SpOC 4 Challenges 1 Beginner: Permutation Encodings and Challenge 3: Simulated Annealing for the Luna Tomato Advertising

Mingxuan Li
mingxuan.li@se23.qmul.ac.uk
BUPT International School,
Queen Mary University of London
London, United Kingdom

Thomas Weise
tweise@hfuu.edu.cn
School of Artificial Intelligence and Big Data,
Hefei University
Hefei, Anhui, China

Sina Abdipoor
sinaa@sunway.edu.my
Department of Data Science and Artificial Intelligence,
Faculty of Engineering and Technology,
Sunway University
Selangor, Malaysia (corresponding author)

Zhize Wu
wuzz@hfuu.edu.cn
School of Artificial Intelligence and Big Data,
Hefei University
Hefei, Anhui, China

Abstract

The 2026 Space Optimization Competition (SpOC 4) poses a unique problem in Challenge 3: designing a lunar advertising constellation that generates a specific Morse code light-curve via artificial occultation. To navigate the mixed continuous-discrete search space, we enhance a baseline Simulated Annealing algorithm with Cross-Correlation Phase Alignment and an L-BFGS-B continuous polishing operator. Supported by an equi-cost replacement strategy and adaptive reheating, our method effectively breaks through topological traps. This framework achieves state-of-the-art performance, finding a constellation with $N = 34$ satellites with an optimal MSE of 0.0487. In turn, the beginner variant of Challenge 1 asks us to select a subset of m possible transfers in order to create a transport network with maximum throughput. Each transfer concerns an Earth orbit, a Moon orbit, and a destination on the surface of the Moon. However, no orbit or destination can be used twice. We approach this with a permutation-based representation which limits the search space to feasible solutions only.

CCS Concepts

• **Computing methodologies** → **Discrete space search**; **Randomized search**; *Heuristic function construction.*

Keywords

SpOC, Space Optimisation Competition, Graph Theory, Orbits

ACM Reference Format:
Mingxuan Li, Thomas Weise, Sina Abdipoor, and Zhize Wu. 2026. SpOC 4 Challenges 1 Beginner: Permutation Encodings and Challenge 3: Simulated Annealing for the Luna Tomato Advertising. In *Genetic and Evolutionary Computation Conference (GECCO Companion '26), July 13–17, 2026, San Jose, Costa Rica.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3795101.3815558

ACM ISBN 979-8-4007-2488-6/2026/07
https://doi.org/10.1145/3795101.3815558

1 Challenge 3

Challenge 3 of the 2026 Space Optimization Competition (SpOC 4) introduces a mathematically formidable lunar occultation advertising problem. The objective is to select a minimal subset of spacecraft from a candidate pool—comprising Distant Retrograde Orbits (DRO), Lyapunov, and Axial orbits—and optimally tune their initial phases. Their cumulative light-curve must synthesize a specific Morse code signal ("Tomatoes for sale") with high fidelity, encoded as:

`- -- -- .- - -- . ... / ..-. -- .-. / ... .- .-.. .`

The goal is to minimize the number N of spacecraft, selected from M candidate orbits, while satisfying the MSE constraint. Due to artificial occultation, the synthesized signal $S_{synth}(t)$ is constrained by a clipping function:

$$S_{synth}(t) = \min \left\{ \sum_{i=1}^{N} s(t; o_i, \phi_i), 1 \right\} \tag{1}$$

where $s(t; o_i, \phi_i)$ is the individual light-curve contribution. The objective requires $\text{MSE}(S_{synth}, S_{target}) \leq 0.05$.

Signal clipping limits the perceived light intensity to a normalized maximum. This truncation, alongside a strict MSE boundary, creates a highly non-linear, non-convex, and rugged fitness landscape [1]. Simultaneous combinatorial orbit selection and continuous phase optimization ($\phi \in [0, 2\pi)$) further complicate the search space. We develop an enhanced simulated annealing framework driven by physics-informed heuristics and hybrid continuous optimization, achieving the best known configuration.

Statistical analysis of 779 candidate orbits revealed that Lyapunov and DRO families offer broader period coverage, establishing bounds for the topological search. To navigate the mixed search space, we first use basic Simulated Annealing (SA) [3]. The active configuration is encoded as $A = \{(o_1, \phi_1), \ldots, (o_N, \phi_N)\}$. Standard neighborhood operators were implemented, such as *Add*, *Remove*, and *Shift*. To unify the infeasible region (MSE $>$ 0.05) and the minimization objective, a penalty-based fitness function $F(A)$ was utilized:

$$F(A) = \begin{cases} N, & \text{if MSE} \leq 0.05 \\ N + \lambda\sqrt{\text{MSE}}, & \text{if MSE} > 0.05 \end{cases} \quad (2)$$

While effective at finding initial feasible states, this baseline approach frequently stagnated, reaching a feasible constellation size of $N = 48$ spacecraft. To break the minimization bottleneck, we introduced the following improvements:

Cross-Correlation Phase Alignment: Instead of random phase assignment, we calculate the instantaneous residual signal $R(t) = S_{target}(t) - S_{synth}(t)$. A discrete grid search maximizes the cross-correlation score$(\phi) = \sum_t R(t) \cdot s(t; o_{new}, \phi)$. This deterministic initialization ensures new satellites precisely fill optical gaps, drastically improving the acceptance rate.

Continuous Polishing & Warm-Start Initialization: We replaced the baseline cliff-penalty with a soft-slope quadratic formulation. Marginal threshold violations trigger a continuous polishing operator based on the Limited-memory Broyden–Fletcher–Goldfarb–Shanno with box constraints (L-BFGS-B) algorithm. This bound-constrained quasi-Newton method provides efficient gradient-based phase fine-tuning to drag configurations back into the feasible domain. Crucially, reducing the constellation $(N - 1)$ is strictly seeded by the optimal N-satellite configuration. A poor randomized initial state would yield a massive MSE, trapping L-BFGS-B in expensive, fruitless gradients.

Equi-Cost Exploration and Reheating: An equi-cost *Replace* operator enables large-step structural mutations. Combined with *Adaptive Reheating* [2] that aggressively restores the system temperature to $T = T_{max} \times 0.1$ upon stagnation, this ensures the algorithm systematically breaks through topological traps using high-quality historical seeds.

The Enhanced SA fundamentally transformed the search dynamics, rapidly breaking through the 40 and 35 satellite barriers. Whenever minimization stalled, L-BFGS-B polishing and replacement operators helped escape local optima without incurring an N-cost penalty. These improvements drove the solution down to the best configuration we could reach: a constellation with $N = 34$ satellites, achieving an overall MSE of 0.0487.

2 Challenge 1 Beginner Mode

Given are the set $E = 1..n$ with n Earth orbits $e \in E$, the set $L = 1..n$ with n lunar orbits $l \in L$, and the set $D = 1..n$ with n destinations $d \in D$ on the moon surface. Cargo can be transported from an Earth orbit $e \in E$ to a Moon orbit $l \in L$ and finally to a destination $d \in D$. We can choose from a subset $T \subseteq E \times L \times D$ of $m = |T|$ combinations $t \in T$ of these orbits. A weight function $w : T \mapsto \mathbb{R}_0^+$ tells us how much cargo can be transported via a route $t \in T$.

There are two instances, **matching-i** with $n = 5000$ and $m = 25000$ and **matching-ii** with 10000 and 92103, respectively. A feasible solution y is a subset of T, i.e., $y \subseteq T$, such that no Earth orbit, Moon orbit, or destination appears twice. We want to find the solution in $\mathbb{Y}$ that maximizes the total cargo weight that can be transported.

Table 1: The sizes $|\mathbb{X}|$ for different designs of search spaces $\mathbb{X}$ on the challenge 1 beginner mode.

representation	$\lvert\mathbb{X}\rvert$	$\rightarrow$	**matching-i**	**matching-ii**
bits $(\{0,1\}^m)$	2^m	$\rightarrow$	$\approx 10^{7526}$	$\approx 10^{27726}$
1 permutation	$m!$	$\rightarrow$	$\approx 10^{99094}$	$\approx 10^{417228}$
3 permutations	$n!^3$	$\rightarrow$	$\approx 10^{48977}$	$\approx 10^{106978}$
2 permutations	$n!^2$	$\rightarrow$	$\approx 10^{32651}$	$\approx 10^{71319}$

This looks like a discrete problem defined over $\mathbb{Y} = \{0, 1\}^m$, where the bit at index i decides whether the i-th possible transfer in T, i.e., $T[i]$, is used or not. Searching in $\mathbb{Y}$ directly would mean that many sampled solutions will be infeasible because they contain some Earth or Moon orbit or destination twice.

We thus first decided to represent solutions as elements x of the set $\mathbb{X}$ of all permutations of the first m natural numbers. Such a permutation x is mapped by a decoding function $g : \mathbb{X} \mapsto \mathbb{Y}$ to a solution y. Initially, y is empty. We then process x from beginning to end. The element $x[i]$ at index i of x then addresses the $x[i]$-th transfer in the transfer list T, i.e., $\tau = T[x[i]]$. If $\tau.e$, $\tau.l$, and $\tau.d$ are not yet used in y, τ is added to y. Otherwise, it is ignored.

The result is a solution $y \in \mathbb{Y}$ which is guaranteed to be feasible. We then consider a representation using three permutations, each of the n first natural numbers. Instead of prioritizing transfer routes, we prioritize the Earth and Moon orbits and destinations separately, since each of them can only be used once. Since we can keep one of the three permutations constant, only two permutations are actually needed as search space. As Table 1 shows, this two-permutation search space is much smaller and searching within it is more efficient, however our experiments are not finished at the time of this writing.

So far, our best results with the 1 permutation encoding are obtained with a hybrid local search utilizing Frequency Fitness Assignment [5, 6] implemented in **moptipy** [4].

References

[1] Sina Abdipoor, Razali Yaakob, Say Leng Goh, and Salwani Abdullah. 2023. Metaheuristic approaches for the University Course Timetabling Problem. *Intelligent Systems with Applications* 19 (2023), 200253. doi:10.1016/j.iswa.2023.200253

[2] Say Leng Goh, Graham Kendall, and Nasser R. Sabar. 2017. Improved local search approaches to solve the post enrolment course timetabling problem. *European Journal of Operational Research* 261, 1 (2017), 17–29. doi:10.1016/j.ejor.2017.01.040

[3] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. 1983. Optimization by Simulated Annealing. *Science* 220, 4598 (1983), 671–680. arXiv:https://www.science.org/doi/pdf/10.1126/science.220.4598.671 doi:10.1126/science.220.4598.671

[4] Thomas Weise and Zhize Wu. 2023. Replicable Self-Documenting Experiments with Arbitrary Search Spaces and Algorithms. In *Genetic and Evolutionary Computation Conference Companion (GECCO'2023 Companion), July 15-19, 2023, Lisbon, Portugal.* ACM, New York, NY, USA. doi:10.1145/3583133.3596306

[5] Thomas Weise, Zhize Wu, Xinlu Li, and Yan Chen. 2021. Frequency Fitness Assignment: Making Optimization Algorithms Invariant under Bijective Transformations of the Objective Function Value. *IEEE Transactions on Evolutionary Computation* 25, 2 (2021), 307–319. doi:10.1109/TEVC.2020.3032090

[6] Thomas Weise, Zhize Wu, Xinlu Li, Yan Chen, and Jörg Lässig. 2023. Frequency Fitness Assignment: Optimization without Bias for Good Solutions can be Efficient. *IEEE Transactions on Evolutionary Computation* 27, 4 (Aug. 2023), 980–992. doi:10.1109/TEVC.2022.3191698