

# On Treewidth, Separators and Yao’s Garbling

Chethan Kamath<sup>1\*</sup>, Karen Klein<sup>2\*\*</sup>, and Krzysztof Pietrzak<sup>2\*\*</sup>

<sup>1</sup> ckamath@protonmail.com

<sup>2</sup> IST Austria, {kklein,pietrzak}@ist.ac.at

**Abstract.** We show that Yao’s garbling scheme is adaptively indistinguishable for the class of Boolean circuits of size  $S$  and treewidth  $w$  with only a  $S^{O(w)}$  loss in security. For instance, circuits with constant treewidth are as a result adaptively indistinguishable with only a polynomial loss. This (partially) complements a negative result of Applebaum et al. (Crypto 2013), which showed (assuming one-way functions) that Yao’s garbling scheme cannot be adaptively *simulatable*. As main technical contributions, we introduce a new pebble game that abstracts out our security reduction and then present a pebbling strategy for this game where the number of pebbles used is roughly  $O(\delta w \log(S))$ ,  $\delta$  being the fan-out of the circuit. The design of the strategy relies on separators, a graph-theoretic notion with connections to circuit complexity.

## 1 Introduction

Suppose that Alice, who holds a function represented as a Boolean circuit  $C$ , and Bob, who holds an input  $x$  to that function, want to jointly evaluate  $y = C(x)$  such that Alice learns nothing about  $x$  while Bob learns nothing about  $C$  (except for some side-information that is unavoidable). Yao put forward<sup>3</sup> the following elegant solution:

1. Alice first sends  $\tilde{C}$ , a “garbling” of the circuit  $C$ , to Bob,
2. Bob then obtains  $\tilde{x}$ , a “garbling” of his input  $x$ , from Alice via oblivious transfer,
3. Bob finally evaluates  $\tilde{C}$  on  $\tilde{x}$  to learn  $y$  and sends it over to Alice.

Yao showed how the garbling steps above can be carried out using a symmetric-key encryption (SKE) scheme (and hence one-way functions). This has been ever since referred to as Yao’s garbling scheme, and is the focus of this work. We describe it next in slightly more details.

---

\* Most of the work was done while the author was at Northeastern University, supported by the IARPA grant IARPA/2019-19-020700009, and Charles University, funded by project PRIMUS/17/SCI/9.

\*\* Funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (682815 - TOCNeT).

<sup>3</sup> According to [7], the idea was first presented by Yao in oral presentations on secure function-evaluation [43,44] but formally described only in [22].

*Yao’s garbling scheme.* Let  $(\text{Enc}, \text{Dec})$  be a (special) SKE. To garble a circuit  $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$  with fan-in 2 and arbitrary fan-out:

1. Alice first samples a pair of secret keys  $(k_w^0, k_w^1)$  for each wire  $w$  in  $C$ .
2. For every gate  $g : \{0, 1\}^2 \rightarrow \{0, 1\}$  with left input wire  $u$ , right input wire  $v$ , and output wire  $w$ , she then computes a *garbling table*  $\tilde{g}$  consisting of the four ciphertexts listed in Table 1.(a) in random order.
3. Finally, she constructs the *output mapping*  $\mu$  which, for each output wire  $w$ , maps each of the keys  $(k_w^0, k_w^1)$  to the bit it “encodes”.

The *garbled circuit*  $\tilde{C}$  consists of all the garbling tables  $\tilde{g}$  and the output map  $\mu$ , Alice sends it over to Bob. This constitutes the *offline* phase of the protocol. To garble an input  $x = x_1 \parallel \dots \parallel x_n$ , Alice simply gives out, for each input wire  $w_i$ , the key  $k_{w_i}^{x_i}$  corresponding to the bit  $x_i$ . This constitutes the *online* phase of the protocol. To evaluate the garbled circuit on the garbled input, the encryption scheme must satisfy a *special correctness property*: for each ciphertext  $c \leftarrow \text{Enc}_k(m)$  there should exist a single key (i.e.,  $k$ ) such that decryption passes. Using the keys in the garbling input, Bob can now evaluate  $C$  “over the encryption” as follows:

1. Starting from the input level and in some topological order, he progressively decrypts each garbling table in  $\tilde{C}$  by trying the two keys in hand on all the four ciphertexts for each garbling table. Thus, in each step, he learns one of the secret keys corresponding to the output wire of the gate in consideration.
2. At the end of this process, Bob recovers exactly one of the two keys associated with each output wire of the circuit. This allows him to use the output map  $\mu$  to “decode” the revealed output keys to the output string  $y \in \{0, 1\}^\ell$ .

The scheme as described above is what is regarded to be the original formulation of Yao’s garbling scheme [33,28]. A slight variant in which Alice defers sending the output map  $\mu$  to the online phase (along with  $\tilde{x}$ ) is also of interest [28], although it suffers from a higher online complexity compared to the original formulation. To avoid confusion, we refer to the original scheme as *Yao’s offline garbling scheme* and the modified scheme as *Yao’s online garbling scheme* or, in short, Online Yao and Offline Yao respectively. Our work concerns the security of Offline Yao.

*Security.* Even though garbling schemes found several applications (see [7]), its security was formally analysed much later in [33]. They consider a simulation-based notion<sup>4</sup> captured by the following experiment:

1. The adversary submits a circuit-input pair  $(C, x)$  to the challenger.

---

<sup>4</sup> This is an equivalent formulation of the definition in [33] and is taken from [28]. Our overview of the proof in [33] to be discussed in Section 1.2 has been adapted accordingly.

$E_{k_u^0}(E_{k_v^0}(k_w^{g(0,0)}))$	$E_{k_u^0}(E_{k_v^0}(k_w^0))$	$E_{k_u^0}(E_{k_v^0}(k_w^1))$
$E_{k_u^1}(E_{k_v^0}(k_w^{g(1,0)}))$	$E_{k_u^0}(E_{k_v^1}(k_w^0))$	$E_{k_u^0}(E_{k_v^1}(k_w^1))$
$E_{k_u^0}(E_{k_v^1}(k_w^{g(0,1)}))$	$E_{k_u^1}(E_{k_v^0}(k_w^0))$	$E_{k_u^1}(E_{k_v^0}(k_w^1))$
$E_{k_u^1}(E_{k_v^1}(k_w^{g(1,1)}))$	$E_{k_u^1}(E_{k_v^1}(k_w^0))$	$E_{k_u^1}(E_{k_v^1}(k_w^1))$
(a)	(b)	(c)

**Table 1.** Garbling tables for (a) general gate  $g$  (b) constant-0 gate and (c) constant-1 gate.  $u$  and  $v$  denote the two input wires and  $w$  denotes the output wire. The two keys associated with (say) the wire  $u$  are denoted by  $k_u^0$  and  $k_u^1$ .

2. The challenger responds either with the real garbling  $(\tilde{C}, \tilde{x})$  (i.e., real game or Real) or with a “simulated” garbling where a constant-0 circuit is used instead of  $C$  (i.e., simulated game or Sim). The constant-0 circuit has the same topology as  $C$  but with all its gates replaced by constant-0 gates.
3. The adversary wins if it guesses which case it is.

Then they gave a reduction from the (special) indistinguishability of the underlying SKE for *offline* Yao. Note that the adversary in the above security game must *select* the garbling input  $x$  at the same time as the circuit  $C$ . This is in conflict with the online-offline nature of the actual scheme where Bob (a potential adversary) sees  $\tilde{C}$  *before* he commits to  $x$ . Hence Bob could have chosen the input, *adaptively*, based on  $\tilde{C}$ . In fact, such a scenario does arise in applications such as one-time programs and secure outsourcing [6]. Therefore it is natural to consider strengthening the above selective definition of simulatability to an adaptive definition where A gets to choose the input after it sees the garbling of a circuit of its choice. Unfortunately, this is too strong a notion to attain for Offline Yao: it was shown in [4] that the *online complexity* of a garbling scheme (or, more generally, a randomised encoding scheme) in the adaptive setting *must* exceed the output-size of the circuit (given that one-way functions exist).<sup>5</sup> Jafargholi and Wichs [28] observed that this negative result does not apply to Online Yao since the output map there gets sent in the online phase, and even managed to prove adaptive *simulatability* of Online Yao. Security of other variants of Yao’s garbling scheme was also proved [23,27]. However, the case of Offline Yao was largely ignored.

<sup>5</sup> On a high level, their argument for Offline Yao proceeds as follows. Consider garbling a PRG  $G$  (of sufficient stretch) and suppose that  $|\tilde{x}|$  is less than that of output of the PRG. Then the (online) simulator for such a scheme can be used to break the security of  $G$ . This exploits the fact that the  $\tilde{x}$  loses information when it encodes an element in the co-domain of  $G$  that is not an image (but not for an image). It is possible to circumvent this result in the random-oracle model [6].

## 1.1 Our Results

Although the negative result in [4] rules out adaptive *simulatability* of Offline Yao, it is not clear if it also applies to its adaptive *indistinguishability* [7], which is defined by the following experiment:

1. The adversary submits a pair of circuit  $(C_0, C_1)$  of the same topology to the challenger
2. The challenger. flips a coin  $b$  and responds with  $\tilde{C}_b$ .
3. The adversary then submits a pair of inputs  $(x_0, x_1)$  such that  $C_0(x_0) = C_1(x_1)$  and the challenger responds with  $\tilde{x}_b$ .
4. The adversary wins if it guesses the bit  $b$  correctly.

Although it is a weaker notion of security, adaptive indistinguishability suffices for certain applications (e.g., adaptively-indistinguishable symmetric-key functional encryption [27]).

	Selective		Adaptive	
	Offline Yao	Online Yao	Offline Yao	Online Yao
Simulatability	[33]		[4]	[28]
Indistinguishability			This work	

**Table 2.** Security of Yao garbling and its variants. The (only) negative result is highlighted in red.

*Our results.* We help (partially) complete the landscape for security of Yao’s garbling (see Table 2). To this end, we characterise the adaptive indistinguishability of Offline Yao in terms of the treewidth<sup>6</sup> of the circuit. Our main results are informally stated below.

**Theorem (main).** Consider the class of Boolean circuits  $\mathcal{C}$  of size  $S$  with treewidth  $w = w(S)$ . Offline Yao is adaptively indistinguishable for  $\mathcal{C}$  with  $S^{O(w)}$  loss in security.<sup>7</sup>

For Boolean circuits of constant (resp., poly-logarithmic) treewidth, we obtain the following corollary.

**Corollary.** Offline Yao is adaptively indistinguishable for Boolean circuits of size  $S$  and  $O(1)$  (resp., **polylog**( $S$ )) treewidth with a polynomial (resp., quasi-polynomial) in  $S$  loss in security.

<sup>6</sup> Since treewidth is defined for undirected graphs, whenever we refer to the treewidth of a directed graph (or a circuit) we refer to the treewidth of the graph obtained by ignoring the direction of its edges.

<sup>7</sup> Consider a reduction  $R$  from a problem  $P$  to another problem  $Q$ . Suppose that  $R$  uses an  $(\epsilon, T)$ -adversary  $A$  that breaks  $Q$  in order to  $(\epsilon', T')$ -break  $P$ . Informally, the *loss in security* incurred by  $R$  can be defined as the ratio  $(\epsilon T')/(\epsilon' T)$ .

Finally, exploiting a classical result from graph theory – that planar circuits of size  $S$  have treewidth at most  $\sqrt{S}$  [34] – we obtain the following corollary.

**Corollary.** Offline Yao is adaptively indistinguishable for planar Boolean circuits with a sub-exponential loss in security.

**Interpreting our results.** Treewidth is a notion from algorithmic graph theory that has found several applications in parametrised and circuit complexity (see Section 1.3) Intuitively, it is a (graph) property that measures how “far” the circuit is from a formula (and, more generally, how far a graph is from a tree): in particular, the smaller the treewidth the closer the circuit is to a formula. Therefore, it is not surprising that having a low treewidth limits how powerful a circuit can be. A precise characterisation of this (from above) was given in [20]: every circuit of size  $S$  and treewidth  $w = w(S)$  can be simulated in *depth*  $w \log(S)$ . Thus, e.g., circuits of constant treewidth can be simulated in  $\mathbf{NC}^1$ . Whether the converse is true in general – i.e., whether  $\mathbf{NC}^i$  can be simulated using circuits with treewidth  $O(\log^{i-1}(S))$  – is an open problem to the best of our knowledge.<sup>8</sup> However it is partially true: namely,  $\mathbf{NC}^1$  circuits can be simulated using polynomial-sized Boolean formulae (which, by definition, have treewidth 1) [42,13]. Consequently, the first corollary applies to functions computable in  $\mathbf{NC}^1$ .

Given the aforementioned negative result from [4], we find any proof of adaptive security for Offline Yao rather surprising. Nevertheless, there are scenarios where our results also lead to improvements in concrete efficiency (even after the loss in security is taken into account). We describe one such scenario next. Recall from the discussion above that for functions computable in  $\mathbf{NC}^1$ , we show security of Offline Yao at only a polynomial loss. Moreover, the online complexity of garbling such a function using Offline Yao depends only on the input length  $n$  (times the security parameter  $\lambda$ ). Now, note that PRGs of arbitrary stretch (say  $n^c$  for a constant  $c \in \mathbb{N}$ ) exist in  $\mathbf{NC}^1$  [16,24]. However, if one were to use Online Yao, then the online complexity is substantial ( $n^c \times \lambda$ ). This example is particularly interesting since Offline Yao for such a function is not simulatable at all as a consequence of the negative result.

Finally, a remark on the optimality of our upper bound: it was recently shown in [30] that any *black-box* reduction that proves indistinguishability of Offline Yao (or, for that matter, Online Yao) must lose security by a factor that is sub-exponential in the depth of the circuit. Therefore, there remains a gap between the lower bound proved there and the upper bound shown here.

*Implications to simulatability of Online Yao.* It is worth pointing out that our results may also imply tighter reductions for simulatability of Online Yao. The reduction for simulatability of Online Yao from [28] loses a factor that is exponential in the *width* of a circuit: our approach can be seen as an extension of their techniques. Since treewidth is bounded from above by width, in cases where there

<sup>8</sup> See [this](#) question (48504) posted on CSTheory, Stack Exchange.

is a gap between treewidth and width for a circuit class, our approach would lead to a tighter reduction for simulatability of Online Yao compared to [28]. A more detailed explanation follows later in Remark 1.

*Comparison with [27].* We conclude the section by comparing our result with [27], which is also concerned with adaptively-indistinguishable garbled circuits. The construction in [27] builds on [23] and therefore has Offline Yao as its basis. However, it requires (i) applying an additional layer of *somewhere equivocal encryption* to the garbling table and (ii) modifying the circuit to be garbled in order to make the security proof go through. These modifications lead to their construction being less efficient compared to plain Offline Yao, but it does allow them to prove adaptive indistinguishability. It is not clear if any of the ideas employed there can be used to argue the indistinguishability of Offline Yao (this is, in fact, posed as an open question there).

## 1.2 Technical Overview

*Outline.* Our starting point is the reduction proving adaptive simulatability of Online Yao [28]. The key idea in [28] is to abstract out the hybrid argument using a pebble game on the circuit, which we call the black-gray (BG) pebble game (Definition 9). To be precise, they showed that if a circuit allows a BG pebbling strategy of length  $\tau$  that uses  $\sigma$  (black) pebbles, then there exists a reduction proving adaptive simulatability of Online Yao with a loss in security at most  $O(\tau 2^\sigma)$ . This allows us to shift the focus from security reductions to the conceptually-cleaner task of coming up with “pebble-efficient” strategies. We start off below by describing this connection and then explain why this approach falls short when it comes to arguing adaptive indistinguishability (or simulatability) of Offline Yao. Next we show how this issue can be remedied, key to it is a new pebble game, which we call the black-gray-red (BGR) pebble game (Definition 12). Analogous to [28], we prove that if there exists a BGR pebbling strategy of length  $\tau$  that uses  $\sigma$  (“grayscale”, i.e. black or gray) pebbles, then there exists a reduction for adaptive indistinguishability of Offline Yao where the loss in security is at most  $O(\tau 2^\sigma)$  (Theorem 6). Finally to complete the proof – and as our main technical contribution – we describe a pebble-efficient strategy for the BGR pebble game in which the number of (grayscale) pebbles used grows only with the treewidth of the circuit (Theorem 5). The strategy has a divide-and-conquer flavour and crucially relies on the notion of *separators* from graph theory. We next elaborate on each of the steps above.


**Pebble game and hybrids.** The reduction in [28] builds on the reduction for selective simulatability of Offline Yao [33]. Both these works follow a sophisticated hybrid argument which can be described abstractly using a BG pebbling strategy.

*Pebbles and garbling modes.* The BG pebble game (formally defined in Definition 9), as its name suggests, uses two types of pebbles: black and gray. A pebble configuration  $\mathcal{P}$  for a circuit  $C$  determines how the garbled circuit  $\tilde{C}$  is simulated in the hybrid  $H_{\mathcal{P}}$ . To be more precise, the pebble configuration  $\mathcal{P}$  can associate each gate  $g$  in  $C$  with a black or gray pebble. In order to translate  $\mathcal{P}$  to the garbling  $\tilde{C}$ , the simulator in hybrid  $H_{\mathcal{P}}$  does the following:

- if  $g$  carries no pebble in  $\mathcal{P}$ , then the corresponding garbling table in  $\tilde{C}$  consists of an honest garbling table of  $g$  (Table 1.(a))
- if  $g$  carries a gray pebble, then the garbling table encodes a constant-0 gate (Table 1.(b)).
- if  $g$  carries a black pebble then the garbling table encodes either a constant-0 or a constant-1 gate (Table 1.(c)) *depending* on the value of (the output wire of)  $g$  when  $C$  is run on the garbling input  $x$ .

The three *modes* above of simulating individual gates are named real, simulated and input-dependent modes respectively or, in short, Real, Sim and Input, respectively (Table 3.(a)). Note that the real garbling game corresponds to the empty pebble configuration (since all the gates are honestly garbled), whereas the simulated game will correspond to the all-gray configuration (since all the gates have been replaced by the constant-0 gate).

*Pebbling rules.* Note that any arbitrary configuration of pebbles  $\mathcal{P}$  describes a valid hybrid  $H_{\mathcal{P}}$ . The role of the pebbling rules is to model indistinguishability of neighbouring hybrids. To be more precise, if a pebble configuration  $\mathcal{Q}$  can be obtained from another configuration  $\mathcal{P}$  by a valid pebbling move (or vice versa) then the hybrids  $H_{\mathcal{P}}$  and  $H_{\mathcal{Q}}$  should be indistinguishable. Consequently a BG pebbling strategy  $\mathcal{P}$ , which must start from an empty configuration and end with the all-gray configuration, leads to a valid sequence of hybrids that establishes that the real garbling game and simulated garbling game are indistinguishable, proving the security of the garbling scheme. In the BG pebble game, the following moves (see Figure 1) are allowed:



**Fig. 1.** Rules for BG pebble game.

1. a black pebble can be placed on or removed from a gate  $g$  if and only if  $g$ 's predecessor gates are pebbled black; and
2. a black pebble on a gate  $g$  can be replaced by a gray pebble if  $g$ 's *successor* gates are pebbled, either black or gray.

To understand the rationale behind the two rules, one needs to take a closer look at the structure of a garbling table in Yao’s scheme. Since this is not that relevant to the current discussion, we refer the readers interested in more details to Section 3.

*Selective Simulatability of Offline Yao.* Observe that in order to simulate  $\tilde{C}$  in a hybrid  $H_{\mathcal{P}}$ , the simulator only needs to know the output value of those gates that are pebbled black in  $\mathcal{P}$  (i.e., the gates in `Input` mode). In the selective setting, since the adversary commits to the garbling input  $x$  in the offline phase, the value of *all* the gates is available beforehand. Hence, in this case the simulator has the luxury of using as many pebbles as it needs. Therefore the pebbling strategy (implicitly) employed in [33] is the following:

1. starting from the input gates, pebble the circuit completely black in some topological order, and then
2. starting from the output gates and in reverse topological order, replace each black pebble with a gray pebble.

To complete the description of the hybrid  $H_{\mathcal{P}}$  in the selective setting, one thing remains to be addressed. For concreteness, let’s consider the simulated game, which corresponds to the all-gray pebble configuration (the argument for other hybrids is analogous). Note that it is not possible to send the honestly-generated output map  $\mu$  in  $H_{\mathcal{P}}$  since this will lead to the output being mapped to the all-0 string. However, since  $x$  is available in the offline phase, [33] resolved this issue by *programming* the output map to map the zero-keys of the output wires to  $C(x)$ . The adversary cannot tell this from the honest output map since the change is information-theoretic.

Since the above pebbling strategy takes at most  $2S$  moves (and uses  $S$  black pebbles) the corresponding hybrid argument only loses a  $2S$  factor. It is possible to further reduce to adaptive simulatability via *random guessing*, but this incurs an additional loss in security that is exponential in the length of  $x$ .

*Adaptive Simulatability of Online Yao.* In order to avoid this exponential loss in the adaptive setting, [28] had to mainly tackle two issues, both arising from the fact that the garbling input  $x$  is now only available in the online phase.

1. Firstly, simulating the hybrids could not rely on the knowledge of the values of too many gates in  $C$ .
2. Secondly, the output map could no longer be programmed in the offline phase since the output  $C(x)$  is only determined in the online phase.

The first issue was resolved in [28] by employing BG pebbling strategies that were more frugal in terms of the number of black pebbles used. To this end, they proved that if there exists a BG pebbling strategy of length  $\tau$  that uses  $\sigma$  black pebbles, then the loss in the resulting security is at most  $O(\tau 2^\sigma)$ . Here, loosely speaking, the  $2^\sigma$  factor is the cost of randomly guessing the output values of the gates pebbled black, which they require in order to carry out the simulation of the hybrids (as well as the reduction).<sup>9</sup> To complete their proof, [28] described

<sup>9</sup> This is one of the earliest applications of the piecewise-guessing framework [25].



two (generic) pebbling strategies: one where  $\sigma$  grows only with the width of the circuit and another where  $\sigma$  grows only with the depth of the circuit. A consequence of the latter is the adaptive simulatability of log-depth (i.e.,  $\mathbf{NC}^1$ ) circuits with a *polynomial* loss in security.

The second issue, on the other hand, was basically side-stepped by *modifying* the garbling scheme to defer the sending of the output map to the online phase, i.e., by resorting to *online* Yao. This tweak allowed [28] to carry out a “deferred programming” of the output map since the garbling input is available in the online phase. The cost is an increased online complexity which is now dependent also on the output size.

**Indistinguishability of Offline Yao: Our Approach.** Unfortunately, given the negative result from [4], it is unlikely that a result as strong as [28] could be shown for adaptive simulatability of Offline Yao.<sup>10</sup> However, as we will see, relaxing the security requirement to adaptive indistinguishability offers some wiggle room. The key to exploiting this, as we explain next, is to discard the simulated garbling mode (Sim) in the hybrids altogether, which allows us to argue security *without* having to program the output map.

*Bypassing the simulated mode.* A standard way to show that a simulation-based definition implies an indistinguishability-based definition (e.g., think of semantic security and IND-CPA) is to use a two-step hybrid argument where the simulated game acts as an intermediary between the “left” and “right” indistinguishability games. If one attempts to use this approach in our context and use the result from [28] to argue adaptive indistinguishability of *offline* Yao garbling, we immediately run into the issue with programming the output map. Thus it seems that the necessity to program the output map is tied to the simulated game, and hence to the simulated mode of garbling. The main idea behind our reduction is therefore to avoid the simulated mode and instead only work with the real and input-dependent modes, which do not require programming the output map. Thus in all our hybrids, the output map is simply the honestly-generated output map and therefore can be generated in the offline phase itself.

*Our approach.* Our idea is to directly replace – gate by gate – the honest garbling table of gates in  $C_0$  ( $\text{Real}_0$ ) with that of gates in  $C_1$  ( $\text{Real}_1$ ). Since the luxury of programming the output map is no longer available, it is crucial to ensure that the evaluation of the garbled circuit in all intermediate hybrids is correct at all times: even though  $C_0(x_0) = C_1(x_1)$  holds (by definition) there is no guarantee that the output of the internal gates of  $C_0$  and  $C_1$  match. An error propagated as a result of one circuit influencing the computation of another may render the hybrids trivially distinguishable to the adversary (via evaluation of

---

<sup>10</sup> Since pseudo-random generators (of arbitrary stretch) exist in  $\mathbf{NC}^1$  [16,24], the result in [4] rules out reductions with polynomial loss for *offline* Yao. This is in stark contrast to the aforementioned positive result from [28] for *online* Yao for  $\mathbf{NC}^1$  circuits.

the garbling). To this end, we employ the input-dependent modes for  $(C_0, x_0)$  and  $(C_1, x_1)$  (resp.,  $\text{Input}_0$  and  $\text{Input}_1$ ). In more details, in all our hybrids, we ensure that a gate in  $\text{Real}_0$  mode is *never* adjacent to another gate in the  $\text{Real}_1$  mode. This is accomplished by maintaining a “frontier” of gates in  $\text{Input}_0$  and  $\text{Input}_1$  mode in between the gates in real mode (see Figure 5). This separation of the left ( $\text{Real}_0$  and  $\text{Input}_0$ ) and right ( $\text{Real}_1$  and  $\text{Input}_1$ ) modes guarantees that the computations belonging to the two circuits do not “corrupt” each other. We point out that this is reminiscent of (circuit) simulation strategies adopted in certain works in circuit complexity [20] (see Section 1.3).


The design of our black-gray-red (BGR) pebble game is carried out keeping the above blueprint in mind. Looking ahead, one can think of it as a symmetrised formulation of the BG pebble game. Our proof that a BGR strategy implies a valid sequence of hybrids is mostly similar to that in [28]: we show that if there exists a BGR pebbling strategy of length  $\tau$  that uses  $\sigma$  grayscale (i.e. black or gray) pebbles, then there exists a reduction to adaptive indistinguishability of Offline Yao with a loss in security at most  $O(\tau 2^\sigma)$  (Theorem 6).<sup>11</sup> The bulk of our technical work goes into coming up with pebble-efficient strategies for the BGR pebble game. This task turns out to be considerably more involved than for the BG pebble game (primarily due to the constraints introduced by the additional rules in the BGR game). The best strategy we could come up with exploits the treewidth  $w$  of the circuit, and as a result the number of (grayscale) pebbles used is roughly  $\sigma := w\delta \log(S)$ , where  $S$  is the size of the circuit and  $\delta$  its fan-out. The strategy has a divide-and-conquer flavour and crucially relies on the notion of separators from graph theory [40,11]. In the remainder of the technical overview, we informally present the BGR pebble game and then briefly explain the treewidth-based BGR strategy.

*BGR pebble game.* Let  $g$  denote the location of a gate in  $G := \Phi(C_0) = \Phi(C_1)$ , the directed acyclic graph (DAG) underlying the circuits, and let  $g_0$  (resp.,  $g_1$ ) denote the corresponding gate in  $C_0$  (resp.,  $C_1$ ). The BGR pebble game (formally defined in Definition 12), as its name suggests, uses three types of pebbles: black, gray and red. In order to translate a BGR pebble configuration  $\mathcal{P}$  to the garbling  $\tilde{C}$ , the simulator in hybrid  $\text{H}_{\mathcal{P}}$  does the following for all internal gates  $g$ :

- if  $g$  carries no pebble in  $\mathcal{P}$ , then its garbling table in  $\tilde{C}$  will be the honest garbling table of  $g_0$ ,
- if  $g$  carries a black pebble then the honest garbling table will be replaced by that of constant-0 or constant-1 gate *depending* on the output value of  $g_0$  when  $C_0$  is run on  $x_0$ ,
- if  $g$  carries a gray pebble, then the simulation is the same as in previous case except that the garbling *depends* on the output value of  $g_1$  when  $C_1$  is run on  $x_1$ ,
- if  $g$  carries a red pebble, then its garbling table in  $\tilde{C}$  will be the honest garbling table of  $g_1$ .

<sup>11</sup> We use the piecewise-guessing framework [25] instead of a direct argument as in [28].

The input is then garbled as follows: For the  $i$ -th input gate, if this gate carries no pebble or a black pebble, then the  $i$ -th key in  $\tilde{x}$  is the key corresponding to the  $i$ th bit of  $x_0$ , otherwise it is the key corresponding to the  $i$ th bit of  $x_1$ . (The pebbles on the output gates are simply ignored.) The four modes of simulation above are real and input-dependent modes for the left and right game respectively or, in short,  $\text{Real}_0$ ,  $\text{Input}_0$ ,  $\text{Input}_1$  and  $\text{Real}_1$  respectively (see Table 3.(b)). Note that the semantics of gates that carry no pebble or a black pebble is the same as in the BG pebble game (if one sets  $(C_0, x_0) = (C, x)$ ), but a gray pebble is now interpreted differently. A BGR pebbling strategy starts off with a configuration with all gates empty (i.e., honest garbling of  $C_0$ ) but the goal is now to pebble them all red (i.e., honest garbling of  $C_1$ ). Thus the extreme hybrids correspond to the left and right games in the adaptive indistinguishability game. The pebbling rules, listed below (see Figure 2), are designed keeping the above discussion in mind and so that indistinguishability of neighbouring hybrids can be argued (Lemma 1):



**Fig. 2.** Rules for BGR pebble game.

1. a black pebble can be placed on or removed from a gate  $g$  if and only if  $g$ 's predecessor gates are pebbled black; and
2. a black pebble on a gate  $g$  can be *swapped* with a gray pebble if  $g$ 's successor gates are pebbled, either black or gray; and
3. a gray pebble on a gate  $g$  can be *swapped* with a red pebble if  $g$ 's predecessor gates are pebbled gray.

Note that the dynamic between no pebbles and black pebbles is similar to the dynamic between red and gray pebbles (hence the reason we consider it to be a symmetric version of the BG pebble game). Since the output values of gates which carry a black or gray pebble in  $\mathcal{P}$  need to be known to carry out the simulation of  $H_{\mathcal{P}}$ , the goal here is to minimise the number of such “grayscale” pebbles.

*Treewidth, separators and BGR pebbling strategies.* Compared to the BG pebble game, pebble-efficient strategies for the BGR pebble game are harder to come by.

(This is not surprising, in hindsight, given the negative result [4].) In particular, the generic pebbling strategies used in [28] no longer work without incurring a blow-up in the number of pebbles employed.<sup>12</sup> Below we briefly explain our treewidth-based strategy, the best (generic) strategy we could come up with.

Crucial to our strategy is the notion of *separators*. Informally, a separator for a circuit  $C$  of size  $S$  is a subset of gates  $\mathcal{S}$  such that removing  $\mathcal{S}$  (and the edges incident on it) from  $C$  partitions  $C$  into sub-circuits of “comparable” size. Slightly more formally,  $\mathcal{S}$  partitions  $C$  into sub-circuits  $C_1, \dots, C_p$  such that for every sub-circuit  $C_i$ ,  $|C_i| \leq 2S/3$  (say). In a classical result from graph theory, it is shown that the size of separator of a graph (and therefore a circuit) is at most its treewidth [40,11]. Since treewidth is a monotonous property – i.e., removing wires or gates from  $C$  can only decrease its treewidth – the process of decomposition into sub-circuits using separators can be recursively carried out further (using a different separator each time) till one ends up with constant-size sub-circuits. Such a recursive decomposition is also carried out in the simulation in [20, Theorem 2] (also see [9,35]).

Our pebbling strategy exploits this recursive decomposition to minimise the number of grayscale pebbles used. To this end, the pebbling strategy maintains *long-term* grayscale pebbles only at the separators. These pebbles help reduce the task at hand to that of (recursively) pebbling the resulting sub-circuits, one at a time *reusing* pebbles in that process. Therefore, our pebbling strategy can be recursively described as follows:

- place grayscale pebbles at the separator  $\mathcal{S}$  of  $G$ ,
- recursively, one at a time, place red pebbles on each subcircuit  $C_i$ ,
- replace the grayscale pebbles on  $\mathcal{S}$  with red pebbles.

Since the depth of the recursion is bounded by  $O(\log S)$  (thanks to the property of the separator), the hope would be that the number of grayscale pebbles maintained overall does not blow up. We show that this is indeed the case as our main technical contribution (Theorem 5).

**Theorem (main).** Any circuit  $C$  of size  $S$ , fan-out  $\delta$  and treewidth  $w$  can be BGR pebbled using  $O(\delta w \log(S))$  grayscale pebbles.

Translating the above divide-and-conquer approach into an actual pebbling strategy (Section 4) turns out to be tricky due to the intricate nature of the BGR pebbling rules. Therefore, before presenting the strategy for arbitrary graphs (Section 4.2), we present a BGR pebbling strategy for path graphs ( $P$ ) as a warm-up (Section 4.1) – since paths have treewidth 1, the latter can be viewed as a special case of the former. We refer the readers to Section 4 for the details.

*Epilogue.* It is instructive to review the above pebbling strategy in terms of the actual simulation. The (garbling tables of) circuit  $C_0$  is being progressively,

<sup>12</sup> The width-based BG strategy from [28,23] can be modified to obtain a comparable BGR strategy for *levelled* circuits. However, the resulting security bounds do not yield any advantage over simply guessing the input (which we want to avoid).

piece by piece, replaced by the (garbling tables of) circuit  $C_1$  as dictated by the recursion, with the bulk of the replacement happening at the base of the recursion. It is exactly those long-term grayscale pebbles placed on the separators which act as the frontier between the pieces of  $C_0$  and  $C_1$ . This ensures that computations of the two circuits are insulated from each other (see Figure 5).

*Remark 1.* We remark that our result on the BGR pebbling complexity can also be used to prove tighter security bounds for simulatability of Online Yao for circuit classes where the treewidth is smaller as the width. This is true since any BGR sequence with complexity  $\sigma$  implies a BG sequence with complexity at most  $\sigma$ : simply consider the BG sequence obtained from a BGR sequence by substituting all the red pebbles with a gray pebble, and note that for BG pebbling only the number of black pebbles is counted.

### 1.3 Related Work

**Garbling.** Most of the works on garbling that are relevant to our paper have already been discussed in Section 1. In addition to them, [5,26,21,31] pertain to adaptively-secure garbling and are also worth pointing out. Besides, there are several constructions of garbling schemes which aim to exploit structured primitives to improve upon other aspects of garbling like, e.g., online complexity (e.g., [12,3]). We refer the readers to [7] for an excellent exposition on both the historical and technical aspects of Yao’s garbling.

**Treewidth, Separators and Computational Complexity.** Treewidth [40] has its roots in algorithmic graph theory. Many hard graph-theoretical problems become tractable when one restricts to graphs of bounded treewidth. In some cases, this leads to even  $\mathbf{NC}$  algorithms for problems which are otherwise known to be  $\mathbf{NP}$ -complete (e.g., [9]). More often than not, this is because bounding the treewidth leads to divide-and-conquer algorithms, sometimes via separators (see [10] for an instructive survey). Unsurprisingly, this also has several consequences in circuit complexity (e.g., [1,2,36,34]), and perhaps the most relevant to our work are [29,20]. It was shown in [29] that circuits with constant treewidth can be simulated in  $\mathbf{NC}^1$ ; [20] extended this result by using separators to show that circuits of size  $S$  and treewidth  $w$  can be simulated in depth  $w \log(S)$ . Both these results can be regarded to be a generalisation of Spira’s theorem that Boolean formulae can be simulated by  $\mathbf{NC}^1$  circuits [42].

**Computing Separators.** The problem of computing (balanced) separators in its full generality is  $\mathbf{NP}$ -complete [19,37]; finding *minimal* separators is  $\mathbf{NP}$ -hard [14]. The parameterised complexity of this problem is well-studied and it is  $\mathbf{W}[1]$ -hard (in both the size of the separator and size of the components) [37]. However, when restricted to *constant-degree* graphs, the problem becomes fixed-parameter tractable [37]. For results pertaining to approximation algorithms for computing (balanced) separators, see [17,18,19].

## 2 Preliminaries

### 2.1 Notation

By  $[a, b]$ , we denote the sequence of integers  $a, a+1, \dots, b-1, b$ . All our logarithms are base two.

*Notation for graphs.* For a graph  $G = (\mathcal{V}, \mathcal{E})$  and a subset  $\mathcal{S} \subseteq \mathcal{V}$ ,  $G_{|\mathcal{S}}$  denotes the subgraph of  $G$  obtained by *restricting* to the set of vertices in  $\mathcal{S}$ . That is  $G_{|\mathcal{S}} = (\mathcal{S}, \mathcal{E}_{|\mathcal{S}})$  where  $\mathcal{E}_{|\mathcal{S}} := \{(u, v) \in \mathcal{E} : u, v \in \mathcal{S}\}$ . For a directed graph  $G$ , a vertex  $u \in \mathcal{V}$  is a predecessor (resp., successor) of another vertex  $v \in \mathcal{V}$  if  $(u, v) \in \mathcal{E}$  (resp.,  $(v, u) \in \mathcal{E}$ ). We say that  $u$  is adjacent to  $v$  if it is either a predecessor or a successor of  $v$ . These definitions can be naturally extended to a set of vertices  $\mathcal{S}$  by taking a union over all the vertices in  $\mathcal{S}$ . The degree  $\delta$  of a vertex is the number of vertices adjacent to it. The in-degree  $\delta_{in}$  (resp., out-degree  $\delta_{out}$ ) of a vertex is its number of predecessors (resp., successors). The degree, in-degree and out-degree of a graph is obtained by taking the corresponding maximum over all its vertices.

*Notation for circuits.* We consider Boolean circuits with explicit input and output gates, associated with the input and output wires respectively. For a circuit  $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$  with  $S$  gates (including the  $n$  input and  $\ell$  output gates) and  $W$  wires of which  $n$  (resp.,  $\ell$ ) are input (resp., output) wires, we denote the DAG that represents the topology of the circuit  $C$  by  $\Phi(C)$ . That is,  $\Phi(C)$  is a graph with  $\mathcal{V} = [1, S]$  obtained by:

1. assigning the input (resp., output) gates to the vertices  $[1, n]$  (resp.,  $[S - \ell + 1, S]$ ),
2. assigning the internal gates to the vertices  $[n + 1, S - \ell]$ , and
3. assigning the wires of the circuit to the edges.

The wires are assigned an index from  $[1, W]$ , with the input (resp., output) wires indexed from  $[1, n]$  (resp.,  $[W - \ell + 1, W]$ ). An internal gate of a circuit is represented by a four-tuple  $(g, u, v, w)$  where  $g : \{0, 1\}^2 \rightarrow \{0, 1\}$  denotes the predicate implemented, and  $u, v$  and  $w$  denote the left input, right input and output wires, respectively. We use  $V_0(w)$  (resp.,  $V_1(w)$ ) as a short-hand for  $V_0(C_0, x_0, w)$  (resp.,  $V_1(C_1, x_1, w)$ ), the function that returns the *value* of the wire  $w$  when the circuit  $C_0$  (resp.,  $C_1$ ) is evaluated on the input  $x_0$  (resp.,  $x_1$ ).

### 2.2 Garbling

The formal definition of syntax and security of garbling schemes is originally from [7]. Our definitions are taken mostly from [27].

**Definition 1 (Indistinguishability).** *A function  $\epsilon : \mathbb{N} \rightarrow [0, 1]$  is negligible if for every polynomial  $p(\lambda)$  there exists an  $\lambda_0 \in \mathbb{N}$  such that  $\epsilon(\lambda) \leq 1/p(\lambda)$  for all  $\lambda \geq \lambda_0$ . Let  $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$  and  $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$  be two distribution ensembles indexed by a security parameter  $\lambda$ . We say that  $X_\lambda$  and  $Y_\lambda$  are  $(\epsilon(\lambda), T(\lambda))$ -indistinguishable if for any adversary  $A$  of size at most  $T(\lambda)$ ,*

$$|\Pr_{a \leftarrow X_\lambda}[A(a) = 1] - \Pr_{a \leftarrow Y_\lambda}[A(a) = 1]| \leq \epsilon(\lambda).$$

**Definition 2 (Garbling Scheme).** A garbling scheme  $\mathbf{GC}$  is a tuple of PPT algorithms  $(\text{GCircuit}, \text{GInput}, \text{GEval})$  with syntax and semantics defined as follows.

$(\tilde{C}, K) \leftarrow \text{GCircuit}(1^\lambda, C)$ . On inputs a security parameter  $\lambda$  and a circuit  $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ , the garble-circuit algorithm  $\text{GCircuit}$  outputs the garbled circuit  $\tilde{C}$  and key  $K$ .

$\tilde{x} \leftarrow \text{GInput}(K, x)$ . On input an input  $x \in \{0, 1\}^n$  and key  $K$ , the garble-input algorithm  $\text{GInput}$  outputs  $\tilde{x}$ .

$y = \text{GEval}(\tilde{C}, \tilde{x})$ . On input a garbled circuit  $\tilde{C}$  and a garbled input  $\tilde{x}$ , the evaluate algorithm  $\text{GEval}$  outputs  $y \in \{0, 1\}^\ell$ .

*Correctness.* There is a negligible function  $\epsilon = \epsilon(\lambda)$  such that for any  $\lambda \in \mathbb{N}$ , any circuit  $C$  and input  $x$  it holds that

$$\Pr \left[ C(x) = \text{GEval}(\tilde{C}, \tilde{x}) \right] = 1 - \epsilon(\lambda),$$

where  $(\tilde{C}, K) \leftarrow \text{GCircuit}(1^\lambda, C)$ ,  $\tilde{x} \leftarrow \text{GInput}(K, x)$ .

**Definition 3 (Adaptive Simulatability).** A garbling scheme  $\mathbf{GC}$  is  $(\epsilon, T)$ -adaptively-simulatable for a class of circuits  $\mathcal{C}$ , if there exists a PPT time simulator  $S = (\text{SCircuit}, \text{SInput})$  such that, for any probabilistic adversary  $A$  of size  $T = T(\lambda)$ ,

$$\left| \Pr [F_{A, \mathbf{GC}, S}(1^\lambda, 0) = 1] - \Pr [F_{A, \mathbf{GC}, S}(1^\lambda, 1) = 1] \right| \leq \epsilon(\lambda).$$

where the experiment  $F_{A, \mathbf{GC}, S}(1^\lambda, b)$  is defined as follows:

1. The adversary  $A$  specifies  $C \in \mathcal{C}$  and gets  $\tilde{C}$  created as follows:
  - if  $b = 0$ :  $(\tilde{C}, K) \leftarrow \text{GCircuit}(1^\lambda, C)$ ,
  - if  $b = 1$ :  $(\tilde{C}, z) \leftarrow \text{SCircuit}(1^\lambda, \Phi(C))$ .
2. The adversary  $A$  specifies  $x$  and gets  $\tilde{x}$  created as follows:
  - if  $b = 0$ ,  $\tilde{x} \leftarrow \text{GInput}(K, x)$ ,
  - if  $b = 1$ ,  $\tilde{x} \leftarrow \text{SInput}(C(x), z)$ .
3. Finally, the adversary outputs a bit  $b'$ , which is the output of the experiment.

**Definition 4 (Adaptive Indistinguishability).** A garbling scheme  $\mathbf{GC}$  is  $(\epsilon, T)$ -adaptively-indistinguishable for a class of circuits  $\mathcal{C}$ , if for any probabilistic adversary  $A$  of size  $T = T(\lambda)$ ,

$$\left| \Pr [G_{A, \mathbf{GC}}(1^\lambda, 0) = 1] - \Pr [G_{A, \mathbf{GC}}(1^\lambda, 1) = 1] \right| \leq \epsilon(\lambda).$$

where the experiment  $G_{A, \mathbf{GC}, S}(1^\lambda, b)$  is defined as follows:

1.  $A$  selects two circuits  $C_0, C_1 \in \mathcal{C}$  such that  $\Phi(C_0) = \Phi(C_1)$  and receives  $\tilde{C}_b$  where  $(\tilde{C}_b, K) \leftarrow \text{GCircuit}(1^\lambda, C_b)$ .
2.  $A$  specifies  $x_0, x_1$  such that  $C_0(x_0) = C_1(x_1)$  and receives  $\tilde{x}_b \leftarrow \text{GInput}(K, x_b)$ .
3. Finally,  $A$  outputs a bit  $b'$ , which is the output of the experiment.

In the selective counterparts of Definitions 3 and 4, the adversary has to select (along with the circuit) the input also in the first step. For self-containment, we provide the definition of selective indistinguishability in Definition 5.

**Definition 5 (Selective Indistinguishability).** *A garbling scheme  $\mathbf{GC}$  is  $(\epsilon, T)$ -selective-indistinguishable for a class of circuits  $\mathcal{C}$ , if for any probabilistic adversary  $\mathbf{A}$  of size  $T = T(\lambda)$ ,*

$$|\Pr [\mathbf{H}_{\mathbf{A}, \mathbf{GC}}(1^\lambda, 0) = 1] - \Pr [\mathbf{H}_{\mathbf{A}, \mathbf{GC}}(1^\lambda, 1) = 1]| \leq \epsilon(\lambda).$$

where the experiment  $\mathbf{G}_{\mathbf{A}, \mathbf{GC}, S}(1^\lambda, b)$  is defined as follows:

1.  $\mathbf{A}$  selects two circuits  $\mathbf{C}_0, \mathbf{C}_1 \in \mathcal{C}$  and two inputs  $x_0, x_1$  such that  $\Phi(\mathbf{C}_0) = \Phi(\mathbf{C}_1)$  and  $\mathbf{C}_0(x_0) = \mathbf{C}_1(x_1)$ . It receives  $(\tilde{\mathbf{C}}_b, \tilde{x}_b)$  where  $(\tilde{\mathbf{C}}_b, K) \leftarrow \mathbf{GCircuit}(1^\lambda, \mathbf{C}_b)$  and  $\tilde{x}_b \leftarrow \mathbf{GInput}(K, x_b)$
2.  $\mathbf{A}$  outputs a bit  $b'$ , which is the output of the experiment.

*Remark 2.* A few remarks concerning Definitions 3 to 5 are in order:

1. We call the experiments corresponding to  $b = 0$  and  $b = 1$  in Definition 3 “real” and “simulated” experiments, respectively. We call the experiments corresponding to  $b = 0$  and  $b = 1$  in Definitions 4 and 5 the “left” and “right” experiments, respectively.
2. When the context is clear, we use the simpler notation  $\mathbf{F}^0$  and  $\mathbf{F}^1$  to denote the experiments  $\mathbf{F}_{\mathbf{A}, \mathbf{GC}, S}(1^\lambda, 0)$  and  $\mathbf{F}_{\mathbf{A}, \mathbf{GC}, S}(1^\lambda, 1)$ , respectively. Similarly, we use  $\mathbf{G}^0, \mathbf{G}^1, \mathbf{H}^0$  and  $\mathbf{H}^1$  for the experiments in Definitions 4 and 5
3. We use  $T_{\mathbf{G}} = T_{\mathbf{G}}(\lambda)$  (resp.,  $T_{\mathbf{H}} = T_{\mathbf{H}}(\lambda)$ ) to denote the time taken to run experiment  $\mathbf{G}$  (resp.,  $\mathbf{H}$ ).

*Offline Yao.* In Algorithm 1 we describe Yao’s original garbling scheme  $\mathbf{YGC}_{\mathbf{SKE}}$  for Boolean circuits with fan-in 2, where the output map is sent along with the garbled circuit in the offline phase. Recall that in the variant from [28], the output map is sent along with the *garbled input* in the online phase. In addition to satisfying the standard notion of security for SKE (IND-CPA), the SKE needs to satisfy the following property for correctness of the garbling schemes to hold.

**Definition 6 (Special Correctness [28]).** *We say that an SKE  $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$  with message space  $\mathcal{M}$  satisfies special correctness if for every security parameter  $\lambda$ , every key  $k \leftarrow \mathbf{Gen}(1^\lambda)$ , every message  $m \in \mathcal{M}$ , and encryption  $c \leftarrow \mathbf{Enc}_k(m)$ ,  $\mathbf{Dec}_{k'}(c) = \perp$  holds for all  $k' \neq k$  with overwhelming probability.*

### 2.3 Pebble Games

In this section, we formally define the pebble games that are relevant to our discussion.



---

**Algorithm 1** Offline Yao scheme  $\mathbf{YGC}_{\mathbf{SKE}}$  based on a symmetric-key encryption scheme  $\mathbf{SKE} := (\mathbf{K}, \mathbf{E}, \mathbf{D})$ .

---

```

1: procedure GCircuit( $1^\lambda, \mathbf{C}$ )
  input:
    1. Security parameter  $\lambda$  in unary
    2. Circuit  $\mathbf{C} : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$  with  $S$  gates and  $W$  wires
  2: for  $j \in [1, W]$  and  $\gamma \in \{0, 1\}$  do  $k_j^\gamma \leftarrow \mathbf{K}(1^\lambda)$  end for ▷ Sample wire keys
  3: Set  $\mu = \{(k_j^0 \rightarrow 0, k_j^1 \rightarrow 1)\}_{j \in [W-\ell+1, W]}$  ▷ Set output map
  4: for  $j \in [n+1, S-\ell]$  do ▷ Garble circuit
    5: Let  $(g_j, u, v, w)$  denote the  $j$ -th gate in  $\mathbf{C}$ 
    6:  $\tilde{g}_j \leftarrow \mathbf{Real}(g_j, \{k_u^\gamma, k_v^\gamma, k_w^\gamma\}_{\gamma \in \{0,1\}})$  ▷ Real defined in Table 3
  7: end for
  8: return  $(\tilde{\mathbf{C}}, K = (\mathbf{k}, \mu))$ , where  $\tilde{\mathbf{C}} := \{\tilde{g}_j\}_{j \in [n+1, S-\ell]}$  and  $\mathbf{k} := \{(k_j^0, k_j^1)\}_{j \in [1, n]}$ 
9: end procedure

10: procedure GInput( $K, x$ )
  input:
    1. Garbling key  $K$  parsed as  $(\mathbf{k}, \mu)$  as in Lines 3 and 8
    2. Input  $x \in \{0, 1\}^n$  with bit decomposition  $\{x_1, \dots, x_n\}$ 
  11: Set  $\mathbf{k}_x := \{k_j^{x_j}\}_{j \in [1, n]}$  ▷ Select input keys from  $\mathbf{k}$ 
  12: return  $\tilde{x} := (\mathbf{k}_x, \mu)$ 
13: end procedure

14: procedure GEval( $\tilde{\mathbf{C}}, \tilde{x}$ )
  input:
    1. Garbled circuit  $\tilde{\mathbf{C}}$  parsed as  $\{\tilde{g}_j\}_{j \in [n+1, S-\ell]}$ 
    2. Garbled input  $\tilde{x}$  parsed as  $(\mathbf{k}_x, \mu)$ 
  15: Parse  $\mathbf{k}_x$  as  $\{k_1, \dots, k_n\}$ 
  16: for  $j \in [n+1, S-\ell]$  do ▷ Decode circuit
    17: Let  $(u, v, w)$  denote the wires of  $g_j$  ▷ The topology is assumed to be public
    18: Parse  $\tilde{g}_j$  as  $(\{c_1, \dots, c_4\})$  ▷ Parse the gate table
    19: for  $l \in [1, 4]$  do ▷ Decrypt each double ciphertext till successful
      20: Let  $m := \mathbf{D}_{k_u}(\mathbf{D}_{k_v}(c_l))$ 
      21: if  $m \neq \perp$  then set  $k_w := m$  ▷ Use special correctness: Definition 6
    22: end for
  23: end for
  24: Parse  $\mu$  as  $\{(k_j^0 \rightarrow 0, k_j^1 \rightarrow 1)\}_{j \in [W-\ell+1, W]}$ 
  25: for  $j \in [W-\ell+1, W]$  do ▷ Decode output
    26: Set  $y_j = 0$  if  $k_j = k_j^0$ ; else set  $y_j = 1$ 
  27: end for
  28: return  $y$ 
29: end procedure

```

---

**Definition 7 (Reversible black pebble game [8,38]).** Consider a DAG  $G = (\mathcal{V}, \mathcal{E})$  with  $\mathcal{V} = [1, S]$  and let  $\mathcal{X}_B = \{\perp, B\}$ . Let  $\mathcal{T} \subseteq \mathcal{V}$  denote the sinks of  $G$ . Consider a sequence  $\mathcal{P} := (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  of pebble configurations for  $G$ , where  $\mathcal{P}_i \in \mathcal{X}_B^\mathcal{V}$  for all  $i \in [0, \tau]$ . We call such a sequence a reversible black pebbling strategy<sup>13</sup> for  $G$  if (i) every vertex is empty in the initial configuration (i.e.,  $\mathcal{P}_0 = (\perp, \dots, \perp)$ ), (ii) every sink is black-pebbled in the final configuration (i.e.,  $\mathcal{P}_\tau(j) = B$  for all  $j \in \mathcal{T}$ ), and (iii) every configuration is obtained by applying the following rule to its preceding configuration:

1.  $\perp \leftrightarrow B$ : a black pebble can be placed on or removed from a vertex if its predecessors are black-pebbled. In particular, a black pebble can be placed on or removed from a source vertex at any time. More formally,  $\exists! j^* \in \mathcal{V}$  such that
  - $(\mathcal{P}_{i+1}(j^*) = B \text{ and } \mathcal{P}_i(j^*) = \perp)$  or  $(\mathcal{P}_{i+1}(j^*) = \perp \text{ and } \mathcal{P}_i(j^*) = B)$ ,
  - $\forall j \in \text{pre}_G(j^*) : \mathcal{P}_i(j) = B$ , and
  - $\forall j \in \mathcal{V} \setminus \{j^*\} : \mathcal{P}_{i+1}(j) = \mathcal{P}_i(j)$ .

The space-complexity of a reversible black pebbling strategy  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  for a DAG  $G$  is defined as the maximum number of black pebbles used at any point in the strategy:

$$\sigma_G(\mathcal{P}) := \max_{i \in [0, \tau]} |\{j \in [1, S] : \mathcal{P}_i(j) = B\}|.$$

**Definition 8.** If  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  is a black pebbling strategy of space-complexity  $\sigma$  for a graph  $G$ , we say that  $\mathcal{P}$  is a  $(\sigma, \tau)$ -strategy for  $G$ . We say that a class of graphs  $\mathcal{G}$  has a  $(\sigma, \tau)$ -strategy if every graph  $G \in \mathcal{G}$  has a  $(\sigma, \tau)$ -strategy. Similarly, we say that a class of circuits  $\mathcal{C}$  has a  $(\sigma, \tau)$ -strategy if for every circuit  $C \in \mathcal{C}$ ,  $\Phi(C)$  has a  $(\sigma, \tau)$ -strategy.

*Remark 3.* Similar definitions apply to the rest of the pebble games considered in the paper.

**Definition 9 (Black-gray (BG) pebble game [23,28]).** Consider a DAG  $G = (\mathcal{V}, \mathcal{E})$  with  $\mathcal{V} = [1, S]$  and let  $\mathcal{X}_{BG} = \{\perp, B, G\}$  denote the set of colours of the pebbles. Consider a sequence  $\mathcal{P} := (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  of pebble configurations for  $G$ , where  $\mathcal{P}_i \in \mathcal{X}_{BG}^\mathcal{V}$  for all  $i \in [0, \tau]$ . We call such a sequence a black-gray pebbling strategy for  $G$  if (i) every vertex is empty in the initial configuration (i.e.,  $\mathcal{P}_0 = (\perp, \dots, \perp)$ ), (ii) every vertex is gray-pebbled in the final configuration (i.e.,  $\mathcal{P}_\tau = (G, \dots, G)$ ) and (iii) every configuration is obtained by applying one of the following rules to its preceding configuration (see Figure 1):

1.  $\perp \leftrightarrow B$ : a black pebble can be placed on or removed from a vertex if its predecessors are black-pebbled. In particular, a black pebble can be placed on or removed from a source vertex at any time. More formally,  $\exists! j^* \in \mathcal{V}$  such that

<sup>13</sup> To be precise, such a pebbling strategy is said to be *persistent* [38] since the final configuration consists of the sinks pebbled. In this paper, we only deal with persistent strategies.

- $(\mathcal{P}_{i+1}(j^*) = B \text{ and } \mathcal{P}_i(j^*) = \perp) \text{ or } (\mathcal{P}_{i+1}(j^*) = \perp \text{ and } \mathcal{P}_i(j^*) = B),$
  - $\forall j \in \text{pre}_G(j^*) : \mathcal{P}_i(j) = B, \text{ and}$
  - $\forall j \in \mathcal{V} \setminus \{j^*\} : \mathcal{P}_{i+1}(j) = \mathcal{P}_i(j).$
2.  $B \mapsto G$ : a black pebble on a vertex  $v \in \mathcal{V}$  can be replaced with a gray pebble if  $v$ 's successors are pebbled (either black or gray). In particular, a black pebble on a sink can be replaced by a gray pebble at any time. More formally,  $\exists! j^* \in \mathcal{V}$  such that
- $\mathcal{P}_{i+1}(j^*) = G \text{ and } \mathcal{P}_i(j^*) = B,$
  - $\forall j \in \text{suc}_G(j^*) : \mathcal{P}_i(j) \in \{B, G\}, \text{ and}$
  - $\forall j \in \mathcal{V} \setminus \{j^*\} : \mathcal{P}_{i+1}(j) = \mathcal{P}_i(j).$

The space-complexity of a BG pebbling strategy  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  for a DAG  $G$  is defined as the maximum number of black pebbles used at any point in the strategy:

$$\sigma_G(\mathcal{P}) := \max_{i \in [0, \tau]} |\{j \in [1, S] : \mathcal{P}_i(j) = B\}|.$$

*Remark 4.* The rule to place or remove a black pebble in Definition 9 (Rule 9.1) is the same as in Definition 7 (Rule 7.1). Therefore the BG pebble game can be thought of as an extension of the RB pebble game (with a different goal).

## 2.4 Graph theory

We recall the definition of treewidth and graph separators, and then state a crucial theorem connecting them, which will be exploited in our pebbling strategy. We emphasise that understanding the definition of treewidth is *not* essential to understanding our pebbling strategies: it is the notion of separators, along with Theorem 1, which is key.

**Definition 10** ([40,11]). A tree decomposition of a graph  $G = (\mathcal{V}, \mathcal{E})$  is a tree,  $T$ , with nodes  $\mathcal{X}_1, \dots, \mathcal{X}_p$ , where each  $\mathcal{X}_i \subseteq \mathcal{V}$ , satisfying the following properties:

1. Each graph vertex is contained in at least one tree node (i.e.,  $\cup_{i \in [1, p]} \mathcal{X}_i = \mathcal{V}$ ).
2. For every edge  $(v, w) \in \mathcal{E}$ , there exists a node  $\mathcal{X}_i$  that contains both  $v$  and  $w$ .
3. The tree nodes containing a vertex  $v$  form a connected subtree of  $T$ .

The width of a tree decomposition is the size of its largest node  $\mathcal{X}_i$  minus one. Its treewidth  $w(G)$  is the minimum width among all possible tree decompositions.

**Definition 11** ([40]). For a graph  $G = (\mathcal{V}, \mathcal{E})$ , a set  $\mathcal{S} \subseteq \mathcal{V}$  is said to be a separator if the graph  $G_{|\overline{\mathcal{S}}}$  has at least two components, and each of these components has size at most  $2|\mathcal{V}|/3$ .<sup>14</sup>

**Theorem 1** ([40,11]). A graph  $G$  with treewidth  $w(G)$  has a separator of size at most  $w(G)$ .

<sup>14</sup> To be precise, such a separator is called ‘‘balanced’’ [20]. In this paper, we only consider balanced separators.

### 3 Hybrid Argument and the BGR Pebble Game

In this section, we formally show that black-gray-red (BGR) pebbling strategies lead to security reductions for Offline Yao. We start off in Section 3.1 by formally defining the BGR pebble game and then explain the semantics of its pebbles, described already (albeit informally) in Section 1.2. This enables us to define a hybrid  $H_{\mathcal{P}}$  in terms of a pebble configuration  $\mathcal{P}$  (Algorithm 2). Then, in Section 3.2, we justify the pebble rules by proving that neighbouring pebble configurations can indeed be proved indistinguishable (Lemma 1). Finally, we put these two steps together in Section 3.3 and show that BGR strategies imply adaptive indistinguishability of Offline Yao (Theorem 4) using the piecewise-guessing framework [25]. Since most of the ideas in Sections 3.2 and 3.3 are similar to pre-existing works [28,25], we skip detailed proofs and resort to high-level sketches.

#### 3.1 Pebble Configurations and Hybrids

The BGR pebble game is a symmetric version of the BG pebble game. In addition to the ones in BG pebble game (Rules 9.1 and 9.2), there are additional rules (Rules 12.2 and 12.3) which govern how the red pebbles interact with the gray pebbles. Intuitively speaking, the dynamic between no pebbles and black pebble (Rule 12.1) is similar to the dynamic between red pebbles and gray pebbles (Rule 12.3): see Remark 5. A more formal definition of the game is given next.

**Definition 12 (Black-Gray-Red (BGR) pebble game).** *Consider a DAG  $G = (\mathcal{V}, \mathcal{E})$  with  $\mathcal{V} = [1, S]$  and let  $\mathcal{X}_{BGR} = \{\perp, B, G, R\}$  denote the set of colours of the pebbles. A pebble is called grayscale if it is black or gray. Consider a sequence  $\mathcal{P} := (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  of pebble configurations for  $G$ , where  $\mathcal{P}_i \in \mathcal{X}_{BGR}^{\mathcal{V}}$  for all  $i \in [0, \tau]$ . We call such a sequence a BGR pebbling strategy for  $G$  if (i) every vertex is empty in the initial configuration (i.e.,  $\mathcal{P}_0 = (\perp, \dots, \perp)$ ), (ii) and every vertex is red-pebbled in the final configuration (i.e.,  $\mathcal{P}_\tau = (R, \dots, R)$ ) and (iii) every configuration is obtained by applying one of the following rules to its preceding configuration (see Figure 2):*

1.  $\perp \leftrightarrow B$ : a black pebble can be placed on or removed from a vertex if its predecessors are black-pebbled. In particular, a black pebble can be placed on or removed from a source vertex at any time. More formally,  $\exists! j^* \in \mathcal{V}$  such that
  - $(\mathcal{P}_{i+1}(j^*) = B \text{ and } \mathcal{P}_i(j^*) = \perp)$  or  $(\mathcal{P}_{i+1}(j^*) = \perp \text{ and } \mathcal{P}_i(j^*) = B)$ ,
  - $\forall j \in \text{pre}_G(j^*) : \mathcal{P}_i(j) = B$ , and
  - $\forall j \in \mathcal{V} \setminus \{j^*\} : \mathcal{P}_{i+1}(j) = \mathcal{P}_i(j)$ .
2.  $B \leftrightarrow G$ : a black pebble on a vertex  $v \in \mathcal{V}$  can be swapped with a gray pebble if  $v$ 's successors carry grayscale pebbles (i.e., either black or gray). In particular, a black pebble on a sink vertex can be swapped with a gray pebble at any time. More formally,  $\exists! j^* \in \mathcal{V}$  such that
  - $(\mathcal{P}_{i+1}(j^*) = G \text{ and } \mathcal{P}_i(j^*) = B)$  or  $(\mathcal{P}_{i+1}(j^*) = B \text{ and } \mathcal{P}_i(j^*) = G)$ ,

- $\forall j \in \text{succ}_G(j^*) : \mathcal{P}_i(j) \in \{B, G\}$ , and
  - $\forall j \in \mathcal{V} \setminus \{j^*\} : \mathcal{P}_{i+1}(j) = \mathcal{P}_i(j)$ .
3.  $G \leftrightarrow R$ : a gray pebble can be swapped with a red if its predecessors are gray-pebbled. In particular, a gray pebble on a source vertex can be swapped with a red pebble at any time. More formally,  $\exists! j^* \in \mathcal{V}$  such that
- $(\mathcal{P}_{i+1}(j^*) = G \text{ and } \mathcal{P}_i(j^*) = R) \text{ or } (\mathcal{P}_{i+1}(j^*) = R \text{ and } \mathcal{P}_i(j^*) = G)$ ,
  - $\forall j \in \text{pre}_G(j^*) : \mathcal{P}_i(j) = G$ , and
  - $\forall j \in \mathcal{V} \setminus \{j^*\} : \mathcal{P}_{i+1}(j) = \mathcal{P}_i(j)$ .

The space-complexity of a BGR pebbling strategy  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  for a DAG  $G$  is defined as the maximum number of grayscale pebbles used at any point in the strategy:

$$\sigma_G(\mathcal{P}) := \max_{i \in [0, \tau]} |\{j \in [1, S] : \mathcal{P}_i(j) \in \{B, G\}\}|.$$

*Remark 5.* A few remarks on the BGR pebble game are in order:

1. Note that Rules 12.1 and 12.2 (the  $B \mapsto G$  part) correspond to the rules in the BG pebble game. The end goals in the two games are however different.
2. When restricted to either black and empty (Rule 12.1) or gray and red pebbles (Rule 12.3), the BGR pebble game simplifies to the reversible black pebble game of Bennett [8] defined in Definition 7. This is obvious for the black pebbles since the BGR pebble game is an extension of the BG pebble game which, in turn, is an extension of the reversible black pebble game. To see why this is the case for gray and red pebbles, simply think of vertices with red pebbles as being empty (i.e.,  $R = \perp$ ) and gray pebbles as black pebbles (i.e.,  $G = B$ ), and note that Rule 12.3 is now the same as Rule 12.1. Therefore, if one starts with an all-red (i.e., empty) configuration, the gray pebbles can be placed using reversible pebbling rules. Some of the reversible pebbling strategies will serve as crucial subroutines in the BGR pebbling strategies in the coming sections.
3. When restricted to black and gray pebbles, the BGR pebble game again simplifies to the reversible pebble game played on the graph with the direction of the edges *flipped*. However, we do not make use of this observation.
4. Only black pebbles can be placed on empty vertices. Gray (resp., red) pebbles have to replace black or red (resp., gray) pebbles, respectively.
5. By the pebbling rules, in any strategy a vertex that is empty can *never* end up adjacent to another vertex with red pebble in any BGR pebbling strategy. Moreover, a vertex with gray (resp., black) pebble cannot be a predecessor of a vertex with no (resp. red) pebble; the converse is however possible. These properties will turn out to be important sanity checks in ensuring the validity of BGR pebbling strategies in the later sections. Moreover, they ensure correctness of the simulations they represent.

	Real ( $\perp$ )	Input (B)	Sim (G)
$c_{0,0}$	$E_{k_u^0}(E_{k_v^0}(k_w^{g(0,0)}))$	$E_{k_u^0}(E_{k_v^0}(k_w^{V(w)}))$	$E_{k_u^0}(E_{k_v^0}(k_w^0))$
$c_{0,1}$	$E_{k_u^0}(E_{k_v^1}(k_w^{g(0,1)}))$	$E_{k_u^0}(E_{k_v^1}(k_w^{V(w)}))$	$E_{k_u^0}(E_{k_v^1}(k_w^0))$
$c_{1,0}$	$E_{k_u^1}(E_{k_v^0}(k_w^{g(1,0)}))$	$E_{k_u^1}(E_{k_v^0}(k_w^{V(w)}))$	$E_{k_u^1}(E_{k_v^0}(k_w^0))$
$c_{1,1}$	$E_{k_u^1}(E_{k_v^1}(k_w^{g(1,1)}))$	$E_{k_u^1}(E_{k_v^1}(k_w^{V(w)}))$	$E_{k_u^1}(E_{k_v^1}(k_w^0))$

(a)

	Real <sub>0</sub> ( $\perp$ )	Input <sub>0</sub> (B)	Input <sub>1</sub> (G)	Real <sub>1</sub> (R)
$c_{0,0}$	$E_{k_u^0}(E_{k_v^0}(k_w^{g_0(0,0)}))$	$E_{k_u^0}(E_{k_v^0}(k_w^{V_0(w)}))$	$E_{k_u^0}(E_{k_v^0}(k_w^{V_1(w)}))$	$E_{k_u^0}(E_{k_v^0}(k_w^{g_1(0,0)}))$
$c_{0,1}$	$E_{k_u^0}(E_{k_v^1}(k_w^{g_0(0,1)}))$	$E_{k_u^0}(E_{k_v^1}(k_w^{V_0(w)}))$	$E_{k_u^0}(E_{k_v^1}(k_w^{V_1(w)}))$	$E_{k_u^0}(E_{k_v^1}(k_w^{g_1(0,1)}))$
$c_{1,0}$	$E_{k_u^1}(E_{k_v^0}(k_w^{g_0(1,0)}))$	$E_{k_u^1}(E_{k_v^0}(k_w^{V_0(w)}))$	$E_{k_u^1}(E_{k_v^0}(k_w^{V_1(w)}))$	$E_{k_u^1}(E_{k_v^0}(k_w^{g_1(1,0)}))$
$c_{1,1}$	$E_{k_u^1}(E_{k_v^1}(k_w^{g_0(1,1)}))$	$E_{k_u^1}(E_{k_v^1}(k_w^{V_0(w)}))$	$E_{k_u^1}(E_{k_v^1}(k_w^{V_1(w)}))$	$E_{k_u^1}(E_{k_v^1}(k_w^{g_1(1,1)}))$

(b)

**Table 3.** (a) Garbling modes in [28]. The gate is denoted by  $g$  and the value of its output wire  $w$  when run on input  $x$  is denoted by  $V(w)$ . (b) Garbling modes in our case. The gates  $g_0$  and  $g_1$  are the gates in the same position in the circuits  $C_0$  and  $C_1$ , respectively. The value  $V_0(w)$  (resp.,  $V_1(w)$ ) denotes the bit going over the wire  $w$  in the computation  $C_0(x_0)$  (resp.,  $C_1(x_1)$ ).

**Template for Hybrids.** A pebble configuration  $\mathcal{P} \in \mathcal{X}_{BGR}^V$  is used to encode a selective hybrid  $H_{\mathcal{P}}$ . For an *internal gate*  $v$ , the translation is carried out as described below:

- if  $v$  carries no pebble ( $\perp$ ) in  $\mathcal{P}$  then  $g$  is garbled as in the left game (Real<sub>0</sub>),
- a black pebble (B) on  $v$  indicates that the garbling of  $g$  is input-dependant on  $x_0$  and  $C_0$ , (Input<sub>0</sub>)
- a gray pebble (G) on  $v$  indicates that the garbling of  $g$  is input-dependant on  $x_1$  and  $C_1$  (Input<sub>1</sub>)
- a red pebble (R) on  $v$  indicates  $g$  is garbled as in the right game (Real<sub>1</sub>).

The distributions corresponding to the four garbling modes – Real<sub>0</sub>, Input<sub>0</sub>, Input<sub>1</sub> and Real<sub>1</sub> – are formally defined in Table 3(b). (Note that the semantics of gray pebbles is different from that in the BG pebble game.) This information is sufficient to construct the garbled circuit  $\tilde{C}$ . What remains to complete the description of  $H_{\mathcal{P}}$ , is describing how to generate the input garbling  $\tilde{x}$  and the output map. If an input gate carries no (resp., a red) pebble then the garbling key for  $x_0$  (resp.,  $x_1$ ) is selected in that hybrid. The output map, on the other hand, is simply the default one prescribed in the scheme and therefore the pebbles on the output gates are ignored. We refer the readers to Algorithm 2 for a formal definition of  $H_{\mathcal{P}}$ .

---

**Algorithm 2** Selective hybrids: the adversary commits to the garbling inputs at the beginning of the game (boxed).

---

1: **procedure**  $H_{\mathcal{P}}(\lambda)$

**parameters:**

1. Circuit parameters  $S$  (gates),  $W$  (wires),  $n$  and  $\ell$  (input and output size)
2. Selective adversary  $A$
3. Symmetric encryption scheme  $\mathbf{SKE} := (K, E, D)$

**index:** A BGR pebble configuration  $\mathcal{P} \in \mathcal{X}_{BGR}^{[1,S]}$

2: Receive circuits  $C_0$  and  $C_1$  from  $A$  ▷  $H \leftarrow A$

3: Receive inputs  $x_0, x_1 \in \{0, 1\}^n$  as commitment from  $A$  ▷  $H \leftarrow A$

4: **for**  $j \in [1, W]$  and  $\gamma \in \{0, 1\}$  **do**  $k_j^\gamma \leftarrow K(1^\lambda)$  **end for** ▷ Sample wire keys

5: Set  $\mu = \{(k_j^0 \rightarrow 0, k_j^1 \rightarrow 1)\}_{j \in [W-\ell+1, W]}$  ▷ Set output map

6: **for**  $j \in [n+1, S-\ell]$  **do** ▷ Garble circuit

7:   For  $b \in \{0, 1\}$ , let the  $j$ -th gate in  $C_b$  be  $(g_{b,j}, u, v, w)$

8:   **if**  $\mathcal{P}(n+j) = \perp$  **then**  $\tilde{g}_j \leftarrow \text{Real}_0(g_{0,j}, \{k_u^\gamma, k_v^\gamma, k_w^\gamma\}_{\gamma \in \{0,1\}})$  ▷ No pebble

9:   **else if**  $\mathcal{P}(n+j) = B$  **then**

10:      $\tilde{g}_j \leftarrow \text{Input}_0(g_{0,j}, \{k_u^\gamma, k_v^\gamma\}_{\gamma \in \{0,1\}}, k_w^{V_0(w)})$  ▷ B-pebbled

11:   **else if**  $\mathcal{P}(n+j) = G$  **then**

12:      $\tilde{g}_j \leftarrow \text{Input}_1(g_{1,j}, \{k_u^\gamma, k_v^\gamma\}_{\gamma \in \{0,1\}}, k_w^{V_1(w)})$  ▷ G-pebbled

13:   **else**  $\tilde{g}_j \leftarrow \text{Real}_1(g_{1,j}, \{k_u^\gamma, k_v^\gamma, k_w^\gamma\}_{\gamma \in \{0,1\}})$  ▷ R-pebbled

14:   **end if**

15: **end for**

16: **for**  $j \in [1, n]$  **do** ▷ Garble input

17:   **if**  $\mathcal{P}(j) \in \{\perp, B\}$  **then** Set  $\tilde{x}_j = k_j^{x_0, j}$

18:   **else** Set  $\tilde{x}_j = k_j^{x_1, j}$

19:   **end if**

20: **end for**

21: Send  $(\tilde{C}, \mu, \tilde{x})$  to  $A$  where  $\tilde{C} := \{\tilde{g}_j\}_{j \in [n+1, S-\ell]}$  and  $\tilde{x} := \{\tilde{x}_j\}_{j \in [1, n]}$  ▷  $H \rightarrow A$

22: Receive a bit  $b'$  from  $A$  ▷  $H \leftarrow A$

23: **return**  $b'$

24: **end procedure**

---

*Sequence of hybrids.* A pebbling strategy  $\mathcal{P} = \{\mathcal{P}_0, \dots, \mathcal{P}_\tau\}$  will give rise to a sequence of selective hybrids

$$H^0 = H_{\mathcal{P}_0}, \dots, H_{\mathcal{P}_\tau} = H^1, \quad (1)$$

Note that the extreme games correspond to the left selective experiment  $H^0 = H_{\mathbf{A}, \mathbf{GC}}(1^\lambda, 0)$  (since  $\mathcal{P}_0 = (\perp, \dots, \perp)$ ) and right selective experiment  $H^1 = H_{\mathbf{A}, \mathbf{GC}}(1^\lambda, 1)$  (since  $\mathcal{P}_\tau = (\mathbf{R}, \dots, \mathbf{R})$ ), respectively. The exact pebbling strategy will be discussed later in Section 4. In the next section, we prove the indistinguishability of two neighbouring hybrids in such a sequence.

### 3.2 Indistinguishability of Neighbouring Hybrids

**Lemma 1 (neighbouring indistinguishability).** *Let  $\mathcal{P}$  and  $\mathcal{Q}$  denote two neighbouring configurations in a BGR pebbling strategy. If the underlying encryption scheme **SKE** is  $(\epsilon, T)$ -IND-CPA secure, then  $H_{\mathcal{P}}$  and  $H_{\mathcal{Q}}$  (as defined in Algorithm 2) are  $(3\epsilon, T - T_{\mathbf{H}})$ -indistinguishable, i.e., for any adversary  $\mathbf{A}$  of size at most  $T - T_{\mathbf{H}}$*

$$|\Pr[\langle \mathbf{A}, H_{\mathcal{P}} \rangle = 1] - \Pr[\langle \mathbf{A}, H_{\mathcal{Q}} \rangle = 1]| \leq 3\epsilon.$$

*Proof.* Recall that hybrids correspond to pebble configurations and that two neighbouring hybrids differ by a single pebble. We split the proof into three cases which correspond to the pebbling moves  $\perp \leftrightarrow \mathbf{B}$ ,  $\mathbf{B} \leftrightarrow \mathbf{G}$  and  $\mathbf{R} \leftrightarrow \mathbf{G}$  respectively. The reduction in the first and last cases is similar, and relies on the indistinguishability of the underlying encryption scheme (similar to [28, Lemma 1]). Therefore in the claim below we focus on the first case. In the second case, we argue that the hybrids are identically distributed (similar to [28, Lemma 2]). Moreover, since the proofs are similar to those in [28], we only provide a sketch here.

*Claim (Rule 12.1:  $\perp \leftrightarrow \mathbf{B}$ ).* If the underlying encryption scheme **SKE** is  $(\epsilon, T)$ -IND-CPA secure, and if  $\mathcal{Q}$  is obtained from  $\mathcal{P}$  using Rule 12.1 then the hybrids  $H_{\mathcal{P}}$  and  $H_{\mathcal{Q}}$  are  $(3\epsilon, T - T_{\mathbf{H}})$ -indistinguishable.

*Proof (Sketch).* Let's consider the case where  $\mathcal{Q}$  is obtained from  $\mathcal{P}$  using  $\perp \mapsto \mathbf{B}$  part of Rule 12.1, i.e.  $\mathcal{Q}$  is obtained by placing a black pebble on an empty vertex  $g^*$  in  $\mathcal{P}$  – the complementary case of  $\mathbf{B} \mapsto \perp$  can be argued in a symmetric manner. Recall that, by the BGR pebbling rules, there are two possible cases:

1.  $g^*$  is a source vertex. In this case  $H_{\mathcal{P}}$  and  $H_{\mathcal{Q}}$  are identical experiments.
2.  $g^*$  is an internal vertex – let us denote the corresponding gate in  $\mathbf{C}_0$  by  $(g^*, u^*, v^*, w^*)$ . In this case, the predecessors of  $g^*$  are guaranteed to be pebbled black in the configuration  $\mathcal{P}$ . By the semantics of the pebbles, this means that both the predecessors of  $g^*$  are in  $\mathbf{Input}_0$  mode in hybrids  $H_{\mathcal{P}}$  and  $H_{\mathcal{Q}}$ . This, in turn, means that one of the output keys corresponding to the incoming wire  $u^*$  is not used in the garbling table of that predecessor gate; the same holds for the wire  $w^*$ . Therefore these keys are free and the



reduction can set them as challenge key of the underlying encryption scheme in order to switch the garbling table of  $g^*$  from  $\text{Real}_0$  to  $\text{Input}_0$  in three steps. This corresponds to a black pebble on  $g^*$  and therefore the hybrid  $H_Q$ .  $\square$

*Claim (Rule 12.2:  $B \leftrightarrow G$ ).* If  $Q$  is obtained from  $\mathcal{P}$  using Rule 12.2 then the hybrids  $H_{\mathcal{P}}$  and  $H_Q$  are identically distributed.

*Proof (Sketch).* Let's restrict to the case where  $Q$  is obtained from  $\mathcal{P}$  using  $B \mapsto G$  part of Rule 12.2, i.e.,  $Q$  is obtained by replacing a black pebble on a vertex  $g^*$  in  $\mathcal{P}$  with a gray pebble – as in the previous proof the argument in the second part  $G \mapsto B$  is symmetric. By the BGR pebbling rules, there are two possible cases:

1.  $g^*$  is a sink vertex. Since  $C_0(x_0) = C_1(x_1)$  and  $g^*$  is in input-dependant mode in both  $H_{\mathcal{P}}$  ( $\text{Input}_0$ ) and  $H_Q$  ( $\text{Input}_1$ ) the simulations are identical.
2.  $g^*$  is an internal vertex – let us denote the corresponding gates in  $C_0$  and  $C_1$  by  $(g_0^*, u^*, v^*, w^*)$  and  $(g_1^*, u^*, v^*, w^*)$  respectively. This means that  $g^*$  is in  $\text{Input}_0$  mode in the hybrid  $H_{\mathcal{P}}$  and  $\text{Input}_1$  mode in the hybrid  $H_Q$  (with all the successors in  $\text{Input}_0$  or  $\text{Input}_1$  mode in both the hybrids). In case that the value of the output wire  $w^*$  is the same for both hybrids (i.e.,  $V_0(w^*) = V_1(w^*)$ ) then we are done by the same argument as the case above in Item 1. So let's focus on the case where the values are different. Since both the keys corresponding to  $w^*$  are treated symmetrically in the successor gates, swapping them in the garbling table results in identical distributions and therefore  $H_{\mathcal{P}}$  and  $H_Q$  are also identically distributed.  $\square$

This proves the lemma.  $\square$

*Selective indistinguishability.* Combining Lemma 1 with the semantics of the pebbles (Table 3) yields (via the standard hybrid argument) selective indistinguishability of Offline Yao.

**Theorem 2.** *Suppose that a class of circuits  $\mathcal{C}$  has a  $(\sigma, \tau)$ -BGR pebbling strategy. If the encryption scheme **SKE** is  $(\epsilon, T)$ -secure then  $\mathbf{YGC}_{\mathbf{SKE}}$  is  $(3\tau\epsilon, T - T_H)$ -selectively-indistinguishable for  $\mathcal{C}$ .*

*Remark 6.* Theorem 2 was already proved in a stronger form in [33]: it was shown there that Offline Yao is selectively-simulatable under similar assumptions. Here it only serves as a stepping stone to the main theorem (Theorem 4) to be proved in the next section.

### 3.3 Adaptive Indistinguishability via Piecewise Guessing

Observe that in the hybrid  $H_{\mathcal{P}}$  described in Algorithm 2, the knowledge of the committed garbling inputs  $x_0$  and  $x_1$  (Line 3) is used to compute the output value of gates that carry grayscale pebbles in the configuration  $\mathcal{P}$  (Lines 10 and 12). So, in principle, the simulation of  $H_{\mathcal{P}}$  can be carried out if this information

is available as an “advice”. Moreover, the indistinguishability of two successive hybrids can be shown (Lemma 1) if such advice for *both* the hybrids is available. In case the number of grayscale pebbles is small, the size of this advice could potentially be smaller than the size of garbling inputs  $x_0$  and  $x_1$ . This means that it is possible to apply the piecewise-guessing framework [25]. We explain this in detail next.

*Applying the piecewise-guessing framework.* The main theorem in [25] is stated below in Theorem 3 after having been simplified and tailored for our application to circuit garbling. The result of applying Theorem 3 to Offline Yao is stated in Theorem 4. Furthermore, exploiting the properties of the pebbling strategies we design, we provide an optimised version of Theorem 4 later in Section 4.3 (Theorem 6).

**Theorem 3 (Theorem 2 in [25] tailored to Definitions 4 and 5).** *Let  $G^0, G^1, H^0$  and  $H^1$  be as in Definitions 4 and 5. Furthermore, let  $H^0 = H_{\mathcal{P}_0}, \dots, H_{\mathcal{P}_\tau} = H^1$  be the sequence of hybrids from Equation (1) and suppose that every pebbling configuration  $\mathcal{P}_i$  in the strategy  $\mathcal{P}_0, \dots, \mathcal{P}_\tau$  can be computed in time  $T_p$ . Assume that for each  $i \in [0, \tau - 1]$ , there exists a function  $\alpha_i: \{0, 1\}^* \rightarrow \{0, 1\}^\sigma$  such that the hybrids  $H_{\mathcal{P}_i}$  and  $H_{\mathcal{P}_{i+1}}$  are  $(\epsilon, T)$ -indistinguishable when  $A$  commits to  $\alpha_i(C_0, C_1, x_0, x_1)$  as advice at the beginning of the experiment (instead of  $(x_0, x_1)$ ).<sup>15</sup> Then  $G^0$  and  $G^1$  are  $(\epsilon \cdot \tau \cdot 2^\sigma, T - (T_\sigma + T_p))$ -indistinguishable where  $T_\sigma$  denotes the time to sample a string in  $\{0, 1\}^\sigma$  uniformly at random.*

**Theorem 4.** *Suppose that a class of circuits  $\mathcal{C}$  has a  $(\sigma, \tau)$ -BGR pebbling strategy. If the encryption scheme **SKE** is  $(\epsilon, T)$ -secure then **YGC<sub>SKE</sub>** is  $(\tau 2^\sigma \cdot 3\epsilon, T - (T_H + T_\sigma + T_p))$ -adaptively-indistinguishable for  $\mathcal{C}$ .*

*Proof (Sketch).* As already observed, the advice function  $\alpha_i$  should return the values of the output wires of all those gates that carry grayscale pebbles in  $\mathcal{P}_i$  and  $\mathcal{P}_{i+1}$ . Therefore, in Theorem 3 we set

$$\begin{aligned} \alpha_i(C_0, C_1, x_0, x_1) := & (V_0(w) : (g, u, v, w) \in C_0 \text{ and } \mathcal{P}(g) = \mathbf{B}) \parallel \\ & (V_1(w) : (g, u, v, w) \in C_1 \text{ and } \mathcal{P}(g) = \mathbf{G}) \end{aligned} \quad (2)$$

where  $\mathcal{P} := \mathcal{P}_{i+1}$  if  $\mathcal{P}_{i+1}$  is obtained from  $\mathcal{P}_i$  by adding a grayscale pebble; and  $\mathcal{P} := \mathcal{P}_i$  otherwise.<sup>16</sup> The length of the advice is therefore smaller than in the selective hybrid in case the pebbling complexity of  $G = \Phi(C_0) = \Phi(C_1)$  is smaller than the input length. What remains is to show that indistinguishability of two consecutive hybrids can be shown relying only on  $\alpha_i := \alpha_i(C_0, C_1, x_0, x_1)$ . To

<sup>15</sup> To show this formally, we need to show that for every pair of neighbouring hybrids  $H_i$  and  $H_{i+1}$ , there exist *partially-selective* hybrids  $\hat{H}_{i,0}$  and  $\hat{H}_{i,1}$  and a function  $\alpha_i$  such that (i)  $H_i$  and  $H_{i+1}$  are “partially-selectivised” versions of  $\hat{H}_{i,0}$  and  $\hat{H}_{i,1}$  using the function  $\alpha_i$  and (ii)  $\hat{H}_{i,0}$  and  $\hat{H}_{i,1}$  are indistinguishable.

<sup>16</sup> Recall from the proof of Lemma 1 that for pebbling configurations  $\mathcal{P}_i$  and  $\mathcal{P}_{i+1}$  that differ by a pebbling move  $\mathbf{B} \leftrightarrow \mathbf{G}$ , the corresponding hybrids  $H_{\mathcal{P}_i}$  and  $H_{\mathcal{P}_{i+1}}$  are *identically* distributed.

see this note that the knowledge of the committed garbling inputs  $x_0$  and  $x_1$  (Line 3) is used to compute the output value of gates that carry grayscale pebbles in the configuration  $\mathcal{P}$  (Lines 10 and 12). Since these are already present in the hint, the reduction algorithm can simply extract these values from  $\alpha_i$  and use them instead of explicitly computing  $V_0(\cdot)$  and  $V_1(\cdot)$ . Following the arguments in Lemma 1, we get that if the encryption scheme is  $(\epsilon, T)$ -secure then the experiments  $H_{\mathcal{P}_i}$  and  $H_{\mathcal{P}_{i+1}}$  are  $(3\epsilon, T - T_H)$ -indistinguishable when  $A$  commits to  $\alpha_i$ , and the proof now follows Theorem 3.  $\square$

## 4 BGR Pebbling Strategy

In this section, we describe our main strategy for the BGR pebble game. As a warm up, we describe a BGR pebbling strategy (Section 4.1) for paths on  $S \in \mathbb{N}$  vertices that uses  $O(\log(S))$  grayscale pebbles. It is instructive to go through it as it serves the following two purposes.

1. Firstly, it highlights the fact that coming up with space-efficient BGR pebbling strategies is non-trivial.
2. More importantly, it serves as a stepping stone to the main strategy: paths can be thought of as graphs of (tree)width one and our main strategy (Section 4.2) extends this strategy using the notion of separators.

Finally, we discuss the implications of our pebbling strategy to the security of Offline Yao (Section 4.3).

### 4.1 Warm up: BGR Strategy for Paths

Before describing BGRPath, the BGR strategy for paths, we describe RBPath, Bennett's [8] space-optimal RB pebbling strategy, as it is used as a sub-strategy in BGRPath. RBPath places a black pebble on the sink of a path of length  $S = 2^s$  using  $s + 1$  black pebbles (in  $S^{\log(3)}$  steps). Since the strategy is reversible, by  $\text{RBPath}^{-1}$  we denote the reverse strategy that *removes* the black pebble on the sink. This strategy, thanks to the observation in Remark 5.2, will also be used to both place or remove a gray pebble on an all-red-pebbled graph.

**Space-Optimal RB Pebbling.** For ease of presentation, let's assume that the path has  $S = 2^s$  vertices  $[1, S]$  – the strategy can easily be extended to paths of arbitrary length by splitting it up appropriately. The strategy  $\text{RBPath}_s$  is described recursive in  $s$ . Since the strategy is reversible, we denote its inverse by  $\text{RBPath}_s^{-1}$ .

For the case  $s = 0$ ,  $\text{RBPath}_1$  simply places a pebble on  $v^*$  (Rule 7.1).

For  $s > 0$  and path  $P_{2^s}$ , let's refer to  $P_{2^s|_{[1, 2^{s-1}]}}$  and  $P_{2^s|_{[2^{s-1}+1, 2^s]}}$  as the first and second halves respectively.  $\text{RBPath}_s$  then does the following.

1. Recursively place a black pebble on the first half  
 $(\text{RBPath}_{s-1}(P_{2^s|_{[1, 2^{s-1}]}}))$

2. Recursively place a black pebble on the second half  
( $\text{RBPPath}_{s-1}(P_{2^s|_{[2^{s-1}, 2^s]}})$ )
3. Recursively remove the black pebble on the first half  
( $\text{RBPPath}_{s-1}^{-1}(P_{2^s|_{[1, 2^{s-1}]}})$ )

**Lemma 2.**  $P_{2^s}$  can be RB-pebbled using at most  $\sigma := s + 1$  black pebbles in at most  $\tau := 3^s = 2^{s \log(3)}$  steps.

*Proof.* For the base case of  $s = 0$ ,  $\text{RBPPath}_0$  simply places a pebble on  $v^*$  (Rule 7.1), hence requires 1 pebble and 1 step. Let's assume that  $\text{RBPPath}_{s-1}$  can Dpebble  $P_{2^{s-1}}$  using  $\log(2^{s-1}) + 1 = s$  pebbles in  $3^{s-1}$  steps. Now, to pebble the sink of  $P_{2^s}$ ,  $\text{RBPPath}_s$  first pebbles the sink of the first half, i.e., the midpoint, then leaves the midpoint pebbled and pebbles the second half, and finally keeps the sink pebbled while removing the pebble on the midpoint. Thus,  $\text{RBPPath}_s$  requires at most  $s + 1$  pebbles and takes at most  $3 \cdot 3^{s-1} = 3^s$  steps. This proves the lemma.  $\square$

**Lemma 3.** Starting from the all-red configuration,  $P_{2^s}$  can be gray-pebbled using at most  $\sigma := s + 1$  gray pebbles in at most  $\tau := 3^s = 2^{s \log(3)}$  steps.

**Space-Optimal BGR Strategy.** Let's denote the pebbling strategy by  $\text{BGRPath}$ . For ease of presentation, we make two simplifying assumptions:

1. the path has  $S = 2^s - 1$  vertices  $[1, S]$  – the strategy can easily be extended to paths of arbitrary length; and
2. the final configuration produced by  $\text{BGRPath}$  is not the all-red configuration as prescribed in Definition 12 but the *penultimate* one with a gray pebble on the source and the rest pebbled red (i.e.,  $(\text{G}, \text{R}, \dots, \text{R})$ ) – getting to the all-red configuration from that is possible in one move by applying Rule 12.3.


The strategy  $\text{BGRPath}$  is described below recursively in  $s$ .

For the base case of  $s = 2$ ,  $\text{BGRPath}_2$  works as follows (see Figure 3.(a)):

1. Place a black pebble on 1, 2 and 3 (in that order) (Rule 12.1:  $\perp \mapsto \text{B}$ )
2. Replace the black pebbles on 3, 2 and 1 (in that order) with a gray pebble (Rule 12.2:  $\text{B} \mapsto \text{G}$ )
3. Replace the gray pebble on 3 and 2 (in that order) with a red pebble (Rule 12.3:  $\text{G} \mapsto \text{R}$ )

For  $s > 2$ , let's refer to the vertex  $2^{s-1}$  as the *midpoint* of  $P_S$ , the path with  $S = 2^s - 1$  vertices. Moreover, let's refer to  $P_{S|_{[1, 2^{s-1}-1]}}$  and  $P_{S|_{[2^{s-1}+1, S]}}$  as the first and second halves of  $P_S$  respectively. Then  $P_S$  can be switched from the empty configuration to the penultimate configuration as follows (see Figure 3).

1. Recursively place a black pebble on the midpoint  $2^{s-1}$   
( $\text{RBPPath}(P_{S|_{[1, 2^{s-1}]}})$ )
2. Recursively switch the second half to its penultimate configuration  
( $\text{BGRPath}_{s-1}(P_{S|_{[2^{s-1}+1, S]}})$ )



**Fig. 3.** Recursive step in BGR pebbling strategy  $\text{BGRPath}$ . We start from an empty configuration (leftmost path) and apply Item 1 through Item 6 to end up with the penultimate configuration (rightmost path).

3. Replace the black pebble on the midpoint  $2^{s-1}$  with a gray pebble (Rule 12.2)
4. Replace the gray pebble on  $2^{s-1} + 1$  with a red pebble (Rule 12.3)
5. Recursively switch the first half to its penultimate configuration ( $\text{BGRPath}_{s-1}(P_{S|[1,2^{s-1}]})$ )
6. Recursively replace the gray pebble on the midpoint  $2^{s-1}$  with a red pebble ( $\text{RGRPath}^{-1}(P_{S|[2,2^{s-1}]})$ )

**Lemma 4.** For  $S = 2^s - 1$ ,  $P_S$  can be BGR-pebbled using at most  $\sigma := s + 2$  grayscale pebbles in at most  $\tau := 2^{s+1} \cdot 3^s = 2^{s(\log(3)+1)+1}$  steps.

*Proof.* For the base case of  $s = 2$ ,  $\text{BGRPath}_2$  requires at most  $s + 2 = 3$  grayscale pebbles and a red pebble is never placed on the start vertex 1. Now, let's assume that on input  $P_{2^{s-1}-1}$ , the path with  $2^{s-1} - 1$  vertices,  $\text{BGRPath}_{s-1}$

1. switches  $P_{2^{s-1}-1}$  from the empty configuration to the penultimate configuration using  $s - 1 + 2 = s + 1$  grayscale pebbles; and
2. never places a red pebble on the start vertex 1.

For the induction step, first note that the inductive step in Item 2 is only possible thanks to the black pebble left as a “marker” on  $2^{s-1}$  in the previous step. This ensures that whenever a black pebble needs to be placed on  $2^{s-1} + 1$  this can be done thanks to the marker. Moreover, the induction hypothesis ensures that a red pebble never has to be placed on  $2^{s-1} + 1$  throughout the execution of  $\text{BGRPath}_{s-1}(P_{S|[2^{s-1}+1, S]})$  (since  $2^{s-1} + 1$  is the start vertex of the second half). Analogously, in the inductive step in Item 5 the (now) gray pebble on  $2^{s-1}$  acts as the marker: whenever a black or gray pebble needs to be placed on  $2^{s-1} - 1$  this can be done thanks to the marker.

To complete the proof, we have to argue that (i) the number of grayscale pebbles used in  $\text{BGRPath}_s$  is one more than that required for the inductive step

BGRPath<sub>s-1</sub>; (ii) the number of steps used in BGRPath<sub>s</sub> is  $2^{s+1} \cdot 3^s$ ; and (iii) the first vertex never carries a red pebble. To see why (i) is true, note that the number of pebbles required for BGRPath<sub>s</sub> is governed by the recursion

$$\sigma(s) = \max(2, 1 + \sigma(s - 1), s + 2) \leq s + 2$$

where the  $(s + 2)$  term is to account for the reversible pebbling carried out in Item 6 (while keeping the first node gray-pebbled), and the last inequality follows by the induction hypothesis. For (ii), write the time-complexity of BGRPath<sub>s</sub> as

$$\tau(s) \leq 2 \cdot \tau(\text{RBPPath}_{s-1}) - 2 + 2 \cdot \tau(s - 1) + 2 = 2 \cdot (3^{s-1} + \tau(s - 1)),$$

where we used the fact that in Item 6 the first node remains gray-pebbled, which clearly saves at least two steps. For (iii), note that the first node is either not pebbled or black-pebbled in Item 1, not pebbled throughout Items 2 to 4, never pebbled red by induction hypothesis in Item 5, and left gray-pebbled in Item 6.  $\square$


*Extending to layered graphs.* The above strategy can easily be extended to layered graphs of depth  $D$  and width  $w$  by pebbling the whole layer in which a pebble lies. This will require  $O(w \log(D))$  grayscale pebbles. However, since the number of pebbles that this strategy requires is proportional to the width of the circuit (and therefore the input-length), the resulting security bounds in Theorem 4 do not yield any advantage over simply guessing the input. We partially remedy this in the next sections.

## 4.2 BGR Pebbling via Separators

The strategy we describe, BGRSwitch, is implicit in the simulation in [20]. Intuitively, it is obtained by substituting the notion of “midpoint” in BGRPath, the strategy for paths discussed in Section 4.1, with that of a graph separator from Definition 11. As a consequence of Theorem 1, a graph  $G$  with treewidth  $w(G)$  can be recursively decomposed using separators of size at most  $w(G)$  into smaller and smaller “component” sub-graphs till the sub-graph is of a manageable (constant) size.<sup>17</sup> As a result, one gets a “component tree” out of the graph, starting with the whole graph at the root and ending with manageable-sized sub-graphs as leaves (see Figure 4). For a graph with  $S$  vertices and degree  $\delta$ , the depth of the component tree is at most  $O(\log S)$  and its out-degree is at most  $\delta \cdot |S|$  (since each vertex in  $\mathcal{S}$  can be connected to at most  $\delta$  components). The pebbling strategy using separators exploits this recursive structure to minimise the number of grayscale pebbles employed.

*Remark 7.* Note that Theorem 1 does not provide any guarantees on whether such a sequence of separators can be found *efficiently*. This becomes crucial when simulating the hybrids since it determines the factor  $T_p$ . We address this question at the end of this section.

<sup>17</sup> This exploits the fact that treewidth is a monotone property: removing vertices or edges from a graph can only decrease its treewidth.



**Fig. 4.** Binary component tree of a graph  $G$  (for simplicity). The graphs  $C_1$  and  $C_2$  on the second level denote the two components that result from the separator  $S_1$  for  $G$ . Similarly, the vertices that form the leaves result from further splitting of  $C_1$  and  $C_2$  using  $S_1$  and  $S_2$ , respectively.

**RB Pebbling via Separators.** As in the case of BGRPath, we first describe RBTreewidth, a space-efficient RB pebbling strategy that will be used as a subroutine in BGRSwitch. RBTreewidth places a black pebble on any vertex on a graph  $G$  of size  $S$  and treewidth  $w$  using  $\sigma := O(w \log(S))$  pebbles. To the best of our knowledge, this strategy is new and might be of independent interest.<sup>18</sup> Since the strategy is reversible, by  $\text{RBTreewidth}^{-1}$  we denote the reverse strategy that *removes* a black pebble. This strategy, thanks to the observation in Remark 5.2, will also be used to both place or remove a gray pebble on an all-red-pebbled graph (RGRTreewidth).

**Lemma 5.** *Every node in a DAG  $G$  with  $S$  vertices, in-degree  $\delta_{in}$  and treewidth  $w$  can be black-pebbled following the RB pebbling rules using at most  $\sigma := O((\delta_{in} + w) \log(S))$  black pebbles in at most  $\tau := (\delta_{in} w)^{O(\log(S))}$  steps.*

**Lemma 6.** *Starting from the all-red configuration, every node in a DAG  $G$  with  $S$  vertices, in-degree  $\delta_{in}$  and treewidth  $w$  can be gray-pebbled following the BGR pebbling rules using at most  $\sigma := O((\delta_{in} + w) \log(S))$  gray(scale) pebbles in at most  $\tau := (\delta_{in} w)^{O(\log(S))}$  steps.*

*Proof (of Lemma 5).* We denote the pebbling strategy by RBTreewidth and it takes as input a graph (component)  $C$  and a vertex  $v^*$  to be pebbled. It uses the same recursive decomposition into components as will be in BGRSwitch (i.e., the component tree). The base case is when the graph  $C$  is of small enough size (i.e., with  $O(1)$  vertices) and here RBTreewidth simply places a black pebble on  $v^*$  using as many black pebbles as needed; i.e.:

<sup>18</sup> It is possible to bound the space-complexity of RB pebbling on DAGs of treewidth  $w$  using existing results. First, use the fact that the RB pebbling number of a graph of size  $S$  is upper bounded by the *plain* black pebbling [39] number with a multiplicative  $\log(S)$  factor [32]. Second, use the fact that the black pebbling number is upper bounded by treewidth  $w$  (via so-called pathwidth) with another multiplicative  $\log(S)$  factor [11, Theorem 2, Corollary 24]. Consequently, we get that the RB pebbling number is at most  $w \log^2(S)$ . But this is a worse bound compared to what we show directly in Lemma 5.

1. place black pebbles on all vertices in  $C$  in topological order (Rule 7.1); and then
2. remove the black pebbles on the ancestors of  $v^*$  in reverse topological order (Rule 7.1).

Otherwise,  $\text{RBTreewidth}$  splits  $C$  into smaller components using its separator, recursively places a black pebble on every vertex in the separator in *topological order*, places a black pebble on  $v^*$  by recursing on the component  $C^*$  that contains  $v^*$ . Finally, it recursively removes the black pebbles on the separator in *reverse topological order*. The details are given below.

1. Decomposes  $C$  into its components  $C_1, \dots, C_p$  using its separator  $\mathcal{S} \subseteq \mathcal{C}$ , where  $\mathcal{C} = \mathcal{S} \cup \mathcal{C}_1 \cup \dots \cup \mathcal{C}_p$  and  $C_i := C|_{\mathcal{C}_i}$ .
2. Recursively place black pebbles on the vertices in  $\mathcal{S}$  in topological order (analogous to Item 1). That is, for each vertex  $s \in \mathcal{S}$  chosen in topological order:
  - (a) recursively place a black pebble on each predecessor of  $s$  in topological order,
  - (b) place a black pebble on  $s$  (Rule 7.1), and
  - (c) recursively remove the black pebbles on each predecessors of  $s$  in reverse topological order.
3. Recursively pebble the component  $C^* \in C_1, \dots, C_p$  which contains  $v^*$ . (Analogous to Item 2)
4. Undo Item 2 by recursively removing the black pebbles on the separator in reverse topological order (analogous to Item 3). That is, for each vertex  $s \in \mathcal{S}$  chosen in reverse topological order:
  - (a) recursively place a black pebble on each predecessor of  $s$  in topological order,
  - (b) remove the black pebble on  $s$  (Rule 7.1), and
  - (c) recursively remove the black pebbles on each predecessors of  $s$  in reverse topological order.

As we explain next, carrying out Items 2 and 4 in topological and reverse topological order, respectively, is crucial for the efficiency (and correctness) of  $\text{RBTreewidth}$ . Recall that the property of the separator  $\mathcal{S}$  guarantees that the components  $C_1, \dots, C_p$  are themselves of small enough size (see Definition 11). Therefore, once  $\mathcal{S}$  is pebbled black  $\text{RBTreewidth}$  can be called on all the resulting components as there are no edges between the components. However, pebbling a vertex  $s$  in  $\mathcal{S}$  itself is tricky: the predecessors of  $s$  could very well be in different components (since there are no guarantees for the vertices in the separator). However, we do have the guarantee that all the predecessors of a predecessor of  $s$  belong to the same component or the separator, and are reachable via either source vertices *or* vertices belonging to  $\mathcal{S}$ . Therefore, as long as the vertices in  $\mathcal{S}$  are black-pebbled in topological order,  $\mathcal{S}$  can be completely pebbled in Item 2 by recursing on small enough components. A similar argument applies when the black pebbles on the separator are removed in Item 4.



With this in mind, we now analyse `RBTreeWidth`. The reason this strategy requires at most  $\sigma := O(w \log(S))$  black pebbles is similar to what we will see in the proof of Theorem 5. The number of pebbles is governed by the expression

$$\sigma(i) \leq (w + \delta_{in}) + \sigma(i + 1), \quad (3)$$

where the index  $i$  is the depth of the recursion of `RBTreeWidth`. The factor  $(w + \delta_{in})$  is the cost of black pebbles placed on the separator in Item 2 and the factor  $\sigma(i + 1)$  is the cost of recursions in Items 2 to 4. Note that since the size of the components in each these recursive calls is at most  $2/3$  of the size of the original component  $C$ , the overall depth of the recursion remains  $O(\log(S))$ . The upper bound on the number of pebbles claimed in the lemma follows on solving Equation (3).

As for the number of steps, it is governed by the expression

$$\tau(i) \leq \tau(i + 1)O(\delta_{in}w). \quad (4)$$

since `RBTreeWidth` is recursively called at most  $O(\delta_{in}w)$  times on the (sub-) components. As in the case of space-complexity, since we end up with constant-size components at the end of the recursion, the base cost is  $O(1)$ . The lemma follows on solving Equation (4).  $\square$

**Recursive Switching.** We are now primed to describe `BGRSwitch`. It takes as input:

1. the original graph  $G = (\mathcal{V}, \mathcal{E})$  that is to be pebbled
2. the vertices  $\mathcal{C} \subseteq \mathcal{V}$  that define the graph component  $C = G_{|\mathcal{C}}$  being currently considered
3. the “higher” separator  $\mathcal{U}$ , which is the union of all the separators in the “higher” recursive calls that resulted in the creation of the current component  $C$ .

Note that  $\mathcal{C}$  and  $\mathcal{U}$  are disjoint sets by definition. Throughout the execution of `BGRSwitch`, we maintain a few pebbling properties as invariants:

- At the start of the execution of `BGRSwitch` on the current component  $C$ , it is guaranteed that the vertices in  $\mathcal{U}$  are all black-pebbled. This, in some sense, “isolates”  $C$  from the rest of the graph and, as a result, it can be pebbled independently of the rest.
- At the end of the execution of `BGRSwitch`, we guarantee that the vertices  $\mathcal{C}$  in  $C$  are red-pebbled (via black and then gray), *except* for the children of the higher separator  $\mathcal{U}$ , which will be left gray-pebbled.

Next, let’s see what happens in `BGRSwitch` when called on  $(G, \mathcal{C}, \mathcal{U})$  (in the first call  $\mathcal{C} = \mathcal{V}$  and  $\mathcal{U} = \emptyset$ ). (For ease of understanding, in parenthesis we refer to the analogous step in `BGRPath`, the BGR pebbling strategy for paths that was described in Section 4.1.) The base case is when the current component  $C := G_{|\mathcal{C}}$  is of small enough size (i.e., with  $O(1)$  vertices). Here `BGRSwitch` simply switches  $C$  to red by using as many pebbles as needed; i.e.:

1. place black pebbles on all vertices in  $C$  (Rule 12.1),
2. replace them with gray pebbles (Rule 12.2), and
3. replace the gray pebbles with red pebbles (Rule 12.3) *except* if the vertex is a child of the upper separator  $\mathcal{U}$ .

Otherwise, BGRSwitch does the switching from no pebbles to red pebbles for  $C$  by recursively splitting into smaller components using the separator for  $C$  as follows (see Figure 5).


1. Decompose  $C = G|_{\mathcal{C}}$  into its components  $C_1, \dots, C_p$  using its separator  $\mathcal{S} \subseteq \mathcal{C}$ , where  $\mathcal{C} = \mathcal{S} \cup \mathcal{C}_1 \cup \dots \cup \mathcal{C}_p$  and  $C_i := G|_{\mathcal{C}_i}$ . (Midpoint-separator analogy)
2. Place black pebbles on the vertices in  $\mathcal{S}$  using RBTreewidth. Note that this is possible only because all the vertices that are required to carry this out are either empty or belong to  $\mathcal{U}$  and therefore are black-pebbled. (Analogous to Item 1)
3. Recursively switch each component  $C_1, \dots, C_p$  using BGRSwitch. After all component are switched, all vertices in  $\mathcal{C}$ , except the ones that are children of  $\mathcal{S}$ , are red-pebbled; the children of  $\mathcal{S}$  are left gray-pebbled. (Analogous to Items 2 and 5)
4. Replace the black pebbles on the separator  $\mathcal{S}$  with gray pebbles by using Rule 12.2.
5. Replace the gray pebbles on  $\mathcal{S}$  and its adjacent vertices with red pebbles (except if the vertex is a child of the upper separator) using Rule 12.3 and RGRTreewidth. (Analogous to Item 6)

Note that during the whole strategy, we maintain as invariant a black-gray frontier between the empty and red-pebbled vertices, and this frontier is exactly at the separators. That is, at any point of the pebbling *no* two vertices such that one is empty and the other is red-pebbled are related (see Remark 5.5). As pointed out in the technical overview (Section 1.2), it is this frontier that insulates the computations in the two circuits and help ensure correctness at all times. In the following theorem we formally analyse its space- and time-complexity.

**Theorem 5 (Main theorem).** *Every DAG  $G$  with  $S$  vertices, degrees  $\delta_{in}, \delta_{out} \leq \delta$  and treewidth  $w$  can be BGR-pebbled using at most  $\sigma = O((\delta_{in} + w\delta_{out}) \log(S))$  grayscale pebbles in at most  $\tau := (\delta w)^{O(\log(S))}$  steps.*

*Proof.* To bound the space-complexity of BGRSwitch, first note that the algorithm indeed maintains the invariants stated above:

1. At the start of the execution of BGRSwitch on the current component  $C$ , it is guaranteed that the vertices in  $\mathcal{U}$  are all black-pebbled. Hence, on input a component at depth of recursion  $i \in [0, O(\log(S))]$ , there are  $|\mathcal{U}(i)| = O(iw)$  pebbles that remain black-pebbled.
2. At the end of the execution of BGRSwitch, all the vertices  $\mathcal{C}$  in  $C$  are red-pebbled, *except* for children of the higher separator  $\mathcal{U}$ , which will be left gray-pebbled. Hence, after the execution of BGRSwitch on a component at depth  $i > 0$  there are up to  $\delta_{out} \cdot |\mathcal{U}(i)| = O(i\delta_{out}w)$  many gray pebbles on the graph.



**Fig. 5.** A schematic diagram demonstrating recursive switching using BGRSwitch. The separator  $\mathcal{S}$ , shown as a solid cross, divides the graph (in thick) into four components. The dotted lines indicate vertices adjacent to  $\mathcal{S}$ . The switching starts from an empty graph (1) and then proceeds as follows: (2) pebble  $\mathcal{S}$  black, (3-6) switch bottom-left, top-left, top-right and bottom-right components (in that order), (7) replace black pebbles on  $\mathcal{S}$  with gray pebbles, (8) replace the gray pebbles on vertices adjacent to  $\mathcal{S}$  with red pebbles, and (9) replace the gray pebble on  $\mathcal{S}$  with red pebbles.

Now, in Item 3 there are up to  $\delta w$  many components, among which some are already switched, some not, and one is currently processed. For the former set of components, all nodes within these which are children of  $\mathcal{U} \cup \mathcal{S}$  are pebbled gray. Hence, by the above, there are up to  $\delta_{out} \cdot (|\mathcal{U}(i)| + 1) + |\mathcal{U}(i)| = O(i\delta_{out}w)$  nodes that remain gray- or black-pebbled while BGRSwitch is processed on a lower component. Now, while some node on the separator  $\mathcal{S}'$  in the currently-processed component is pebbled using RBTreewidth or RGRTreewidth (cf. Item 2), there are up to  $|\mathcal{S}'| \leq w$  additional nodes that remain pebbled. By Lemmas 5 and 6, the space-complexity of RBTreewidth/RGRTreewidth is bounded by  $O((\delta_{in} + w) \log(S))$ . Thus, we arrive at

$$\sigma(i) \leq O(i\delta_{out}w) + w + O((\delta_{in} + w) \log(S)) = O((\delta_{in} + w\delta_{out}) \log(S)).$$

As for the number of steps, on input a component at depth of recursion  $i$ , the time-complexity of BGRSwitch is governed by the expression

$$\tau(i) \leq (\delta_{in}w)^{O(\log(S))} \cdot w + \tau(i+1)\delta w. \quad (5)$$

The first factor is the cost of the subroutines used to pebble the separator black: the subroutine is called at most  $|\mathcal{S}| \leq w$  times, each time incurring a cost of at most  $(\delta_{in}w)^{O(\log(S))}$  (Lemma 5). The second factor is the cost of recursively calling BGRSwitch on at most  $\delta w$  (sub-)components. Since we end up with constant-size components at the end of the recursion, the base cost is  $O(1)$ . On solving Equation (5), the theorem follows.  $\square$

*Computing the separators.* Finally, let us return to the question of computing the sequence of separators underlying our pebbling strategy. While we are not aware of an efficient algorithm for computing balanced separators (see discussion in Section 1.3), it suffices for our purpose that a separator of size  $w$  can be found in time at most  $S^{O(w)}$ : since we anyway lose a similar factor in the distinguishing advantage, the overall (asymptotic) loss that the reduction incurs remains similar. Therefore, we simply enumerate all  $w$ -sized subsets of vertices till we find a balanced separator – note that given a separator it is easy to verify that it is indeed one, i.e., the problem lies in **NP**. Since computing any BGR pebbling configuration requires knowledge of at most  $O(\log(S))$  many separators, the total time required to compute a pebbling configuration is at most  $T_p = O(\log(S)S^w)$ .

### 4.3 Optimised Piecewise Guessing

Recall that in Theorem 4, the loss in adaptive security is exponential in the BGR pebbling complexity. This is because the reduction requires as advice the value of the output wire of all the gates that are grayscale pebbled. Therefore when Theorem 4 is used in conjunction with Theorem 5, the loss is exponential in the treewidth *as well as* degree. First, note that for Yao’s garbling scheme, we only consider Boolean circuits with fan-in 2. We argue next that the dependence on

*out-degree* can be removed thanks to the structure of the configurations in the BGRSwitch pebbling strategy. The resulting theorem is stated in Theorem 6.

Let's return to the recursive step in Item 3 which is the cause of dependence on the degree. At the start of this step, all the vertices in the separator  $\mathcal{S}$  have been pebbled. Then each component  $C_i$  is recursively switched to red one at a time. At the end of switching  $C_i$ , each vertex in  $C_i$  is pebbled red, *except* for those vertices that are children of  $\mathcal{S}$  (or  $\mathcal{U}$ ) which are left gray. Therefore we can restrict our focus on those vertices that have its predecessors in the separator – let's consider one such vertex  $v^*$ . Note that in any configuration where  $v^*$  carries a gray pebble, it is guaranteed that its predecessors in the separator are black-pebbled. Therefore, instead of requiring the value of the gate  $g^*$  corresponding to  $v^*$  as an advice, it can simply be computed as a function of the values of its predecessor gates (which *are* included in the advice). To sum up, instead of providing as advice the values of the output wires of *all* the gates that are grayscale pebbled as in Equation (2), it suffices to provide a much smaller advice as outlined above. As a result of this observation, we get the following optimised version of Theorem 4. This leads to the corollaries stated in Section 1.1.

**Theorem 6.** *Suppose that a class of circuits  $\mathcal{C}$  of size  $S$ , fan-in 2 has degree  $\delta$  and treewidth  $w$ . If the encryption scheme **SKE** is  $(\epsilon, T)$ -secure then **YGC<sub>SKE</sub>** is  $(3\tau 2^\sigma \epsilon, T - (T_H + T_\sigma + T_p))$ -adaptively-indistinguishable for  $\mathcal{C}$  where*

$$\tau := (\delta w)^{O(\log(S))}, \quad \sigma := O(w \log(S)) \quad \text{and} \quad T_p := O(\log(S)S^w).$$

## 5 Conclusion and Open Problems

Yao's garbling scheme is one of the most fundamental cryptographic constructions. In this work, we took another step towards completing the landscape of its security. Our result leads to several interesting questions, the most natural being whether the upper bound on loss in security can be improved. To this end, one could look at other (orthogonal) graph properties. Another pressing question is whether there are other applications of treewidth in cryptography (which seems relatively overlooked compared to other fields such as circuit complexity or algorithmic graph theory). This closely concerns the divide-and-conquer approach employed in our security reduction: it seems that the approach of surgically replacing one circuit with another should find use in other scenarios. Our hope is that this work spurs further research in this direction.

**Acknowledgements.** We would like to thank Daniel Wichs for helpful discussions on the landscape of adaptive security of Yao's garbling.

## References

- [1] M. Alekhnovich and A. Razborov. Satisfiability, branch-width and tseitin tautologies. *computational complexity*, 20(4):649–678, 2011.
- [2] E. Allender, S. Chen, T. Lou, P. A. Papakonstantinou, and B. Tang. Width-parametrized SAT: time–space tradeoffs. *Theory Comput.*, 10:297–339, 2014.
- [3] P. Ananth and A. Lombardi. Succinct garbling schemes from functional encryption through a local simulation paradigm. In A. Beimel and S. Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 455–472. Springer, Heidelberg, Nov. 2018.
- [4] B. Applebaum, Y. Ishai, E. Kushilevitz, and B. Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 166–184. Springer, Heidelberg, Aug. 2013.
- [5] M. Bellare, V. T. Hoang, and S. Keelveedhi. Instantiating random oracles via UCEs. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 398–415. Springer, Heidelberg, Aug. 2013.
- [6] M. Bellare, V. T. Hoang, and P. Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In X. Wang and K. Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 134–153. Springer, Heidelberg, Dec. 2012.
- [7] M. Bellare, V. T. Hoang, and P. Rogaway. Foundations of garbled circuits. In T. Yu, G. Danezis, and V. D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, Oct. 2012.
- [8] C. H. Bennett. Time/space trade-offs for reversible computation. *SIAM Journal on Computing*, 18(4):766–776, 1989.
- [9] H. L. Bodlaender. Nc-algorithms for graphs with small treewidth. In J. van Leeuwen, editor, *Graph-Theoretic Concepts in Computer Science, 14th International Workshop, WG '88, Amsterdam, The Netherlands, June 15-17, 1988, Proceedings*, volume 344 of *Lecture Notes in Computer Science*, pages 1–10. Springer, 1988.
- [10] H. L. Bodlaender. A tourist guide through treewidth. *Acta Cybern.*, 11(1-2):1–21, 1993.
- [11] H. L. Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical Computer Science*, 209(1):1 – 45, 1998.
- [12] D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Heidelberg, May 2014.
- [13] R. P. Brent. The parallel evaluation of general arithmetic expressions. *J. ACM*, 21(2):201–206, Apr. 1974.

- [14] T. N. Bui and C. Jones. Finding good approximate vertex and edge partitions is np-hard. *Inf. Process. Lett.*, 42(3):153–159, 1992.
- [15] S. A. Cook and R. Sethi. Storage requirements for deterministic polynomial time recognizable languages. *J. Comput. Syst. Sci.*, 13(1):25–37, 1976.
- [16] M. Cryan and P. B. Miltersen. On pseudorandom generators in NC. In *MFCS*, volume 2136 of *Lecture Notes in Computer Science*, pages 272–284. Springer, 2001.
- [17] G. Even, J. S. Naor, S. Rao, and B. Schieber. Fast approximate graph partitioning algorithms. *SIAM Journal on Computing*, 28(6):2187–2214, 1999.
- [18] U. Feige, M. T. Hajiaghayi, and J. R. Lee. Improved approximation algorithms for minimum-weight vertex separators. In H. N. Gabow and R. Fagin, editors, *37th ACM STOC*, pages 563–572. ACM Press, May 2005.
- [19] U. Feige and M. Mahdian. Finding small balanced separators. In J. M. Kleinberg, editor, *38th ACM STOC*, pages 375–384. ACM Press, May 2006.
- [20] A. Gál and J. Jang. A generalization of Spira’s theorem and circuits with small segregators or separators. *Inf. Comput.*, 251:252–262, 2016.
- [21] S. Garg and A. Srinivasan. Adaptively secure garbling with near optimal online complexity. In J. B. Nielsen and V. Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 535–565. Springer, Heidelberg, Apr. / May 2018.
- [22] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In A. Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.
- [23] B. Hemenway, Z. Jafargholi, R. Ostrovsky, A. Scafuro, and D. Wichs. Adaptively secure garbled circuits from one-way functions. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 149–178. Springer, Heidelberg, Aug. 2016.
- [24] R. Impagliazzo and M. Naor. Efficient cryptographic schemes provably as secure as subset sum. *Journal of Cryptology*, 9(4):199–216, Sept. 1996.
- [25] Z. Jafargholi, C. Kamath, K. Klein, I. Komargodski, K. Pietrzak, and D. Wichs. Be adaptive, avoid overcommitting. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 133–163. Springer, Heidelberg, Aug. 2017.
- [26] Z. Jafargholi and S. Oechsner. Adaptive security of practical garbling schemes. In K. Bhargavan, E. Oswald, and M. Prabhakaran, editors, *Progress in Cryptology – INDOCRYPT 2020*, pages 741–762, Cham, 2020. Springer International Publishing.
- [27] Z. Jafargholi, A. Scafuro, and D. Wichs. Adaptively indistinguishable garbled circuits. In Y. Kalai and L. Reyzin, editors, *TCC 2017, Part II*, volume 10678 of *LNCS*, pages 40–71. Springer, Heidelberg, Nov. 2017.
- [28] Z. Jafargholi and D. Wichs. Adaptive security of Yao’s garbled circuits. In M. Hirt and A. D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 433–458. Springer, Heidelberg, Oct. / Nov. 2016.
- [29] M. J. Jansen and J. Sarma. Balancing bounded treewidth circuits. *Theory Comput. Syst.*, 54(2):318–336, 2014.

- [30] C. Kamath, K. Klein, K. Pietrzak, and D. Wichs. Limits on the adaptive security of Yao’s garbling, *CRYPTO 2021*, To Appear.
- [31] F. Kitagawa, R. Nishimaki, K. Tanaka, and T. Yamakawa. Adaptively secure and succinct functional encryption: Improving security and efficiency, simultaneously. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 521–551. Springer, Heidelberg, Aug. 2019.
- [32] R. Y. Levine and A. T. Sherman. A note on Bennett’s time-space tradeoff for reversible computation. *SIAM Journal on Computing*, 19(4):673–677, 1990.
- [33] Y. Lindell and B. Pinkas. A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, Apr. 2009.
- [34] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- [35] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.
- [36] D. Lokshтанov, I. Mikhailin, R. Paturi, and P. Pudlák. Beating brute force for (quantified) satisfiability of circuits of bounded treewidth. In A. Czumaj, editor, *29th SODA*, pages 247–261. ACM-SIAM, Jan. 2018.
- [37] D. Marx. Parameterized graph separation problems. In R. G. Downey, M. R. Fellows, and F. K. H. A. Dehne, editors, *Parameterized and Exact Computation, First International Workshop, IWPEC 2004, Bergen, Norway, September 14-17, 2004, Proceedings*, volume 3162 of *Lecture Notes in Computer Science*, pages 71–82. Springer, 2004.
- [38] J. Nordström. *New Wine into Old Wineskins: A Survey of Some Pebbling Classics with Supplemental Results*. 2015.
- [39] M. S. Paterson and C. E. Hewitt. Record of the project mac conference on concurrent systems and parallel computation. chapter Comparative Schematology, pages 119–127. ACM, New York, NY, USA, 1970.
- [40] N. Robertson and P. D. Seymour. Graph minors. II. algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.
- [41] R. Santhanam. On separators, segregators and time versus space. In *Computational Complexity Conference*, pages 286–294. IEEE Computer Society, 2001.
- [42] P. Spira. On time-hardware complexity of tradeoffs for boolean functions. In *Proc. 4th Hawaii Symp. System Sciences*, page 525–527. North Hollywood and Western Periodicals, 1971.
- [43] A. C.-C. Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164. IEEE Computer Society Press, Nov. 1982.
- [44] A. C.-C. Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, Oct. 1986.



## A RGB Complexity and Non-Determinism

An alternative way to express our results is through Santhanam’s black-white-red (BWR) pebble game [41]. These were introduced in order to study the simulation of non-deterministic complexity classes using deterministic ones. The BWR pebble game allows us to state our results purely in terms of pebbling complexity instead of graph properties like separators or treewidth. It might also turn out useful to directly study BWR pebbling complexities of DAGs to directly prove adaptive security of Yao’s garbling of the underlying circuit.

**Definition 13 (Black-White-Red (BWR) pebble game [41]).** *Consider a DAG  $G = (\mathcal{V}, \mathcal{E})$  with  $\mathcal{V} = [1, S]$  and let  $\mathcal{X} = \{\perp, B, W, R\}$  denote the set of colours of the pebbles. Consider a sequence  $\mathcal{P} := (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  of pebble configurations for  $G$ , where  $\mathcal{P}_i \in \mathcal{X}^\mathcal{V}$  for all  $i \in [0, \tau]$ . We call such a sequence a BWR pebbling strategy for  $G$  if (i) every vertex is empty in the initial configuration (i.e.,  $\mathcal{P}_0 = (\perp, \dots, \perp)$ ), (ii) every vertex is black-pebbled in the final configuration (i.e.,  $\mathcal{P}_\tau = (B, \dots, B)$ ), and (iii) every configuration is obtained by applying one of the following rules to its preceding configuration:*

1.  $\perp \mapsto W$ : a white pebble can be placed on an empty vertex at any time
2.  $W \mapsto R$ : a white pebble on a vertex  $v$  can be replaced by a red pebble if  $v$ ’s predecessors all carry a pebble, either red or white
3.  $R \mapsto B$ : a red pebble on a vertex  $v$  can be replaced by a black pebble if  $v$ ’s successors all carry a red pebble.

The space-complexity of a strategy  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_\tau)$  for a DAG  $G$  is defined as

$$\sigma_G(\mathcal{P}) := \max_{i \in [0, \tau]} |\{j \in [1, S] : \mathcal{P}_i(j) \in \{W, R\}\}|.$$

Intuitively, the white and red pebbles represent (non-deterministic) guesses. A white pebble (as in the black-white pebble game [15]) represents an unverified guess of the value corresponding to that vertex, while a red pebble represents a guess that is consistent with prior guesses that the vertex depends on. A black pebble represents a guess that need not be stored any longer. We find the following lemma from [41, Lemma 4.3] useful as it allows to state Theorem 5 in terms of BWR pebbling (Theorem 7).

**Lemma 7 ([41]).** *If a graph  $G = (\mathcal{V}, \mathcal{E})$  with  $|\mathcal{V}| = S$  has a strategy for the BWR game that uses space  $\sigma_{BWR}(S)$  then  $G$  has an  $\sigma_{BWR}(S)$  separator.*

**Theorem 7.** *Every DAG  $G$  with  $S$  vertices and degree  $\delta$  that can be BWR-pebbled using  $\sigma_{BWR}$  pebbles can be BGR-pebbled using most  $\sigma = O(\sigma_{BWR} \delta \log(S))$  (grayscale) pebbles in at most  $O(\sigma_{BWR}^{\log(S)})$  steps.*

**Open Question 1** *Can every BWR pebbling strategy be efficiently simulated using a BGR pebbling strategy? This would make the results in Theorem 7 explicit.*