Optimal Shuffle Code with Permutation Instructions

Sebastian Buchwald, Manuel Mohr, and Ignaz Rutter

Karlsruhe Institute of Technology
{sebastian.buchwald, manuel.mohr, rutter}@kit.edu

Abstract. During compilation of a program, register allocation is the task of mapping program variables to machine registers. During register allocation, the compiler may introduce *shuffle code*, consisting of copy and swap operations, that transfers data between the registers. Three common sources of shuffle code are conflicting register mappings at joins in the control flow of the program, e.g, due to if-statements or loops; the calling convention for procedures, which often dictates that input arguments or results must be placed in certain registers; and machine instructions that only allow a subset of registers to occur as operands. Recently, Mohr et al. [8] proposed to speed up shuffle code with special hardware instructions that arbitrarily permute the contents of up to five registers and gave a heuristic for computing such shuffle codes.
In this paper, we give an efficient algorithm for generating optimal shuffle code in the setting of Mohr et al. An interesting special case occurs when no register has to be transferred to more than one destination, i.e., it suffices to permute the contents of the registers. This case is equivalent to factoring a permutation into a minimal product of permutations, each of which permutes up to five elements.

1 Introduction

One of the most important tasks of a compiler during code generation is register allocation, which is the task of mapping program variables to machine registers. During this phase, it is frequently necessary to insert so-called *shuffle code* that transfers values between registers. Common reasons for the insertion of shuffle code are control flow joins, procedure calling conventions and constrained machine instructions.

The specification of a shuffle code, i.e., a description which register contents should be transferred to which registers, can be formulated as a directed graph whose vertices are the registers and an edge (u, v) means that the content of u before the execution of the shuffle code must be in v after the execution. Naturally, every vertex must have at most one incoming edge. Note that vertices may have several outgoing edges, indicating that their contents must be transferred to several destinations, and even loops (u, u), indicating that the content of register u must be preserved. We call such a graph a *Register Transfer Graph* or *RTG*. Two important special types of RTGs are outdegree-1 RTGs where the maximum out-degree is 1 and PRTGs where $\deg^-(v) = \deg^+(v) = 1$ for all vertices v ($\deg^-$ and $\deg^+$ denote the in- and out-degree of a vertex, respectively).

We say that a shuffle code, consisting of a sequence of copy and swap operations on the registers, *implements* an RTG if after the execution of the shuffle code every register whose corresponding vertex has an incoming edge has the correct content. The *shuffle code generation* problem asks for a shortest shuffle code that implements a given RTG.

Fig. 1: Two example RTGs where the optimal shuffle code is not obvious.

The amount of shuffle code directly depends on the quality of copy coalescing, a subtask of register allocation [8]. As copy coalescing is NP-complete [2], reducing the amount of shuffle code is expensive in terms of compilation time, and thus cannot be afforded in all contexts, e.g., just-in-time compilation.

Therefore, it has been suggested to allow more complicated operations than simply copying and swapping to enable more efficient shuffle code. Mohr et al. [8] propose to allow performing permutations on the contents of small sets of up to five registers. The processor they develop offers three instructions to implement shuffle code:

`copy:` copies the content of one register to another one

`permi5:` cyclically shifts the contents of up to five registers

`permi23:` swaps the contents of two registers and performs a cyclic shift of the contents of up to three registers; the two sets of registers must be disjoint.

In fact, the two operations `permi5` and `permi23` together allow to arbitrarily permute the contents of up to five registers in a single operation. A corresponding hardware and a modified compiler that employs a greedy approach to generate the shuffle code have been shown to improve performance in practice [8]. While the greedy heuristic works well in practice, it does not find an optimal shuffle code in all cases.

It is not obvious how to generate optimal shuffle code using the three instructions `copy`, `permi5` and `permi23` even for small RTGs. In the left RTG from Fig. 1, a naive solution would implement edges $(1, 2)$ and $(1, 3)$ using copies and the remaining cycle $(4\,5\,6)$ using a `permi5`. However, using one `permi23` to implement the cycle $(4\,5\,6)$ and swap registers 1 and 2, and then copying register 2 to 3 requires only two instructions. This is legal because the contents of register 1 can be overwritten. The same trick is not applicable for the right RTG in Fig. 1 because of the loop $(1, 1)$ and hence three instructions are necessary to implement that RTG.

A maximum permutation size of 5 may seem arbitrary at first but is a consequence of instruction encoding constraints. In each `permi` instruction, the register numbers and their order must be encoded in the instruction word. Hence, $\lceil \log_2 \left(\binom{n}{k} k! \right) \rceil$ bits of an instruction word are needed to be able to encode all permutations of k registers out of n total registers. As many machine architectures use a fixed size for instruction words, e.g., 32 or 64 bits, and the operation type must also be encoded in the instruction word, space is very limited. In fact, for a 32 bit instruction word, 34 is the maximum number of registers that leave enough space for the operation type.

Related Work. As long as only copy and swap operations are allowed, finding an optimal shuffle code for a given RTG is a straightforward task [6, p. 56–57]. Therefore work in the area of compiler construction in this context has focused on coalescing techniques that reduce the number and the size of RTGs [1, 2, 5, 7].

From a theoretical point of view, the most closely related work studies the case where the input RTG consists of a union of disjoint directed cycles, which can be interpreted as a permutation π. Then, no copy operations are necessary for an optimal shuffle code and hence the problem of finding an optimal shuffle code using `permi23` and `permi5` is equivalent to writing π as a shortest product of permutations of maximum size 5, where a permutation of n elements has size k if it fixes $n - k$ elements.

There has been work on writing a permutation as a product of permutations that satisfy certain restrictions. The factorization problem on permutation groups from computational group theory [9] is the task of writing an element g of a permutation group as a product of given generators S. Hence, an algorithm for solving the factorization problem could be applied in our context by using all possible permutations of size 5 or less as the set S. However, the algorithms do not guarantee minimality of the product. For the case that S consists of all permutations that reverse a contiguous subsequence of the elements, known as the pancake sorting problem, it has been shown that computing a factoring of minimum size is NP-complete [3].

Farnoud and Milenkovic [4] consider a weighted version of factoring a permutation into transpositions. They present a polynomial constant-factor approximation algorithm for factoring a given permutation into transpositions where transpositions have arbitrary non-negative costs. For the case that the transposition costs are defined by a path-metric, they show how to compute a factoring of minimum weight in polynomial time. In our problem, we cannot assign costs to an individual transposition as its cost is context-dependent, e.g., four transpositions whose product is a cycle require one operation, whereas four arbitrary transpositions may require two.

Contribution and Outline. In this paper, we present an efficient algorithm for generating optimal shuffle code using the operations `copy`, `permi5`, and `permi23`, or equivalently, using copy operations and permutations of size at most 5.

We first prove the existence of a special type of optimal shuffle codes whose copy operations correspond to edges of the input RTG in Section 2. Removing the set of edges implemented by copy operations from an RTG leaves an outdegree-1 RTG.

We show that the greedy algorithm proposed by Mohr et al. [8] finds optimal shuffle codes for outdegree-1 RTGs and that the size of an optimal shuffle code can be expressed as a function that depends only on three characteristic numbers of the outdegree-1 RTG rather than on its structure. Since PRTGs are a special case of outdegree-1 RTGs, this shows that GREEDY is a linear-time algorithm for factoring an arbitrary permutation into a minimum number of permutations of size at most 5.

Finally, in Section 4, we show how to compute an optimal set of RTG edges that will be implemented by copy operations such that the remaining outdegree-1 RTG admits a shortest shuffle code. This is done by several dynamic programs for the cases that the input RTG is disconnected, is a tree, or is connected and contains a (single) cycle.

2 Register Transfer Graphs and Optimal Shuffle Codes

In this section, we rephrase the shuffle code generation problem as a graph problem. An RTG that has only self-loops needs no shuffle-code and is called *trivial*.

It is easy to define the effect of a permutation on an RTG. Let G be an RTG and let π be an arbitrary permutation that is applied to the contents of the registers. We define $\pi G = (V, \pi E)$, where $\pi E = \{(\pi(u), v) \mid (u, v) \in E\}$. This models the fact that if v should receive the data contained in u, then after π moves the data contained in u to some other register $\pi(u)$, the data contained in $\pi(u)$ should end up in v. We observe that for two permutations π_1, π_2 of V, it is $(\pi_2 \circ \pi_1)G = \pi_2(\pi_1(G))$, i.e., we have defined a group action of the symmetric group on RTGs. For PRTGs, the shuffle code generation problem asks for a shortest shuffle code that makes the given PRTG trivial.

Unfortunately, it is not possible to directly express copy operations in RTGs. Instead, we rely on the following observation. Consider an arbitrary shuffle code that contains a copy $a \to b$ with source a and target b that is followed by a transposition τ of the contents of registers c and d. We can replace this sequence by a transposition of the registers $\{c, d\}$ and a copy $\tau(a) \to \tau(b)$. Thus, given a sequence of operations, we can successively move the copy operations to the end of the sequence without increasing its length. Thus, for any RTG there exists a shuffle code that consists of a pair of sequences $((\pi_1, \ldots, \pi_p), (c_1, \ldots, c_t))$, where the π_i are permutation operations and the c_i are copy operations. We now strengthen our assumption on the copy operations.

Lemma 1. *Every instance of the shuffle code generation problem has an optimal shuffle code $((\pi_1, \ldots, \pi_p), (c_1, \ldots, c_t))$ such that*
  (i) No register occurs as both a source and a target of copy operations.
  (ii) Every register is the target of at most one copy operation.
  (iii) There is a bijection between the copy operations c_i and the edges of πG that are not loops, where $\pi = \pi_p \circ \pi_{p-1} \circ \cdots \circ \pi_1$.
  (iv) If u is the source of a copy operation, then u is incident to a loop in πG.
  (v) The number of copies is $\sum_{v \in V} \max\{\deg_G^+(v) - 1, 0\}$.

Proof. Consider an optimal shuffle code of the form $((\pi_1, \ldots, \pi_p), (c_1, \ldots, c_t))$ as above and assume that the number t of copy operations is minimal among all optimal shuffle codes.

Suppose there exists a register that occurs as both a source and a target of copy operations or a register that occurs as the target of more than one copy operation. Let k be the smallest index such that in the sequence $c_1, \ldots, c_k$ there is a register occurring as both a source and a target or a register that occurs as a target of two copy operations. We show that we can modify the sequence of copy operation such that the length of the prefix without such registers increases. Inductively, we then obtain a sequence without such registers. Let v and w denote the source and target of c_k, respectively. Let i denote the largest index such that c_i is a copy operation that has w as a source or target or such that c_i is a copy operation with target v. We distinguish three cases based on whether c_i has target v, target w, or source w.

Case 1: The target of c_i is v; see Fig. 2a. Let u denote the source of operation c_i. The sequence first copies a value from u to v and from there to w. We then replace c_k by a copy with source u and target w. (If $u = w$, we omit the operation altogether.) This only changes the outcome of the shuffle code if the value contained in u or v is modified between operations c_i and c_k, i.e., if there exists a copy operation c_j with $i < j < k$ whose target is either u or v. But then already the smaller sequence $c_1, \ldots, c_j$ has u
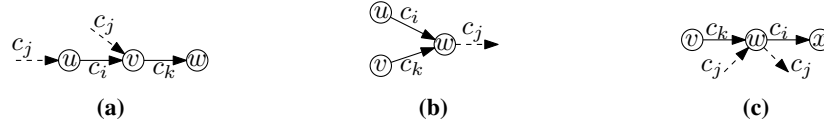
Fig. 2: Illustration of the proof of Lemma 1. The copies c_j with $i < j < k$ along the dashed edges would contradict either the choice of i or k.

occur as both a source and a target or v as a target of two operations, contradicting the minimality of k.

Case 2: The target of c_i is w; see Fig. 2b. In this case the copy operation c_i copies a value to w and later this value is overwritten by the operation c_k. Note that by the choice of i there is no operation c_j with $i < j < k$ with source w. Thus, omitting the copy operation c_i does not change the outcome of the shuffle code. A contradiction to optimality.

Case 3: The source of c_i is w; see Fig. 2c. Let x denote the target of operation c_i. In this case first a value is copied from w to x and later the value in v is copied to w. We claim that no copy operation c_j with $i < j < k$ involves x or w. If x occurs as the source of c_j (as the target of c_j), then x occurs as a source and target (two times as a target) in the sequence $c_1, \ldots, c_j$, contradicting the minimality of k. If w is the target of c_j, then w occurs as a source and a target in the sequence $c_1, \ldots, c_j$, contradicting the choice of k. If w is the source of c_j we have a contradiction to the choice of i. This proves the claim. We can thus, without changing the outcome of the shuffle code move the operation c_i immediately before the operation c_k. Then our sequence contains consecutive copy operations $w \to x$ and $v \to w$. Replace these two operations by a cyclic shift of v, w and x and a copy operation $w \to v$. This decreases the number of copy operations by 1 and thus contradicts the minimality of t.

Altogether, in each case, we have either found a contradiction to the optimality of the shuffle code, to the minimality of the number of copy operations or we have succeeded in producing a shuffle code that has a longer prefix satisfying properties (i) and (ii). Inductively, we obtain a shuffle code satisfying both (i) and (ii). Fix such a code. Since no register is both source and target of a copy operation, the copy operations are commutative and can be reordered arbitrarily without changing the result.

For property (iii) first observe that the only way to transfer a value from u to v is via a copy operation $u \to v$. This is due to the fact that the shuffle code is correct, that no node occurs as both a source and a target of copy operations, and that π only permutes the values in the initial registers but does not duplicate them. Thus, for every edge there must be a corresponding copy operation. Conversely, this number of copy operations certainly suffices for a correct shuffle code for πG.

For property (iv) consider a copy operation from u to v such that u is not incident to a loop. If the in-degree of v in πG were 1, then there would be an incoming edge, which would correspond to a copy operation with target u, which is not possible by property (i). Thus, u has in-degree 0. But then, the contents of u are irrelevant and we can replace the copy from u to v by an operation that swaps the contents of u and v, resulting in a shuffle code with fewer copy operations.

By property (iv) every vertex that is the source of an edge in πG is incident to a loop. Hence $\sum_{v \in V} \max\{\deg_{\pi G}^+(v) - 1, 0\}$ is the number of non-loop edges in πG, which is the same as the number of copy operations by property (iii). Note that by definition π only permutes the out-degrees of the vertices, and hence $\sum_{v \in V} \max\{\deg_{\pi G}^+(v) - 1, 0\} = \sum_{v \in V} \max\{\deg_G^+(v) - 1, 0\}$. This shows property (iv) and finishes the proof. □

We call a shuffle code satisfying the conditions of Lemma 1 *normalized*. Observe that the number of copy operations used by a normalized shuffle code is a lower bound on the number of necessary copy operations since permutations, by definition, only permute values but never create copies of them.

Consider now an RTG G together with a normalized optimal shuffle code and one of its copy operations $u \to v$. Since the code is normalized, the value transferred to v by this copy operation is the one that stays there after the shuffle code has been executed. If v had no incoming edge in G, then we could shorten the shuffle by omitting the copy operation. Thus, v has an incoming edge (u', v) in G, and we associate the copy $u \to v$ with the edge (u', v) of G. In fact, $u' = \pi^{-1}(u)$, where $\pi = \pi_p \circ \cdots \circ \pi_1$. In this way, we associate every copy operation with an edge of the input RTG. In fact, this is an injective mapping by Lemma 1 (ii).

Lemma 2. *Let $((\pi_1, \ldots, \pi_p), (c_1, \ldots, c_t))$ be an optimal shuffle code S for an RTG $G = (V, E)$ and let $C \subseteq E$ be the edges that are associated with copies in S. Then*
  (i) *Every vertex v has $\max\{\deg_G^+(v) - 1, 0\}$ outgoing edges in C.*
 (ii) *$G - C$ is an outdegree-1 RTG.*
(iii) *$\pi_1, \ldots, \pi_p$ is an optimal shuffle code for $G - C$.*

Proof. For property (i) observe that, since permuting the register contents does not duplicate values, it is necessary that at least $\max\{\deg_G^+(v) - 1, 0\}$ of the edges of v are implemented by copy operations and thus are in C. By property (v) of Lemma 1 the number of copy operations is exactly the sum of these values, which immediately implies that equality holds at every vertex.

Property (ii) follows immediately from property (i).

Finally, for property (iii), suppose there is a shorter optimal shuffle code $\pi_1', \ldots, \pi_{p'}'$ with $p' < p$ for $G - C$. Let $\pi' = \pi_{p'}' \circ \cdots \circ \pi_1'$. Then $\pi'G$ has $|C|$ edges that are not loops and by creating a copy operation for each of them we obtain a shorter shuffle code. This is a contradiction to the optimality of the original shuffle code. Hence property (iii) holds. □

Lemma 2 shows that an optimal shuffle code for an RTG G can be found by first picking for each vertex one of its outgoing edges (if it has any) and removing the remaining edges from G, second finding an optimal shuffle code for the resulting outdegree-1 RTG, and finally creating one copy operation for each of the previously removed edges. Fig. 3 shows that the choice of the outgoing edges is crucial to obtain an optimal shuffle code.

In the following, we first show how to compute an optimal shuffle code for an outdegree-1 RTG in Section 3. Afterwards, in Section 4, we design an algorithm for efficiently determining a set of edges to be removed such that the resulting outdegree-1 RTG admits a shuffle code with the smallest number of operations.
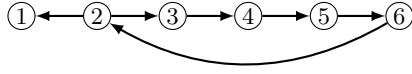
(a) The original RTG G needs one permutation and one copy operation.
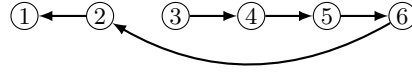
(b) After removing the edge $(2, 3)$, the RTG needs two permutation operations.

Fig. 3: The RTG G obtains the normalized optimal shuffle code (π_1, c_1), where $\pi_1 = (23456)$ and $c_1 = 3 \rightarrow 1$. However, after removing the edge $(2, 3)$ (instead of $(1, 2)$) we cannot achieve an optimal solution anymore.

3 Optimal Shuffle Code for Outdegree-1 RTGs

In this section we prove the optimality of the greedy algorithm proposed by Mohr et al. [8] for outdegree-1 RTGs. Before we formulate the algorithm, let us look at the effect of applying a transposition $\tau = (u\ v)$ to contiguous vertices of a k-cycle $K = (V_K, E_K)$ in a PRTG G, where k-cycle denotes a cycle of size k. Hence, $u, v \in V_K$ and $(u, v) \in E_K$. Then, in τG, the cycle K is replaced by a $(k-1)$-cycle and a vertex v with a loop. We say that τ has reduced the size of K by 1. If τK is trivial, we say that τ resolves K. It is easy to see that `permi5` reduces the size of a cycle by up to 4 and `permi23` reduces the sizes of two distinct cycles by 1 and up to 2, respectively. We can now formulate GREEDY as follows.

1. Complete each directed path of the input outdegree-1 RTG into a directed cycle, thereby turning the input into a PRTG.
2. While there exists a cycle K of size at least 4, apply a `permi5` operation to reduce the size of K as much as possible.
3. While there exist a 2-cycle and a 3-cycle, resolve them with a `permi23` operation.
4. Resolve pairs of 2-cycles by `permi23` operations.
5. Resolve triples of 3-cycles by pairs of `permi23` operations.

We claim that GREEDY computes an optimal shuffle code. Let G be an outdegree-1 RTG and let Q denote the set of paths and cycles of G. For a path or cycle $\sigma \in Q$, we denote by $\text{size}(\sigma)$ the number of vertices of σ. Define $X = \sum_{\sigma \in Q} \lfloor \text{size}(\sigma)/4 \rfloor$ and $a_i = |\{\sigma \in Q \mid \text{size}(\sigma) = i \mod 4\}|$ for $i = 2, 3$. We call the triple $\text{sig}(G) = (X, a_2, a_3)$ the *signature* of G.

Lemma 3. *Let G be an outdegree-1 RTG with* $\text{sig}(G) = (X, a_2, a_3)$. *The number* GREEDY$(G)$ *of operations in the shuffle code produced by the greedy algorithm is* GREEDY$(G) = X + \max\{\lceil (a_2 + a_3)/2 \rceil, \lceil (a_2 + 2a_3)/3 \rceil\}$.

Proof. After the first step we have a PRTG with the same signature as G. Clearly, GREEDY produces exactly X operations for reducing all cycle sizes below 4. Afterwards, only `permi23` operations are used to resolve the remaining cycles of size 2 and 3.

If $a_2 \geq a_3$, then first a_3 operations are used to resolve pairs of cycles of size 2 and 3. Afterwards, the remaining $a_2 - a_3$ cycles of size 2 are resolved by using $\lceil (a_2 - a_3)/2 \rceil$ operations. In total, these are $\lceil (a_2 + a_3)/2 \rceil$ operations.

If $a_3 \geq a_2$, then first a_2 operations are used to resolve pairs of cycles of size 2 and 3. Afterwards, the remaining $a_3 - a_2$ cycles of size 3 are resolved by using $\lceil 2(a_3 - a_2)/3 \rceil$ operations. In total, these are $\lceil (a_2 + 2a_3)/3 \rceil$ operations.

We observe that $(a_2 + a_3)/2 \leq (a_2 + 2a_3)/3$ holds if and only if $a_2 \leq a_3$ and that equality holds for $a_2 = a_3$. Since $\lceil \cdot \rceil$ is a monotone function, this implies that the total cost produced by the last part of the algorithm is $\max\{\lceil (a_2 + a_3)/2 \rceil, \lceil (a_2 + 2a_3)/3 \rceil\}$.

□

In particular, the length of the shuffle code computed by GREEDY only depends on the signature of the input RTG G. In the remainder of this section, we prove that GREEDY is optimal for outdegree-1 RTGs and therefore the formula in Lemma 3 actually computes the length of an optimal shuffle code.

Lemma 4. *Let G, G' be PRTGs with $\mathrm{sig}(G) = (X, a_2, a_3)$, $\mathrm{sig}(G') = (X', a_2', a_3')$ and $\mathrm{GREEDY}(G) - \mathrm{GREEDY}(G') \geq c$, and let $(\Delta_X, \Delta_2, \Delta_3) = \mathrm{sig}(G) - \mathrm{sig}(G')$. If $a_2 \geq a_3$, then $2\Delta_X + \Delta_2 + \Delta_3 \leq -2c + 1$. If $a_3 > a_2$, then $3\Delta_X + \Delta_2 + 2\Delta_3 \leq -3c + 2$.*

Proof. We assume that $\mathrm{GREEDY}(G) - \mathrm{GREEDY}(G') \geq c$ and start with the case that $a_2 \geq a_3$. By Lemma 3 and basic calculation rules for $\lceil \cdot \rceil$, we have the following.

$$\mathrm{GREEDY}(G) = X + \lceil (a_2 + a_3)/2 \rceil \leq X + (a_2 + a_3 + 1)/2$$
$$\mathrm{GREEDY}(G') \geq X' + \lceil (a_2' + a_3')/2 \rceil \geq X + \Delta_X + (a_2 + a_3 + \Delta_2 + \Delta_3)/2$$

Therefore, their difference computes to

$$\mathrm{GREEDY}(G) - \mathrm{GREEDY}(G') \leq -\Delta_X - (\Delta_2 + \Delta_3 - 1)/2$$
$$= -(2\Delta_X + \Delta_2 + \Delta_3 - 1)/2.$$

By assumption, we thus have $-(2\Delta_X + \Delta_2 + \Delta_3 - 1)/2 \geq c$, or equivalently $2\Delta_X + \Delta_2 + \Delta_3 \leq -2c + 1$.

Now consider the case $a_3 > a_2$. By Lemma 3, we have the following.

$$\mathrm{GREEDY}(G) = X + \lceil (a_2 + 2a_3)/3 \rceil \leq X + (a_2 + 2a_3 + 2)/3$$
$$\mathrm{GREEDY}(G') \geq X' + \lceil (a_2' + 2a_3')/3 \rceil \geq X + \Delta_X + (a_2 + 2a_3 + \Delta_2 + 2\Delta_3)/3$$

Similar to above, their difference computes to

$$\mathrm{GREEDY}(G) - \mathrm{GREEDY}(G') \leq -\Delta_X - (\Delta_2 + 2\Delta_3 - 2)/3$$
$$= -(3\Delta_X + \Delta_2 + 2\Delta_3 - 2)/3.$$

Similarly as above, by assumption we have $-(3\Delta_X + \Delta_2 + 2\Delta_3 - 2)/3 \geq c$, which is equivalent to $3\Delta_X + \Delta_2 + 2\Delta_3 \leq -3c + 2$. □

Lemma 4 gives us necessary conditions for when the GREEDY solutions of two RTGs differ by some value c. These necessary conditions depend only on the difference of the two signatures. To study them more precisely, we define $\Psi_1(\Delta_X, \Delta_2, \Delta_3) = 2\Delta_X + \Delta_2 + \Delta_3$ and $\Psi_2(\Delta_X, \Delta_2, \Delta_3) = 3\Delta_X + \Delta_2 + 2\Delta_3$. Next, we study the effect of a single transposition on these two functions.
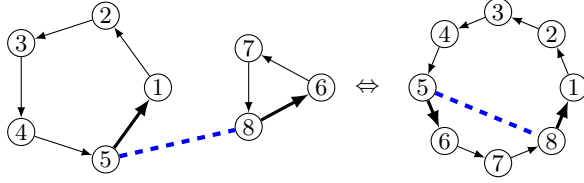
Fig. 4: The transposition $\tau = (5\ 8)$ acting on PRTGs. Affected edges are drawn thick. Read from left to right, the transposition is a merge; read from right to left, it is a split.

		0	1	2	3
0		$(0,0,0)$	$(0,0,0)$	$(0,0,0)$	$(0,0,0)$
1			$(0,1,0)$	$(0,-1,1)$	$(1,0,-1)$
2				$(1,-2,0)$	$(1,-1,-1)$
3					$(1,1,-2)$

(a) Signature change $(\Delta_X, \Delta_2, \Delta_3)$.

		0	1	2	3
0		0	0	0	0
1			1	0	1
2				0	0
3					1

		0	1	2	3
0		0	0	0	0
1			1	1	1
2				1	0
3					0

(b) Values of Ψ_1 (left) and Ψ_2 (right).

Table 1: Signature changes and Ψ values for merges. Row and column are the cycle sizes modulo 4 before the merge.

Let $G = (V, E)$ be a PRTG with $\mathrm{sig}(G) = (X, a_2, a_3)$ and let τ be a transposition of two elements in V. We distinguish cases based on whether the swapped elements are in different connected components or not. In the former case, we say that τ is a *merge*, in the latter we call it a *split*; see Fig. 4 for an illustration.

We start with the merge operations as they are a bit simpler. When merging two cycles of size s_1 and s_2, respectively, they are replaced by a single cycle of size $s_1 + s_2$. Note that removing the two cycles may decrease the values a_2 and a_3 of the signature by at most 2 in total. On the other hand, the new cycle can potentially increase one of these values by 1. The value X never decreases, and it increases by 1 if and only if s_1 mod $4 + s_2$ mod $4 \geq 4$. Table 1a shows the possible signature changes $(\Delta_X, \Delta_2, \Delta_3)$ resulting from a merge. The entry in row i and column j shows the result of merging two cycles whose sizes modulo 4 are i and j, respectively. Table 1b shows the corresponding values of Ψ_1 and Ψ_2. Only entries with $i \leq j$ are shown, the remaining cases are symmetric.

Lemma 5. *Let G be a PRTG with $\mathrm{sig}(G) = (X, a_2, a_3)$ and let τ be a merge. Then* GREEDY$(G) \leq$ GREEDY(τG).

Proof. Suppose GREEDY$(\tau G) <$ GREEDY(G). Then GREEDY$(G) -$ GREEDY$(\tau G) \geq 1$ and by Lemma 4 either $\Psi_1 \leq -1$ or $\Psi_2 \leq -1$. However, Table 1b shows the values of Ψ_1 and Ψ_2 for all possible merges. In all cases it is $\Psi_1, \Psi_2 \geq 0$. A contradiction. $\square$

In particular, the lemma shows that merges never decrease the cost of the greedy solution, even if they were for free. We now make a similar analysis for splits. It is, however, obvious that splits indeed may decrease the cost of greedy solutions. In fact, one can always split cycles in a PRTG until it is trivial.

Fig. 5: Transition graphs for Ψ_1 (left) and Ψ_2 (right).

First, we study again the effect of splits on the signature change $(\Delta_X, \Delta_2, \Delta_3)$. Since a split is an inverse of a merge, we can essentially reuse Table 1a. If merging two cycles whose sizes modulo 4 are i and j, respectively, results in a signature change of $(\Delta_X, \Delta_2, \Delta_3)$, then, conversely, we can split a cycle whose size modulo 4 is $i + j$ into two cycles whose sizes modulo 4 are i and j, respectively, such that the signature change is $(-\Delta_X, -\Delta_2, -\Delta_3)$, and vice versa. Note that given a cycle whose size modulo 4 is s one has to look at all cells (i, j) with $i + j \equiv s \pmod 4$ to consider all the possible signature changes. Since Ψ_1, Ψ_2 are linear, negating the signature change also negates the corresponding value. Thus, we can reuse Table 1b for splits by negating each entry.

Lemma 6. *Let $G = (V, E)$ be a PRTG and let π be a cyclic shift of c vertices in V. Let further $(\Delta_X, \Delta_2, \Delta_3)$ be the signature change affected by π. Then $\Psi_1(\Delta_X, \Delta_2, \Delta_3) \geq -\lceil(c-1)/2\rceil$ and $\Psi_2(\Delta_X, \Delta_2, \Delta_3) \geq -\lceil(3c - 3)/4\rceil$.*

Proof. We can write $\pi = \tau_{c-1} \circ \cdots \circ \tau_1$ as a product of $c - 1$ transpositions such that any two consecutive transpositions τ_i and τ_{i+1} affect a common element for $i = 1, \ldots, c-1$.

Each transposition decreases Ψ_1 (or Ψ_2) by at most 1, but a decrease happens only for certain split operations. However, it is not possible to reduce Ψ_1 (or Ψ_2) with every single transposition since for two consecutive splits the second has to split one of the connected components resulting from the previous splits. To get an overview of the sequences of splits that reduce the value of Ψ_1 (or of Ψ_2) by 1 for each split, we consider the following transition graphs T_k for Ψ_k ($k = 1, 2$) on the vertex set $S = \{0, 1, 2, 3\}$. In the graph T_k there is an edge from i to j if there is a split that splits a component of size $i \mod 4$ such that one of the resulting components has size $j \mod 4$ and this split decreases Ψ_k by 1. The transition graphs T_1 and T_2 are shown in Fig. 5.

For Ψ_1 the longest path in the transition graph has length 1. Thus, the value of Ψ_1 can be reduced at most every second transposition and $\Psi_1(\Delta_X, \Delta_2, \Delta_3) \geq -\lceil(c - 1)/2\rceil$.

For Ψ_2 the longest path has length 3 (vertex 1 has out-degree 0). Therefore, after at most three consecutive steps that decrease Ψ_2, there is one that does not. It follows that at least $\lfloor(c-1)/4\rfloor$ operations do not decrease Ψ_2, and consequently at most $\lceil(3c-3)/4\rceil$ operations decrease Ψ_2 by 1. Thus, $\Psi_2(\Delta_X, \Delta_2, \Delta_3) \geq -\lceil(3c - 3)/4\rceil$. □

Since `permi5` performs a single cyclic shift and `permi23` is the concatenation of two cyclic shifts, Lemmas 6 and 4 can be used to show that no such operation may decrease the number of operations GREEDY has to perform by more than 1.

Corollary 1. *Let G be a PRTG and let π be an operation, i.e., either a* `permi23` *or a* `permi5`. *Then* GREEDY$(G) \leq$ GREEDY$(\pi G) + 1$.

Proof. Assume for a contradiction that $\text{GREEDY}(G) > \text{GREEDY}(\pi G) - 1$. By Lemma 4 we have that either $\Psi_1(\Delta_X, \Delta_2, \Delta_3) \leq -3$ or $\Psi_2(\Delta_X, \Delta_2, \Delta_3) \leq -4$.

We distinguish cases based on whether π is a `permi5` or a `permi23`. If π is a `permi5`, then it is a c-cycle with $c \leq 5$. By Lemma 6, we have that $\Psi_1(\Delta_X, \Delta_2, \Delta_3) \geq -2$ and $\Psi_2(\Delta_X, \Delta_2, \Delta_3) \geq -3$. This contradicts the above bounds from Lemma 4.

If π is a `permi23`, then it is a composition of a 2-cycle and a c-cycle with $c \leq 3$. According to Lemma 6, both cycles contribute at least -1 to Ψ_1, and at least -1 and -2 to Ψ_2. Therefore, we have $\Psi_1(\Delta_X, \Delta_2, \Delta_3) \geq -2$ and $\Psi_2(\Delta_X, \Delta_2, \Delta_3) \geq -3$. This is again a contradiction. $\qquad\square$

Using this corollary and an induction on the length of an optimal shuffle code, we show that GREEDY is optimal for PRTGs; if no operation reduces the number of operations GREEDY needs by more than 1, why not use the operation suggested by GREEDY?

Theorem 1. *Let G be a PRTG. An optimal shuffle code for G takes* $\text{GREEDY}(G)$ *operations. Algorithm* GREEDY *computes an optimal shuffle code in linear time.*

Proof. The proof is by induction on the overall length of an optimal shuffle code. Clearly, GREEDY computes optimal shuffle codes for all instances that have a shuffle code of length 0.

Assume that G admits an optimal shuffle code of length $k + 1$. We show that $\text{GREEDY}(G) = k + 1$. First of all, note that $\text{GREEDY}(G) \geq k + 1$ as it computes a shuffle code of length $\text{GREEDY}(G)$. Let $\pi_1, \ldots, \pi_{k+1}$ be a shuffle code for G. Then obviously $\pi_{k+1} G$ admits an optimal shuffle code of length k, and therefore $\text{GREEDY}(\pi_{k+1} G) = k$ by our inductive assumption. Corollary 1 implies $\text{GREEDY}(G) \leq \text{GREEDY}(\pi_{k+1} G) + 1 = k + 1$; the induction hypothesis is proved.

Clearly, algorithm GREEDY indeed computes a correct, and thus optimal, shuffle code. It can easily be implemented to run in linear time. $\qquad\square$

Moreover, since merge operations may not decrease the cost of GREEDY and any PRTG that can be formed from the original outdegree-1 RTG G by inserting edges can be obtained from the PRTG G' formed by GREEDY and a sequence of merge operations, it follows that the length of an optimal shuffle for G is $\text{GREEDY}(G')$.

Lemma 7. *Let G be an outdegree-1 RTG and let G' be the PRTG formed by completing each directed path into a directed cycle. Then the length of an optimal shuffle code of G is* $\text{GREEDY}(G')$.

Proof. Assume $\pi_1, \ldots, \pi_k$ is an optimal shuffle code for G. Of course, applying $\pi = \pi_k \circ \cdots \circ \pi_1$ to G maps every value of G somewhere, that is, $\pi_1, \ldots, \pi_k$ is actually an optimal shuffle code for some instance G'' that consists of a disjoint union of directed cycles and contains G as a subgraph. It is not hard to see that G'' can be obtained from G' by a sequence of merge operations $\tau_1, \ldots, \tau_t$, i.e., $G'' = \tau_t \circ \cdots \circ \tau_1 G'$. Lemma 5 implies that $\text{GREEDY}(G') \leq \text{GREEDY}(\tau_1 G') \leq \cdots \leq \text{GREEDY}(\tau_t \circ \cdots \circ \tau_1 G') = \text{GREEDY}(G'') = k$, where the last equality follows from Theorem 1, the optimality of GREEDY for PRTGs. $\qquad\square$

By combining Theorem 1 and Lemma 7, we obtain the main result of this section.

Theorem 2. *Let G be an outdegree-1 RTG. Then an optimal shuffle code for G requires* $\text{GREEDY}(G)$ *operations.* GREEDY *computes such a shuffle code in linear time.*

4 The General Case

In this section we study the general case. A *copy set* of an RTG $G = (V, E)$ is a set $C \subseteq E$ such that $G - C = (V, E - C)$ is an outdegree-1 RTG and $|C| = \sum_{v \in V} \max\{\deg^+(v) - 1, 0\}$. We denote by $\mathcal{C}(G)$ the set of all copy sets of G. According to Lemma 2 an optimal shuffle code for G can be found by finding a copy set $C \in \mathcal{C}(G)$ such that the outdegree-1 RTG $G - C$ admits a shortest shuffle code. By Theorem 2 an optimal shuffle code for $G - C$ can be computed with the greedy algorithm and its length can be computed according to Lemma 3.

We thus seek a copy set $C \in \mathcal{C}(G)$ that minimizes the cost function $\text{GREEDY}(G - C) = X + \max\{\lceil (a_2 + a_3)/2 \rceil, \lceil (a_2 + 2a_3)/3 \rceil\}$, where (X, a_2, a_3) is the signature of $G - C$. Such a copy set is called *optimal*. Clearly, this is equivalent to minimizing the function

$$\text{GREEDY}'(G - C) = X + \max\{\frac{a_2 + a_3}{2}, \frac{a_2 + 2a_3}{3}\} = \begin{cases} X + \frac{a_2}{2} + \frac{a_3}{2} & \text{if } a_2 \geq a_3 \\ X + \frac{a_2}{3} + \frac{2a_3}{3} & \text{if } a_2 < a_3 \end{cases}$$

To keep track of which case is used for evaluating GREEDY', we define $\text{diff}(G - C) = a_2 - a_3$ and compute for each of the two function parts and every possible value d a copy set C_d with $\text{diff}(G - C_d) = d$ that minimizes that function.

More formally, we define $\text{cost}^1(G - C) = X + \frac{1}{2}a_2 + \frac{1}{2}a_3$ and $\text{cost}^2(G - C) = X + \frac{1}{3}a_2 + \frac{2}{3}a_3$ and we seek two tables $T_G^1[\cdot], T_G^2[\cdot]$, such that $T_G^i[d]$ is the smallest cost $\text{cost}^i(G - C)$ that can be achieved with a copy set $C \in \mathcal{C}(G)$ with $\text{diff}(G - C) = d$. We observe that $T_G^i[d] = \infty$ for $d < -n$ and for $d > n$. The following lemma shows that the length of an optimal shuffle code can be computed from these two tables.

Lemma 8. *Let $G = (V, E)$ be an RTG. The length of an optimal shuffle code for G is* $\sum_{v \in V} \max\{\deg^+(v) - 1, 0\} + \min\{\min_{d \geq 0} \lceil T_G^1[d] \rceil, \min_{d < 0} \lceil T_G^2[d] \rceil\}$.

Proof. Let $m = \sum_{v \in V} \max\{\deg^+(v) - 1, 0\}$. Consider an optimal normalized shuffle code for G, which, according to Lemma 2, consists of a copy set $C \subseteq E$ and a sequence of k permutation operations, i.e., the length of the shuffle code is $m + k$. Let (X, a_2, a_3) denote the signature of $G - C$ and let $d = a_2 - a_3$. If $a_2 \geq a_3$, or equivalently $d \geq 0$, then according to Theorem 2, we have $k = \text{GREEDY}(G) = X + \lceil (a_2 + a_3)/2 \rceil = \lceil X + (a_2 + a_3)/2 \rceil = \lceil \text{cost}^1(G - C) \rceil$, and therefore the length of the shuffle code is at most $m + T_G^1[d]$. If $a_2 < a_3$, i.e., if $d < 0$, then we have $k = \text{GREEDY}(G) = X + \lceil (a_2 + 2a_3)/3 \rceil = \lceil X + (a_2 + 2a_3)/3 \rceil = \lceil \text{cost}^2(G - C) \rceil$, and therefore the length of the shuffle code is at most $m + T_G^2[d]$. In either case the length of the shuffle code is bounded by the expression given in the statement of the theorem.

Conversely, assume that the minimum of the expression is obtained for some value $T_G^i[d]$. If $d < 0$ (resp. if $d \geq 0$), there exists a copy set C such that $\text{sig}(G - C) = (X, a_2, a_3)$ and $\text{GREEDY}(G - C) = \lceil \text{cost}^2(G - C) \rceil$ (resp. $\text{GREEDY}(G - C) = \lceil \text{cost}^2(G - C) \rceil$) is at most $T_G^1[d]$ (resp. at most $T_G^2[d]$). Then, clearly, the shuffle code defined by C and GREEDY applied to $G - C$ has length at most $m + \lceil T_G^2[d] \rceil$ (resp. $m + \lceil T_G^1[d] \rceil$). $\qquad\square$

In the following, we show how to compute for an RTG G a table $T_G[\cdot]$ with

$$T_G[d] = \min_{\substack{C \in \mathcal{C}(G) \\ \text{diff}(G-C)=d}} \text{cost}(G-C)$$

for an arbitrary cost function $\text{cost}(G-C) = c(\text{sig}(G-C))$, where c is a linear function. This is done in several steps depending on whether G is disconnected, is a tree, or is connected and contains a cycle. Before we continue, we introduce several preliminaries to simplify the following calculations. We denote by P_s a directed path on s vertices.

Definition 1. *A map f that assigns a value to an outdegree-1 RTG is* signature-linear *if there exists a linear function $g\colon \mathbb{R}^3 \to \mathbb{R}$ such that $f(G) = g(\text{sig}(G))$ for every outdegree-1 RTG G. For a signature-linear function f, $\Delta_f(s) = f(P_{s+1}) - f(P_s)$ is the* correction term.

Note that both $\text{cost} = c \circ \text{sig}$ and $\text{diff} = d \circ \text{sig}$ with $d(X, a_2, a_3) = a_2 - a_3$ are signature-linear. The correction term $\Delta_f(s)$ describes the change of f when the size of one connected component is increased from s to $s+1$.

Lemma 9. *Let f be a signature-linear function. Then the following hold:*
  (i) $f(G_1 \cup G_2) = f(G_1) + f(G_2)$ *for disjoint outdegree-1 RTGs G_1, G_2,*
  (ii) Let $G = (V, E)$ be an outdegree-1 RTG and let $v \in V$ with in-degree 0. Denote by s the size of the connected component containing v and let $G^+ = (V \cup \{u\}, E \cup \{(u,v)\})$ where u is a new vertex. Then $f(G^+) = f(G) + \Delta_f(s)$.

Proof. For (i) observe that $\text{sig}(G_1 \cup G_2) = \text{sig}(G_1) + \text{sig}(G_2)$; then the statement follows from the signature-linearity of f.

For (ii) observe that by adding u, we replace a connected component of size s by one of size $s+1$. Thus $\text{sig}(G^+) = \text{sig}(G) - \text{sig}(P_s) + \text{sig}(P_{s+1})$. The statement follows from the signature-linearity of f and the definition of $\Delta_f(s)$. $\qquad\square$

Note that $\Delta_f(s) = \Delta_f(s+4)$ for all values of s and hence it suffices to know the size of the enlarged component modulo 4.

The main idea for computing table $T_G[\cdot]$ by dynamic programming is to decompose G into smaller edge-disjoint subgraphs $G = G_1 \cup \cdots \cup G_k$ such that the copy sets of G can be constructed from copy sets for each of the G_i. We call such a decomposition *proper partition* if for every vertex v of G there exists an index i such that G_i contains all outgoing edges of v. Let $G_1, \ldots, G_k$ be a proper partition of G and let $\mathcal{C}_i \subseteq \mathcal{C}(G_i)$ for $i = 1, \ldots, k$. We define $\mathcal{C}_1 \otimes \cdots \otimes \mathcal{C}_k = \{C_1 \cup \cdots \cup C_k \mid C_i \in \mathcal{C}_i, i = 1, \ldots, k\}$. It is not hard to see that $\mathcal{C}(G_1 \cup \cdots \cup G_k) = \mathcal{C}(G_1) \otimes \cdots \otimes \mathcal{C}(G_k)$.

4.1 Disconnected RTGs

We start with the case that G is disconnected and consists of connected components $G_1, \ldots, G_k$, which form a proper partition of G. The main issue is to keep track of diff and cost. For an RTG G, we define $\mathcal{C}(G; d) = \{C \in \mathcal{C}(G) \mid \text{diff}(G - C) = d\}$. By Lemma 9(i) and the signature-linearity of diff, if $C_i \in \mathcal{C}(G_i; d_i)$ for $i = 1, 2$, then $C_1 \cup C_2 \in \mathcal{C}(G_1 \cup G_2; d_1 + d_2)$. This leads to the following lemma.

13

Lemma 10. *Let G be an RTG and let G_1, G_2 be vertex-disjoint RTGs. Then*
(i) $\mathcal{C}(G) = \bigcup_d \mathcal{C}(G; d)$ and
(ii) $\mathcal{C}(G_1 \cup G_2; d) = \bigcup_{d'} (\mathcal{C}(G_1; d') \otimes \mathcal{C}(G_2; d - d'))$.

Proof. Equation (i) follows immediately from the definition of $\mathcal{C}(G; d)$. For Equation (ii) observe that if $C_1 \in \mathcal{C}(G_1; d')$ and $C_2 \in \mathcal{C}(G_2; d - d')$, then $C = C_1 \cup C_2$ is a copy set of G and by Lemma 9(i) $\mathrm{diff}(G - C) = \mathrm{diff}((G_1 - C_1) \cup (G_2 - C_2)) = \mathrm{diff}(G_1 - C_1) + \mathrm{diff}(G_2 - C_2) = d' + d - d' = d$, and hence $C_1 \cup C_2 \in \mathcal{C}(G; d)$. Conversely, if $C \in \mathcal{C}(G; d)$, define $C_i = C \cap E_i$ where E_i is the edge set of G_i for $i = 1, 2$. Let $d' = \mathrm{diff}(G_1 - C_1)$. As above, it follows from Lemma 9(i) that $d = \mathrm{diff}(G - C) = \mathrm{diff}(G_1 - C_1) + \mathrm{diff}(G_2 - C_2) = d' + \mathrm{diff}(G - C)$, and hence $\mathrm{diff}(G - C) = d - d'$. Thus $C \in \mathcal{C}(G_1; d') \otimes \mathcal{C}(G_2; d - d')$. $\square$

By further exploiting the signature-linearity of cost, we also get $\mathrm{cost}((G_1 \cup G_2) - (C_1 \cup C_2)) = \mathrm{cost}(G_1 - C_1) + \mathrm{cost}(G_2 - C_2)$, allowing us to compute the cost of copy sets formed by the union of copy sets of vertex-disjoint graphs.

Lemma 11. *Let G_1, G_2 be two vertex-disjoint RTGs and let $G = G_1 \cup G_2$. Then $T_G[d] = \min_{d'}\{T_{G_1}[d'] + T_{G_2}[d - d']\}$.*

Proof. Applying the definition of $T_G[\cdot]$ as well as Lemma 10 (ii) and Lemma 9 (i) yields

$$
T_G[d] = \min_{C \in \mathcal{C}(G;d)} \mathrm{cost}(G - C) = \min_{C \in \bigcup_{d'}(\mathcal{C}(G_1;d') \otimes \mathcal{C}(G_2;d-d'))} \mathrm{cost}(G - C)
$$

$$
= \min_{d'}\left\{ \min_{C \in \mathcal{C}(G_1;d') \otimes \mathcal{C}(G_2;d-d')} \mathrm{cost}(G - C)\right\}
$$

$$
= \min_{d'}\left\{ \min_{C_1 \in \mathcal{C}(G_1;d')} \mathrm{cost}(G_1 - C_1) + \min_{C_2 \in \mathcal{C}(G_2;d-d')} \mathrm{cost}(G_2 - C_2)\right\}
$$

$$
= \min_{d'}\{T_{G_1}[d'] + T_{G_2}[d - d']\}. \qquad \square
$$

By iteratively applying Lemma 11, we compute $T_G[\cdot]$ for a disconnected RTG G with an arbitrary number of connected components. In the following, we will analyze the running time needed for the combination of all tables $T_{G_i}[\cdot]$ for the components G_i of G.

Lemma 12. *Let G be an RTG with n vertices and connected components $G_1, \ldots, G_k$. Given the tables $T_{G_i}[\cdot]$ for $i = 1, \ldots, k$, the table $T_G[\cdot]$ can be computed in $O(n^2)$ time.*

Proof. Let n_i denote the number of vertices of G_i. For two graphs H_1 and H_2 with h_1 and h_2 vertices, respectively, computing $T_{H_1 \cup H_2}[\cdot]$ according to Lemma 11 takes time $O(h_1 \cdot h_2)$ and the table size is $O(h_1 + h_2)$. Thus, iteratively combining the table for G_{i+1} with the table for $\bigcup_{j=0}^{i} G_j$ takes time $O(\sum_{i=1}^{k-1} n_{i+1} \sum_{j=1}^{i} n_j)$. It is $\sum_{i=1}^{k-1} n_{i+1} \sum_{j=1}^{i} n_j \leq \sum_{i=1}^{k-1} n_{i+1} n = n \sum_{i=1}^{k-1} n_{i+1} \leq n^2$. Hence, the running time is $O(n^2)$. $\square$

14

4.2 Tree RTGs

For a tree RTG G, we compute $T_G[\cdot]$ in a bottom-up fashion. The direction of the edges naturally defines a unique root vertex r that has no incoming edges and we consider G as a rooted tree. For a vertex v, we denote by $G(v)$ the subtree of G with root v. Let v be a vertex with children $v_1, \ldots, v_k$.

How does a copy set C of $G(v)$ look like? Clearly, $G(v) - C$ contains precisely one of the outgoing edges of v, say (v, v_j). Then $Z_j = \{(v, v_i) \mid i \neq j\} \subseteq C$. Graph $G(v) - Z_j$ has connected components $G(v_i)$ for $i \neq j$, whose union we denote $G_{\neg j}$, and one additional connected component $G^+(v_j)$ that is obtained from $G(v_j)$ by adding the vertex v and the edge (v, v_j). This forms a proper partition of $G(v) - Z_j$. As above, we decompose the copy set $C - Z_j$ further into a union of a copy set $C_{\neg j}$ of $G_{\neg j}$ and a copy set C_j of $G^+(v_j)$. Graph $G_{\neg j}$ is disconnected and can be handled as above. Note that the only child of the root of $G^+(v_j)$ is v_j and hence C_j is a copy set of $G(v_j)$.

For expressing the cost and difference measures for copy sets of $G^+(v_j)$ in terms of copy sets of $G(v_j)$, we use the correction terms Δ_{cost} and Δ_{diff}. By Lemma 9 (ii), $\mathrm{diff}(G^+(v_j) - C_j) = \mathrm{diff}(G(v_j) - C_j) + \Delta_{\mathrm{diff}}(s)$, where s is the size of the *root path* $P(v_j, C_j)$ of $G(v_j) - C_j$, i.e., the size of the connected component of $G(v_j) - C_j$ containing v_j. An analogous statement holds for cost. More precisely, it suffices to know s modulo 4. Therefore, we further decompose our copy sets as follows, which allows us to formalize our discussion.

Definition 2. *For a tree RTG G with root v and children $v_1, \ldots, v_k$, we define* $\mathcal{C}(G; d, s) = \{C \in \mathcal{C}(G; d) \mid |P(v, C)| \equiv s \pmod 4\}$. *We further decompose these by* $\mathcal{C}(G; d, s, j) = \{C \in \mathcal{C}(G; d, s) \mid (v, v_j) \notin C\}$, *according to which outgoing edge of the root is not in the copy set.*

The following lemma gives calculation rules for composing copy sets.

Lemma 13. *Let G be a tree RTG with root v and children $v_1, \ldots, v_k$ and for a fixed vertex v_j, $1 \leq j \leq k$, let $G^+(v_j)$ be the subgraph of G induced by the vertices in $G(v_j)$ together with v. Let further $G_{\neg j} = \bigcup_{i=1, i \neq j}^{k} G(v_i)$ and $Z_j = \{(v, v_i) \mid i \neq j\}$. Then*

  (i) $\mathcal{C}(G; d) = \bigcup_{s=0}^{3} \mathcal{C}(G; d, s)$ and $\mathcal{C}(G; d, s) = \bigcup_{j=1}^{k} \mathcal{C}(G; d, s, j)$.
  (ii) $\mathcal{C}(G^+(v_j); d, s) = \mathcal{C}(G(v_j); d - \Delta_{\mathrm{diff}}(s), s - 1)$.
  (iii) $\mathcal{C}(G; d, s, j) = \bigcup_{d'} (\mathcal{C}(G_{\neg j}; d') \otimes \mathcal{C}(G^+(v_j); d - d', s) \otimes \{Z_j\})$.

Proof. The statements in (i) follow immediately from the definitions of $\mathcal{C}(G; d, s)$ and $\mathcal{C}(G; d, s, j)$. We continue with Statement (ii). Since v in $G^+(v_j)$ has only one child v_j, the edge (v, v_j) is not in any copy set of $G^+(v_j)$. Therefore, the copy sets of $\mathcal{C}(G^+(v_j))$ and $\mathcal{C}(G(v_j))$ are in one-to-one correspondence. We need to understand how the partition into copy sets with difference measure d and root path length s (modulo 4) respects this bijection. Let s be the root path size of $G^+(v_j) - C$ for a copy set $C \in \mathcal{C}(G^+(v_j))$. Obviously, $|P(G(v_j) - C)| = |P(G^+(v_j) - C)| - 1 = s - 1$. Moreover, going from $G^+(v_j) - C$ to $G(v_j) - C$ replaces a connected component of size s by one of size $s - 1$. Therefore $\mathrm{sig}(G(v_j) - C) = \mathrm{sig}(G^+(v_j) - C) - \mathrm{sig}(P_s) + \mathrm{sig}(P_{s+1})$. By the signature-linearity of diff, we have $\mathrm{diff}(G(v_j) - C) = \mathrm{diff}(G^+(v_j) - C) - \Delta_{\mathrm{diff}}(s)$. Note further that $\Delta_{\mathrm{diff}}(s) = \Delta_{\mathrm{diff}}(s + 4)$ for every value of s, and hence it suffices to

know $s \mod 4$. Overall, it follows that a copy set $C \in \mathcal{C}(G^+(v_j); d, s)$ is a copy set of $G(v_j)$ with difference measure $\text{diff}(G^+(v_j) - C) - \Delta_{\text{diff}}(s)$ and root path size modulo 4 being $s - 1$. Thus $C \in C(G(v_j), d - \Delta_{\text{diff}}(s), s - 1)$. And conversely $C \in C(G(v_j), d - \Delta_{\text{diff}}(s), s - 1)$ satisfies $C \in C(G^+(v_j); d, s)$.

Next, we consider Statement (iii). First, observe that the copy sets $\mathcal{C}$ of G whose root path starts with (v, v_j) are exactly those copy sets of G that contain all edges in Z_j. These sets correspond bijectively to copy sets of $G - Z_j$. Thus $\mathcal{C}(G; d, s, j) = \mathcal{C}(G - Z_j; d, s) \otimes \{Z_j\}$. Observe that $G - Z_j = G_{\neg j} \cup G^+(v_j)$ is a proper partition of $G - Z_j$. Furthermore, the root path of any copy set of this graph lies in $G^+(v_j)$. Therefore, Lemma 10 (ii) implies that $\mathcal{C}(G - Z_j; d, s) = \bigcup_{d'}(\mathcal{C}(G_{\neg j}; d') \otimes (\mathcal{C}(G(v_j)^+; d - d', s))$. Combining this with the previously derived description of $\mathcal{C}(G; d, s, j)$ yields Statement (iii). $\qquad \square$

To make use of this decomposition of copy sets, we extend our table T with an additional parameter s to keep track of the size of the root path modulo 4. We call the resulting table $\tilde{T}$. More formally, $\tilde{T}_v[d, s] = \min_{C \in \mathcal{C}(G(v); d, s)} \text{cost}(G(v) - C)$. It is not hard to see that $T_G[\cdot]$ can be computed from $\tilde{T}_r[\cdot, \cdot]$ for the root r of a tree RTG G.

Lemma 14. *Let G be a tree RTG with root r. Then $T_G[d] = \min_s \tilde{T}_r[d, s]$.*

Proof. Using the definitions of $T_G[\cdot]$ and $\tilde{T}_r[\cdot, \cdot]$, we obtain

$$T_G[d] = \min_{C \in \mathcal{C}(G; d)} \text{cost}(G - C) = \min_{s \in \{0, \dots, 3\}} \min_{C \in \mathcal{C}(G; d, s)} \text{cost}(G - C) = \min_{s \in \{0, \dots, 3\}} \tilde{T}_r[d, s].$$

$\qquad \square$

To compute $\tilde{T}_v[\cdot, \cdot]$ in a bottom-up fashion, we exploit the decompositions from Lemma 13 and the fact that we can update the cost function from $G(v_j) - C_j$ to $G^+(v_j) - C_j$ using the correction term Δ_{cost}. The proof is similar to that of Lemma 11.

Lemma 15. *Let G be a tree RTG, let v be a vertex of G with children $v_1, \dots, v_k$, and let $G(v_i) = (V_i, E_i)$ for $i = 1, \dots, k$. Let further $G_{\neg j} = (V_{\neg j}, E_{\neg j}) = \bigcup_{i=1, i \neq j}^{k} G(v_i)$. Then the following equation holds.*
$$\tilde{T}_v[d, s] = \min_{j \in \{1, \dots, k\}} \min_{d'} T_{G_{\neg j}}[d'] + \tilde{T}_{v_j}[d - d' - \Delta_{\text{diff}}(s), (s - 1) \mod 4] + \Delta_{\text{cost}}(s)$$

Proof. According to the definition of $\tilde{T}_v[d, s]$ and Lemma 13 (i), we find that

$$\tilde{T}_v[d, s] = \min_{C \in \mathcal{C}(G; d, s)} \text{cost}(G - C) = \min_j \min_{C \in \mathcal{C}(G; d, s, j)} \text{cost}(G - C) \qquad (1)$$

Using Lemma 13 (iii) yields

$$\min_{C \in \mathcal{C}(G; d, s, j)} \text{cost}(G - C) = \min_{d'} \min_{\substack{X \in \mathcal{C}(G_{\neg j}; d') \\ Y \in \mathcal{C}(G^+(v_j); d - d', s)}} \text{cost}(G - X - Y - Z_j). \qquad (2)$$

Note that $G - Z_j = G_{\neg j} \cup G^+(v_j)$. By Lemma 10, we have that for $X \in \mathcal{C}(G_{\neg j}; d'), Y \in \mathcal{C}(G^+(v_j); d - d', s)$, it is $\text{cost}(G - X - Y - Z_j) = \text{cost}(G_{\neg j} \cup G^+(v_j) - X - Y) =$

$\mathrm{cost}(G_{\neg j} - X) + \mathrm{cost}(G^+(v_j) - Y)$. Therefore,

$$
\min_{\substack{X \in \mathcal{C}(G_{\neg j}; d') \\ Y \in \mathcal{C}(G^+(v_j); d-d', s)}} \mathrm{cost}(G - X - Y - Z_j) \tag{3}
$$
$$
= \min_{X \in \mathcal{C}(G_{\neg j}; d')} \mathrm{cost}(G_{\neg j} - X) + \min_{Y \in \mathcal{C}(G^+(v_j); d-d', s)} \mathrm{cost}(G^+(v_j) - Y).
$$

By definition $\min_{X \in \mathcal{C}(G_{\neg j}; d')} \mathrm{cost}(G_{\neg j} - X) = T_{G_{\neg j}}[d']$. Furthermore, $G^+(v_j)$ is a tree RTG whose root v has the single child v_j. Hence, by Lemma 13 (ii) and Lemma 9 (ii), we find

$$
\min_{Y \in \mathcal{C}(G^+(v_j); d-d', s)} \mathrm{cost}(G^+(v_j) - Y)
$$
$$
= \min_{Y \in \mathcal{C}(G(v_j); d-d'-\Delta_{\mathrm{diff}}(s), s-1)} \mathrm{cost}(G(v_j) - Y) + \Delta_{\mathrm{cost}}(s) \tag{4}
$$
$$
= \tilde{T}_{v_j}[d - d' - \Delta_{\mathrm{diff}}(s), s - 1] + \Delta_{\mathrm{cost}}(s)
$$

Combining Equations 1–4 yields the claim. $\qquad \square$

For leaves v of a tree RTG G, $\tilde{T}_v[0, 1] = 0$ and all other entries are ∞. We compute $T_G[\cdot]$ by iteratively applying Lemma 15 in a bottom-up fashion, using Lemma 14 to compute $T[\cdot]$ from $\tilde{T}[\cdot, \cdot]$ in linear time when needed.

Lemma 16. *Let $G = (V, E)$ be a tree RTG with n vertices and root r. The tables $\tilde{T}_r[\cdot, \cdot]$ and $T_G[\cdot]$ can be computed in $O(n^3)$ time.*

Proof. First observe that given $\tilde{T}_v[\cdot, \cdot]$ for $v \in V$, table $T_{G(v)}[\cdot]$ can be computed in linear time according to Lemma 14. In particular, $T_G[\cdot]$ can be computed from $\tilde{T}_r[\cdot, \cdot]$ in linear time.

We now bound the computation time for $\tilde{T}_r[\cdot, \cdot]$. Let $v \in V$ with children $v_1, \dots, v_k$. Given the tables $\tilde{T}_{v_i}[\cdot, \cdot]$, we can compute $\tilde{T}_v[\cdot, \cdot]$ by Lemma 15. More precisely, for each $j = 1, \dots, k$, we first compute $T_{G_{\neg j}}[\cdot]$ in quadratic time by Lemma 12 followed by $O(n)$ table lookups, one for each value of d'. Hence, processing v takes time $O(\deg^+(v) \cdot n^2)$. Since $\sum_{v \in V} \deg^+(v) = n - 1$, the total processing time to compute $\tilde{T}_r[\cdot, \cdot]$ in a bottom-up fashion is $O(n^3)$. $\qquad \square$

4.3 Connected RTGs Containing a Cycle

We now look at connected RTGs that contain a cycle. We first introduce an additional decomposition for copy sets to simplify the following calculations.

Lemma 17. *Let $G = (V, E)$ be a connected RTG containing a directed cycle K and let $e_1, \dots, e_k$ denote the edges of K whose source has out-degree at least 2. Let further $O = \{(u, v) \in E \mid u \in K, (u, v) \notin K\}$. Then*

$$
\mathcal{C}(G; d) = \mathcal{C}(G - O; d) \otimes \{O\} \cup \bigcup_{i=1}^{k} \mathcal{C}(G - e_i; d) \otimes \{\{e_i\}\}.
$$

Proof. Every copy set $C \in \mathcal{C}(G; d)$ contains either some edge of K or it contains all edges in O. Note that edges of K that are not among $e_1, \ldots, e_k$ are not contained in any copy set. Thus, in the former case, $e_i \in C$ for some $i \in \{1, \ldots, k\}$ and hence $C \in \mathcal{C}(G - e_i; d) \otimes \{\{e_i\}\}$. In the latter case $C \setminus O$ is a copy set of $G - O$, hence $C \in \mathcal{C}(G - O; d) \otimes \{O\}$. Conversely, any copy set in $\mathcal{C}(G - O; d) \otimes \{O\}$ forms a copy set of G and also every copy set in $\mathcal{C}(G - e_i; d) \otimes \{\{e_i\}\}$ for any value of i forms a copy set of G. This finishes the proof. $\qquad \square$

As before, this decomposition can be used to efficiently compute $T_G[\cdot]$ from the tables of smaller subgraphs of a connected RTG G containing a cycle.

Lemma 18. *Let $G = (V, E)$ be a connected RTG containing a directed cycle K and let $e_1, \ldots, e_k$ denote the edges of K whose source has out-degree at least 2. Let further $O = \{(u, v) \in E \mid u \in K, (u, v) \notin K\}$. Then*

$$T_G[d] = \min \left\{ T_{G-O}[d], \min_{i=1}^{k} T_{G-e_i}[d] \right\}.$$

Proof. Using the definition of $T_G[\cdot]$ and Lemma 17, we find that

$$T_G[d] = \min_{C \in \mathcal{C}(G; d)} \mathrm{cost}(G - C)$$
$$= \min_{C \in (\mathcal{C}(G-O; d) \otimes \{O\}) \cup \bigcup_{i=1}^{k}(\mathcal{C}(G-e_i; d) \otimes \{\{e_i\}\})} \mathrm{cost}(G - C).$$

As we minimize cost over a union of sets, we can minimize it over the sets individually and then take the minimum of the results. Hence, we find that

$$\min_{C \in \mathcal{C}(G-O; d) \otimes \{O\}} \mathrm{cost}(G - C) = \min_{C \in \mathcal{C}(G-O; d)} \mathrm{cost}(G - O - C) = T_{G-O}[d]$$

and

$$\min_{C \in \mathcal{C}(G-e_i; d) \otimes \{\{e_i\}\}} \mathrm{cost}(G - C) = \min_{C \in \mathcal{C}(G-e_i; d)} \mathrm{cost}(G - e_i - C) = T_{G-e_i}[d],$$

which together yield the claim. $\qquad \square$

Lemma 19. *Let $G = (V, E)$ be a connected RTG containing a directed cycle. The table $T_G[\cdot]$ can be computed in $O(n^4)$ time.*

Proof. Let $e_1, \ldots, e_k$ be the edges of the cycle K. First, observe that $G - e_i$ is a tree for $i = 1, \ldots, k$. Hence, we can compute each table $T_{G-e_i}[\cdot]$ in $O(n^3)$ time by Lemma 16. Thus, computing all these tables takes $O(n^4)$ time.

Second, let $O = \{(u, v) \in E \mid u \in K, (u, v) \notin K\}$. The graph $G - O$ is the disjoint union of the cycle K and several tree RTGs $G_1, \ldots, G_t$. The table $T_K[\cdot]$ has only one finite entry and can be computed in constant time. The tables $T_{G_i}[\cdot]$ can be computed in $O(n^3)$ time. Using Lemma 12, we then compute $T_{G-O}[\cdot]$ in quadratic time.

With these tables available, we can compute $T_G[\cdot]$ according to Lemma 18. This takes $O(n^2)$ time. The overall running time is thus $O(n^4)$. $\qquad \square$

18

4.4 Putting Things Together

To compute $T_G[\cdot]$ for an arbitrary RTG G, we first compute $T_K[\cdot]$ for each connected component K of G using Lemmas 16 and 19. Then, we compute $T_G[\cdot]$ using Lemma 12 and the length of an optimal shuffle code using Lemma 8. To actually compute the shuffle code, we augment the dynamic program computing $T_G[\cdot]$ such that an optimal copy set C can be found by backtracking in the tables. An optimal shuffle code is then constructed by applying GREEDY to $G - C$ and adding one copy operation for each edge in C.

Theorem 3. *Given an RTG G, an optimal shuffle code can be computed in $O(n^4)$ time.*

Proof. We compute all tables $T_C[\cdot]$, where C is a connected component of G, in $O(n^4)$ time using Lemma 16 and Lemma 19. Using Lemma 12, we then compute $T_G[\cdot]$ in $O(n^2)$ time. From this, we can compute the length of an optimal shuffle code by Lemma 8.

In fact, it is not difficult to modify the dynamic program in a way that, given an entry $T_G[d]$, a corresponding copy set C of G with $\mathrm{cost}(G - C) = T_G[d]$ can be computed by backtracking in the tables. Hence, to compute an optimal shuffle code for G, we first compute an optimal copy set C_{opt} of G in $O(n^4)$ time. Then, we compute an optimal shuffle code $\pi_1, \ldots, \pi_k$ for $G - C_{\mathrm{opt}}$ using GREEDY, which takes linear time according to Theorem 2. Let $\pi = \pi_k \circ \ldots \circ \pi_1$. For each edge $(u, v) \in C_{\mathrm{opt}}$, we define a corresponding copy operation $\pi(u) \to v$. Let $c_1, \ldots, c_t$ be these copy operations in arbitrary order. Then the sequence $S = \pi_1, \ldots, \pi_k, c_1, \ldots, c_t$ is an optimal shuffle code. This can be seen as follows. First, by Lemma 8, the length of S is minimal. It remains to show that S is indeed a shuffle code for G. This is clearly true, as it first shuffles the values in the registers so that a subset of the values is in the correct position and then uses copy operations to transfer the remaining values to their destinations. $\square$

5 Conclusion

We have presented an efficient algorithm for generating optimal shuffle code using copy instructions and permutation instructions, which allow to arbitrarily permute the contents of up to five registers. As an intermediate result, we have proven the optimality of the greedy algorithm for factoring a permutation into a minimal product of permutations, each of which permutes up to five elements. It would be interesting to allow permutations of larger size.

References

1. Blazy, S., Robillard, B.: Live-range unsplitting for faster optimal coalescing. In: Languages, Compilers, and Tools for Embedded Systems (LCTES '09). pp. 70–79. ACM (2009)

2. Bouchez, F., Darte, A., Rastello, F.: On the complexity of register coalescing. In: Code Generation and Optimization (CGO '07). pp. 102–114. IEEE (2007)
3. Caprara, A.: Sorting by reversals is difficult. In: Computational Molecular Biology (RE-COMB'97). pp. 75–83. ACM (1997)
4. Farnoud, F., Milenkovic, O.: Sorting of permutations by cost-constrained transpositions. IEEE Transactions on Information Theory 58(1), 3–23 (2012)
5. Grund, D., Hack, S.: A fast cutting-plane algorithm for optimal coalescing. In: Krishnamurthi, S., Odersky, M. (eds.) Compiler Construction, Lecture Notes in Computer Science, vol. 4420, pp. 111–125. Springer Berlin Heidelberg (2007)
6. Hack, S.: Register Allocation for Programs in SSA Form. Ph.D. thesis, Universität Karlsruhe (2007), http://digbib.ubka.uni-karlsruhe.de/volltexte/documents/6532
7. Hack, S., Goos, G.: Copy coalescing by graph recoloring. SIGPLAN Notices 43(6), 227–237 (2008)
8. Mohr, M., Grudnitsky, A., Modschiedler, T., Bauer, L., Hack, S., Henkel, J.: Hardware acceleration for programs in SSA form. In: Compilers, Architecture and Synthesis for Embedded Systems (CASES '13). ACM (2013)
9. Seress, Á.: Permutation Group Algorithms, vol. 152. Cambridge University Press (2003)