by Lemma 4 (which relies on this model in a crucial way). The key observation is that all levels in the recursive composition may re-use the same global storage, and only the local storage gets added. Consequently, we have the following composition lemma.

Lemma 4 (recursive composition in the global-tape model [2, Lem. 5.10]): *Suppose that there exists a global-tape oracle machine that, for every $i = 1, \dots, t-1$, computes f_{i+1} by making oracle calls to f_i while using a global-tape of length L and a local-tape of length l_i , which also accounts for the machine's state. Then, f_t can be computed by a standard oracle machine that makes calls to f_1 and uses space $L + \sum_{i=1}^{t-1} (l_i + \log_2 l_i)$.*

Proof Sketch: We compute f_t by allocating space for the emulation of the global-tape and the local-tapes of each level in the recursion. We emulate the recursive computation by capitalizing on the fact that all recursive levels use the same global-tape (for making queries and receiving answers). Recall that in the actual recursion, each level may use the global-tape arbitrarily as long as when it returns control to the invoking machine the global-tape contains the correct answer. Thus, the emulation may do the same, and emulate each recursive call by using the space allocated for the global-tape as well as the space designated for the local-tape of this level. The emulation should also store the locations of the other levels of the recursion on the corresponding local-tapes, which is accounted for by the extra $\sum_{i=1}^{t-1} \log_2 l_i$ term. ■

4 Tedious details

Recall that Theorems 1 and 2 were proved by using a composition lemma akin [2, Lem. 5.10], which implies a space bound on a procedure (in the standard model) that computes $\text{TrEv}_{h,\ell}$ when given oracle access to F . As for F , it was assigned the task for computing the $\hat{f}_{u,i}$'s, providing the values of the v_u 's for all $u \in \{0, 1\}^\ell$, and computing the coefficients of the linear combination that underlies the interpolation procedure. We stress that, while the latter v_u 's and all $f_{u,i}$'s appear explicitly in the input to $\text{TrEv}_{h,\ell}$, the $\hat{f}_{u,i}$'s and the interpolation coefficients need to be computed. We address both tasks in the more general setting of the proof of Theorem 2.

¹⁰This means that the prior contents of the global-tape (i.e., the query q) is lost (i.e., it is replaced by the answer $f(q)$). Thus, if we wish to keep such prior contents, then we need to copy it to a local-tape. We also stress that, according to the standard oracle invocation conventions, the head location after the oracle responds is at the left-most cell of the global-tape.

Computing the $\widehat{f}_{u,i}$'s. As stated in Footnote 8, for a prime field \mathcal{K} , the low degree extension of $f_{u,i} : [k]^k \rightarrow [k]$ is given by $\widehat{f} : \mathcal{K}^k \rightarrow \mathcal{K}$ such that

$$\widehat{f}_{u,i}(x_1, \dots, x_k) = \sum_{a_1, \dots, a_k \in [k]} \left(\prod_{s \in [k]} \chi_{a_s}(x_s) \right) \cdot f_{u,i}(a_1, \dots, a_k), \quad (4)$$

where $\chi_a(x) \stackrel{\text{def}}{=} \prod_{b \in [k] \setminus \{a\}} (x - b)/(a - b)$ is a degree $k - 1$ univariate polynomial over \mathcal{K} . Hence, $\widehat{f}_{u,i}$ can be computed by going over all possible $(a_1, \dots, a_k) \in [k]^k$ (and all $s, b \in [k]$), which can be done using $\log_2(k^k \cdot k^2) = (1 + o(1)) \cdot \ell$ space.

Computing the interpolation coefficients. The desired coefficients for the interpolation (of a univariate polynomial (based on m evaluation points)) are the first row of the inverse of an m -by- m Vandermonde matrix over \mathcal{K} . While the Vandermonde matrix has a simple explicit form (i.e., its entries are powers of the evaluation points), its inverse has a simple explicit form only in some cases (i.e., for some structured sequence of evaluation points). We can use such a structured sequence, but prefer not to rely on such low level considerations. Instead, we use the fact that matrix inversion is in NC, and hence can be computed in polylogarithmic space, whereas here we the input length is $m^2 \cdot \log_2 |\mathcal{K}| = \text{poly}(\ell)$. Hence, the inverse of the relevant matrix can be computed in space $\text{poly}(\log(\ell)) = o(\ell)$.

Conclusion. The function F can be evaluated in space that is linear in the length of the main part of its input (i.e., linear in $|\widehat{x}| + |\widehat{y}|$, which is $2\ell \cdot \log_2 |\mathcal{K}| = O(\ell \log \ell)$ in the proof of Theorem 1 and $2k \cdot \log_2 |\mathcal{K}| = O(\ell)$ in the proof of Theorem 2). Composing the (standard model) procedure that computes $\text{TrEv}_{h,\ell}$ (when given oracle access to F) with the algorithm for computing F , we derive the claimed results.

Acknowledgments

I am grateful to Amnon Ta-Shma, Ben Chen, and Madhu Sudan for helpful discussions and comments.

References

- [1] James Cook and Ian Mertz. Tree Evaluation is in Space $O(\log n \cdot \log \log n)$. *ECCC*, TR23-174, 2023.
- [2] Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.
- [3] Omer Reingold. Undirected ST-Connectivity in Log-Space. In *37th ACM Symposium on the Theory of Computing*, pages 376–385, 2005.