

A Omitted proofs

Proof of Theorem 1. Let $f := (f_i)_{i=1}^k$ be the objective function vector, z be any solution, $Z = \{i : f_i(z) = 0\}$, if $|Z| = k$ then z is an extreme solution, so P 1-approximates z . Assume otherwise, we define $\lambda \in (0, 1]^k$ where $\lambda_i := \epsilon/\delta_i$ if $i \in Z$ and $\lambda_i := \epsilon/[(k - |Z|)f_i(z)]$ otherwise for some sufficiently small $\epsilon > 0$. By definition of sufficient solution set and Observation 1, there is $x \in P$ minimizing $\lambda^\top f$, i.e. $\lambda^\top f(x) \leq \lambda^\top f(z) = \epsilon$. If $f_i(x) > 0$ for some $i \in Z$ or $f_i(x) > (k - |Z|)f_i(z)$ for some $i \notin Z$, then since $f(x) \in \mathbb{R}_{\geq 0}^k$, we have $\lambda^\top f(x) > \epsilon$, a contradiction. Therefore, x , and by extension P , $(k - |Z|)$ -approximates z . Since z can assume positive values in all objectives⁶, this factor simplifies to k .

We show tightness by construction. Let $\epsilon \in (0, k)$, $m := k^2$, $\theta_i := \sum_{j=0}^{k-1} e_{ik-j}$ for $i = 1, \dots, k$ where e_j is the j th unit vector in \mathbb{R}^m , we define a non-negative k -objective instance over $\{0, 1\}^m$: $\min_x \{f(x) := (\theta_i^\top x - \epsilon \prod_{j=0}^{k-1} x_{ik-j})_{i=1}^k : |x| \geq k\}$. We see that the set of all supported solutions is precisely $S := \{\theta_i\}_{i=1}^k$. Let $z := \sum_{i=0}^{k-1} e_{ik+1}$ be a solution, for all $i = 1, \dots, k$, $f_i(\theta_i) = k - \epsilon \geq (k - \epsilon)f_i(z)$ (equality holds if $k > 1$). This means S fails to $(k - \epsilon - \epsilon)$ -approximate z for any $\epsilon > 0$, and ϵ can be arbitrarily small. Since S is a complete solution set, the claim follows. \square

Proof of Lemma 1. Let c be any point in $\text{Int}(Q)$, by definition of A , $w_i^{(c)} = w_j^{(c)}$ iff $\delta_{i,j} = 0$, and $w^{(c)}$ admits multiple minima iff they contain different elements among those sharing weights in $w^{(c)}$, while sharing all other elements. Indeed, let x and y be a pair of minima violating this condition, they must contain different sets of weights so for all bijection γ between $x \setminus y$ and $y \setminus x$, there is $u \in x \setminus y$ where $w_u^{(c)} \neq w_{\gamma(u)}^{(c)}$; this leads to a contradiction when combined with the base exchange property. This means these optima share image under w , and bases not having the same image do not minimize $w^{(c)}$.

Let b be any point on the boundary of Q and L be the set of points between b and c excluding endpoints, we show that π_c also sorts $w^{(b)}$. Let $i, j \in E$ where $w_i^{(c)} < w_j^{(c)}$, then $w_i^{(b)} > w_j^{(b)}$ implies $w_i^{(d)} = w_j^{(d)}$ for some $d \in L$, meaning L meets a hyperplane in A , a contradiction as $L \subseteq \text{Int}(Q)$. For all pairs $i, j \in E$ where $w_i^{(c)} = w_j^{(c)}$, $\delta_{i,j} = 0$ so $w_i^{(b)} = w_j^{(b)}$. With this, every pair is accounted for, so π_c sorts $w^{(b)}$. Therefore, since Greedy guarantees optimality, any base minimizing $w^{(c)}$ also minimizes $w^{(b)}$, yielding the claim. \square

Proof of Corollary 1. We see that $|A'| = |H_A|$, which is upper bounded by the number of half-space intersections from hyperplanes in A . Since these are $(k - 2)$ -dimensional hyperplanes, applying the formula in [32] gives $|H_A| \leq \sum_{i=1}^k \binom{|A|}{i-1}$ which is increasing in $|A|$, so the claim follows from $|A| \leq m(m - 1)/2$. We have A' is a sufficient trade-off set following from Lemma 1 and $\bigcup_{Q \in H_A} Q = U$. \square

Proof of Lemma 2. Let $0 < \lambda_c < \lambda_d < 1$ such that $b := (1 - b)a + ba' \in A^*$ and for all $\lambda \in [\lambda_c, \lambda_b)$, $(1 - \lambda)a + \lambda a' \notin A^*$, and let $c := (1 - \lambda_c)a + \lambda_c a'$, then elements sharing weight in $w^{(b)}$ must be mapped to consecutive positions in π_c . Indeed, let $p, q \in E$ ($\pi_c(p) < \pi_c(q)$) where $w_p^{(b)} = w_q^{(b)}$, if there is $o \in E$ where $\pi_c(o) \in (\pi_c(p), \pi_c(q))$ and $w_o^{(b)} \neq w_p^{(b)}$, then since the former implies $w_o^{(c)} \in (w_p^{(c)}, w_q^{(c)})$, we have $w_o^{(d)} = w_p^{(d)}$ or $w_o^{(d)} = w_q^{(d)}$ for some d in the open line segment connecting b and c which implies $d \in A^*$, a contradiction. Each such consecutive sequence of l positions contains $l(l - 1)/2$ pairs. From here, we consider two cases:

- If such a sequence contains no pair (i, j) where $\delta_{i,j} = 0$, then the aforementioned pairs correspond to $l(l - 1)/2$ duplicates of b in A^* . Furthermore, since the weights are transformed linearly w.r.t. trade-off, for all sufficiently small $\epsilon > 0$, these sequences are reversed between π_c and $\pi_{b+\epsilon(b-c)}$, whereas positions not in these sequences are stationary. Reversing l consecutive positions requires $l(l - 1)/2$ adjacent swaps, so the Kendall distance between π_c and $\pi_{b+\epsilon(b-c)}$ equals the multiplicity of b in A^* .

⁶If $|Z| \geq k'$ for all solutions z , the instance is reducible to $(k - k')$ -objective instances, and the guarantee factor is likewise tight.

- If such a sequence contains $h > 1$ elements with the same weight at all trade-off, then these must occupy consecutive positions in π_c . As we assumed, the relative ordering among these elements is fixed, so exactly $h(h-1)/2$ swaps are saved. Furthermore, any pair (i, j) among these elements is such that $\Delta_{i,j} \notin A$, meaning these $h(h-1)/2$ pairs are already subtracted from A^* .

In any case, we can assign to each duplicate of b in A^* a permutation sorting $w^{(b)}$ so that these form a sequence of adjacent swap from π_c to $\pi_{b+\epsilon(b-c)}$ including $\pi_{b+\epsilon(b-c)}$ and not π_c . This directly yields the claim if a and a^* are not in A^* .

Assume otherwise, then for all hyperplanes $\Delta_{i,j}$ containing a , $w_i^{(a)} = w_j^{(a)}$, so for every such pair (i, j) , we arrange π_a so that their pairwise ordering in π_a is the reverse of that in $\pi_{a'}$. We likewise give a' the same treatment⁷. With this, the Kendall distance between π_a and $\pi_{a'}$ is maximized and equal to $|A^*|$. \square

Proof of Lemma 3. Let $E_o := \{a \in E : \tau(a) < \tau(o)\}$ be the set of elements Greedy considers adding to x before $o \in E$ when run on τ , we have $x \cap E_u = x' \cap E_u$. If $v \in x$ or $v \notin x'$ or $u \notin x$ or $u \in x'$ then $x = x'$, as can be seen from how Greedy selects elements:

- If $v \in x$, then $v \in x'$ since Greedy observes v before u when run on τ' . Whether Greedy adds u to x only depends on whether there is a circuit in $(x \cap E_u) \cup \{u\} = (x' \cap E_u) \cup \{u\}$, so it makes the same decision when run on τ' . Afterwards, it proceeds identically on both permutations, leading to the same outcome, so $x = x'$. By symmetry, the same follows from $u \in x'$.
- If $u \notin x$, then there is a circuit in $(x \cap E_u) \cup \{u\} = (x' \cap E_u) \cup \{u\}$, so $u \notin x'$. By the same argument, Greedy makes the same decision regarding v on both permutations, leading to $x = x'$. By symmetry, the same follows from $v \notin x'$.

Assume otherwise, it is a known property of bases [20] that $x \cup \{v\}$ contains a unique circuit C and that $v \in C$. Greedy not adding v to x implies that $C \subseteq (x \cap E_v) \cup \{v\} = (x' \cap E_v) \cup \{u, v\}$. Let v' be the first element after v that x and x' differ at and assume w.l.o.g. $v' \in x \setminus x'$, we have $(x' \cap E_{v'}) \cup \{u\} = (x \cap E_{v'}) \cup \{v\}$ and since v' is not added into x' before Greedy terminates, there must be another circuit in $(x' \cap E_{v'}) \cup \{v'\} \subset x \cup \{v\}$ containing v' , which is distinct from the unique circuit C . The contradiction implies that x and x' do not differ after v , so $x \otimes x' = \{u, v\}$. \square

Proof of Theorem 2. We define $l_c := (1-c)a + cb$ for $c \in [0, 1]$, let $0 \leq \theta \leq \theta' \leq 1$ where π_{l_θ} and $\pi_{l_{\theta'}}$ are an adjacent swap apart⁸ and the Greedy solutions on them, x and x' , are such that $|x \otimes x'| = 2$. Let $u, v \in E$ where $x \cap \{u, v\} = \{u\}$ and $x' \cap \{u, v\} = \{v\}$, Lemma 3 implies $\pi_{l_\theta}(u) < \pi_{l_\theta}(v)$ and $\pi_{l_{\theta'}}(u) > \pi_{l_{\theta'}}(v)$, so $\pi_a(u) < \pi_a(v)$. This means as the trade-off moves from a to b , the Greedy solution minimizing the scalarized weight changes incrementally by having exactly one element shifted to the right on π_a (to a position not occupied by the current solution). Since at most $hm - h(h+1)/2$ such changes can be done sequentially, Greedy produces at most $hm - h(h+1)/2 + 1$ distinct solutions in total across all trade-offs between a and b .

To show this upper bound, we keep track of the following variables as the trade-off moves from a to b . Since each solution contains n elements, let p_i be the i th leftmost position on π_a among those occupied by the current Greedy solution for $i = 1, \dots, n$, we see that upon each change, there is at least a $j \in \{1, \dots, n\}$ where p_j increases. Furthermore, for all i , p_i can increase by at most $m - n$ since it cannot be outside of $[i, m - n + i]$, so the quantity $p := \sum_{i=1}^n p_i$ can increase by at most $n(m - n)$. We see that p increases by l when the change is incurred by a swap in the Greedy solution such that the added element is positioned l to the right of the removed element on π_a , we call this a l -move. Furthermore, each element pair participates in at most one swap, so p can be increased by at most $m - l$ l -moves for every $l = 1, \dots, m - 1$. Therefore, to upper bound the number of moves, we can assume smallest possible distance in each move, and the increase in p from using all possible l -moves for all $l = 1, \dots, h$ is $\sum_{j=1}^h j(m - j) \geq n(m - n)$. This means no more than $\sum_{j=1}^h (m - j) = hm - h(h+1)/2$ moves can be used to increase p by at most $n(m - n)$. \square

⁷This is also done for any $\Delta_{i,j}$ containing both a and a' .

⁸If $\theta = \theta'$, we assume π_{l_θ} is closer to π_a in Kendall distance.

Proof of Lemma 4. First, we observe that a set superset a base iff its rank is n . We see that for all $\lambda \in [0, 1]$ and $x, y \in \{0, 1\}^m$, $r(x) > r(y)$ implies $f_\lambda(x) < f_\lambda(y)$. Thus, for each $\lambda \in \Lambda$, MOEA/D performs (1+1)-EA search toward a base's superset with fitness f_λ , which concludes in $O(m \log n)$ expected steps [26]. The claim follows from the fact that MOEA/D produces $|\Lambda|$ search points in each step. \square

Proof of Theorem 3. We assume each solution in P_λ superset a base for all $\lambda \in \Lambda$; this occurs within expected time $O(|\Lambda|m \log n)$, according to Lemma 4. Since for each $\lambda \in \Lambda$, the best improvement in f_λ is retained in each step, the expected number of steps MOEA/D needs to minimize f_λ is at most the expected time (1+1)-EA needs to minimize f_λ . We thus fix a trade-off λ and assume the behaviors of (1+1)-EA. Note that we use $d_\lambda \cdot w^{(\lambda)}$ in the analysis instead for integral weights; we scale f_λ and OPT_λ accordingly.

We call the bit flips described in Lemma 5 *good flips*. Let s be the current search point, if good 1-bit flips incur larger total weight reduction than good 2-bit flips on s , we call s 1-step, and 2-step otherwise. If at least half the steps from s to the MWB z are 1-steps, Lemma 5 implies the optimality gap of s is multiplied by at most $1 - 1/2(m - n)$ on average after each good 1-bit flip. Therefore, from $f_\lambda(s) \leq d_\lambda(m - n)w_{max} + OPT_\lambda$, the expected difference D_L after L good 1-bit flips is $E[D_L] \leq d_\lambda(m - n)w_{max}(1 - 1/2(m - n))^L$. At $L = \lceil (2 \ln 2)(m - n) \log(2d_\lambda(m - n)w_{max} + 1) \rceil$, $E[D_L] \leq 1/2$ and by Markov's inequality and the fact that $D_L \geq 0$, $\Pr[D_L < 1] \geq 1/2$. Since weights are integral, $D_L < 1$ implies that z is reached. The probability of making a good 1-bit flip is $\Theta((m - n)/m)$, so the expected number of steps before L good 1-bit flips occur is $O(Lm/(m - n)) = O(m(\log(m - n) + \log w_{max} + \log d_\lambda))$. Since 1-steps take up most steps between s and z , the bound holds.

If at least half the steps from s to z are 2-steps, Lemma 5 implies the optimality gap of s is multiplied by at most $1 - 1/2n$ on average after each good 2-bit flip. Repeating the argument with $L = \lceil (2 \ln 2)n \log(2d_\lambda(m - n)w_{max} + 1) \rceil$ and the probability of making a good 2-bit flip being $\Theta(n/m^2)$, we get the bound $O(m^2(\log(m - n) + \log w_{max} + \log d_\lambda))$. Summing this over all $\lambda \in \Lambda$ gives the total bound. \square

Proof of Theorem 4. From Corollary 2, to collect a new point in C , it is sufficient to perform a 2-bit flip on some supported solution. In worst-case, there can be only one trade-off $\lambda \in \Lambda$ such that all non-extreme supported solutions minimize $w^{(\lambda)}$, so the correct solution is mutated with probability at least $1/l$ in each iteration, where l is the number of already collected points. Since $|\Lambda|$ search points are generated in each iteration, the expected number of search points required to enumerate C is $O(|\Lambda|m^2 \sum_{l=1}^{|C|} l) = O(|\Lambda||C|^2 m^2)$. \square