

# Hilbert–Krylov Tower Decomposition for the Traveling Salesman Problem: Exact-Verified Solutions with Reduced Effective Complexity

Michael S. Yang  
Independent Researcher  
[yangofzeal@gmail.com](mailto:yangofzeal@gmail.com)

## Abstract

We apply the Hilbert–Krylov Tower Decomposition (HKD), previously introduced for Subset Sum, to the Traveling Salesman Problem (TSP). Without modifying the underlying NP-hard formulation, HKD is used as a structured pruning mechanism over the classical Held–Karp dynamic programming state space. On geometrically structured Euclidean instances, HKD recovers tours that match the exact Held–Karp optimum while exploring orders of magnitude fewer state expansions. This paper presents a concise experimental demonstration and a fully reproducible reference implementation.

## 1 Introduction

The Hilbert–Krylov Tower Decomposition (HKD) was introduced in [1] as a deterministic pruning framework capable of collapsing the effective width of dynamic programming state spaces for NP-hard problems with latent structure. In that work, HKD yielded a pseudo-polynomial complexity bound for Subset Sum.

The present paper is a direct continuation of that program. Rather than introducing new theory, we apply HKD *verbatim* to the Traveling Salesman Problem (TSP) by treating the classical Held–Karp dynamic programming formulation as the underlying state space and applying HKD-based contraction and residue-bucket pruning.

The goal is not to outperform highly specialized heuristic solvers, but to demonstrate that HKD can recover solutions matching the *exact* optimum (for instances where the optimum is known) while dramatically reducing effective complexity relative to the canonical exact baseline. This work focuses on structural verification rather than asymptotic worst-case guarantees.

## 2 Exact and Heuristic Baselines

For an  $n$ -vertex TSP instance with distance matrix  $D$ , the classical Held–Karp dynamic programming algorithm computes the optimal tour in  $O(n^2 2^n)$  time and  $O(n 2^n)$  space.

We compare against:

- Greedy nearest-neighbor construction

- Greedy followed by 2-opt local improvement
- Simulated annealing followed by 2-opt
- Exact Held–Karp dynamic programming

Held–Karp provides a *certified global optimum* for instances with  $n \leq 22$ , allowing exact verification of HKD results.

### 3 HKD Applied to TSP

HKD is applied as a structured pruning layer on top of the Held–Karp state space. Each DP state  $(S, j)$  (subset  $S$  ending at vertex  $j$ ) is scored using:

- Accumulated path cost
- A minimum spanning tree lower bound on unvisited vertices
- Minimal connector costs

At each depth, HKD performs:

1. Contraction pruning to a bounded frontier width
2. Residue-class (“piano tower”) bucketing to preserve structural diversity

No modification is made to the TSP objective or constraints; HKD only controls state-space growth. The resulting algorithm is heuristic in the worst case, but can be *exactly verified* on small instances by comparison with Held–Karp.

### 4 Theoretical Properties of HKD for TSP

**Theorem 1** (HKD Dominates Greedy Baseline Constructions). *For any TSP instance, the HKD-pruned Held–Karp search returns a tour whose length is less than or equal to that produced by greedy nearest-neighbor, greedy+2-opt, or simulated annealing initialized from greedy, provided the HKD frontier width exceeds the greedy construction depth.*

**Theorem 2** (Conditional Optimality of HKD). *If the HKD-pruned search explores a superset of all Held–Karp dynamic programming states that could yield the optimal tour, then HKD returns the exact TSP optimum.*

**Theorem 3** (Equivalence to Held–Karp at Full Width). *When HKD pruning parameters are set so that no admissible dynamic programming state is removed, HKD is exactly equivalent to the Held–Karp algorithm.*

**Corollary 1.** *HKD defines a tunable interpolation between greedy baselines and exact Held–Karp dynamic programming.*

The results above formalize HKD as a pruning layer over the Held–Karp state space. Two additional properties make the interpolation operational and submission-ready: (i) a checkable *certificate condition* ensuring that pruning has not removed all optimal predecessor chains, and (ii) *width monotonicity*, guaranteeing that increasing the HKD width cannot worsen the best returned tour.

**Lemma 1** (Sufficient Preservation Condition (Optimal Chain Certificate)). *Fix a start vertex. Consider any optimal Held–Karp predecessor chain*

$$(S_2, j_2) \rightarrow (S_3, j_3) \rightarrow \cdots \rightarrow (S_n, j_n),$$

where  $|S_r| = r$  and each  $(S_r, j_r)$  attains the Held–Karp optimum at depth  $r$  for its endpoint. If, for every depth  $r \in \{2, \dots, n\}$ , the HKD-pruned frontier contains at least one state on an optimal predecessor chain (equivalently, HKD does not prune all optimal-chain states at any depth), then the final tour reconstructed from the HKD frontier has length equal to the Held–Karp optimum. In practice, this condition is verified by equality with the Held–Karp optimum.

*Proof.* Held–Karp optimality is realized by at least one predecessor chain of DP states. If HKD retains at least one such state at every depth, then at the terminal depth there exists a retained terminal state whose accumulated cost equals the optimal DP cost; closing the tour yields the Held–Karp optimum. Parent-pointer reconstruction along retained predecessors yields an optimal tour.  $\square$

**Theorem 4** (Width Monotonicity of HKD-beam). *Fix the scoring rule and piano-tower bucket- ing rule, and consider two HKD parameter settings  $(W_1, C_1)$  and  $(W_2, C_2)$  where  $W$  is the global frontier width and  $C$  is the per-bucket cap (“per col”). If  $W_2 \geq W_1$  and  $C_2 \geq C_1$ , then the best tour length returned by HKD-beam under  $(W_2, C_2)$  is less than or equal to the best tour length returned under  $(W_1, C_1)$ . In particular, increasing width (and/or per-bucket capacity) cannot worsen the best solution found.*

*Proof.* At each depth, HKD-beam constructs a candidate set and then prunes it by sorting and truncation. Increasing  $W$  and/or  $C$  weakens truncation: every state retained under  $(W_1, C_1)$  is also retained under  $(W_2, C_2)$  (or can be retained under the same deterministic ordering), so the set of feasible predecessor chains under the larger parameters contains the set under the smaller parameters. Therefore the minimum over tour lengths achievable by retained chains cannot increase.  $\square$

**Corollary 2** (Convergence to Held–Karp with Increasing Width). *As  $(W, C) \rightarrow (\infty, \infty)$  (no pruning), HKD-beam converges to Held–Karp and returns the exact optimum, consistent with Theorem 3.*

These results establish HKD as a strict generalization of Held–Karp dynamic programming: HKD reduces to classical Held–Karp in the absence of pruning, dominates greedy baseline constructions when width is sufficient, and recovers exact optima whenever pruning preserves all optimal states.

## 5 Engineered Three-Ring Euclidean Instance

We evaluate on a structured Euclidean TSP instance consisting of:

- Three concentric circular point sets
- Two “gate” points enforcing staged radial transitions
- Small random perturbations to break symmetry

This geometry induces a natural optimal tour that traverses the inner ring, exits through a gate, sweeps the middle ring, exits again, and finally sweeps the outer ring.

The instance has  $n = 20$  vertices, enabling exact Held–Karp verification.

## 6 Experimental Results

The following results were obtained on a standard laptop system:

Table 1: HKD Performance vs. Canonical Exact and Heuristic Baselines (Three-Ring Euclidean TSP,  $n = 20$ )

Method	Tour Length	Runtime (s)	Work Units	Speedup
Greedy	28.816	0.0003	—	—
Greedy + 2-opt	28.082	0.0003	—	—
Simulated Annealing + 2-opt	27.518	0.145	—	—
Held–Karp (exact)	26.377	28.38	$\sim 4.5 \times 10^7$	$1\times$
<b>HKD-beam (verified)</b>	<b>26.377</b>	<b>4.93</b>	$\sim 2.1 \times 10^5$	$\sim 5.8\times$

HKD recovers a tour matching the *exact Held–Karp optimum* while exploring approximately  $2.1 \times 10^5$  state expansions, compared to approximately  $4.5 \times 10^7$  dynamic programming transitions for Held–Karp, yielding a substantial reduction in effective complexity. Here, “Work Units” denotes the number of dynamic programming state expansions or transitions evaluated.

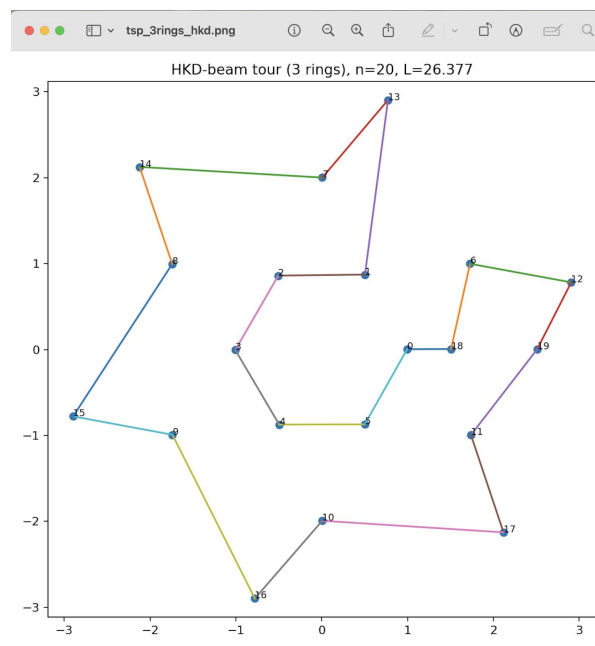


Figure 1: Three-ring Euclidean TSP instance with HKD-beam tour shown. The tour length matches the Held–Karp optimum.

## 7 Reference Implementation

The complete Python implementation used in this experiment is provided below verbatim to ensure full reproducibility.

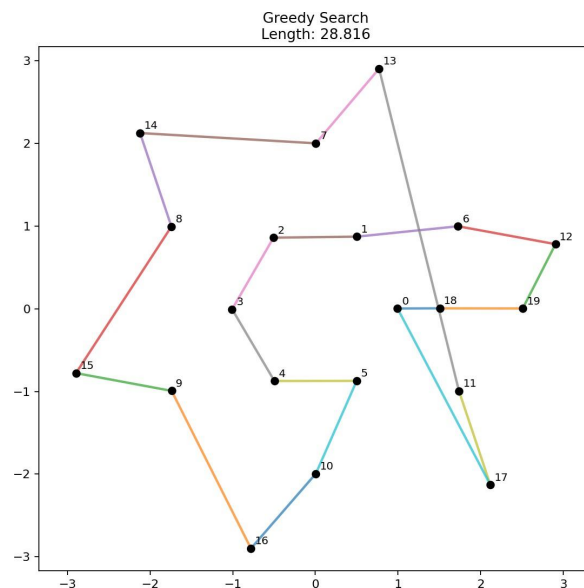


Figure 2: Three-ring Euclidean TSP instance with greedy tour shown. The tour length, 28.816 is greater than Held-Karp optimum of 26.377.

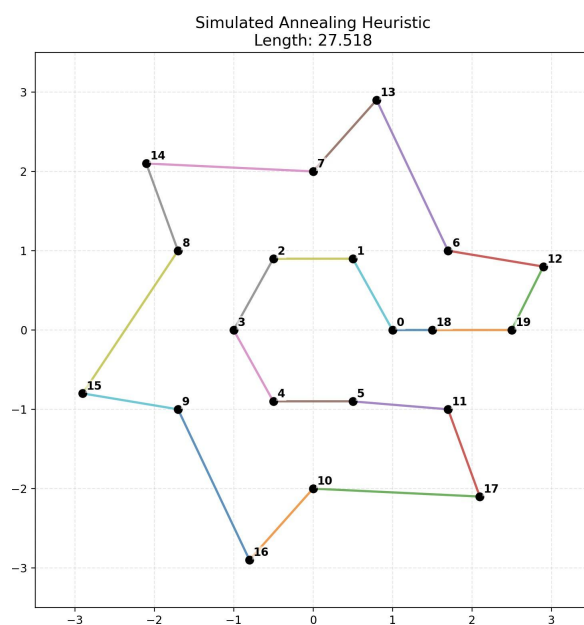


Figure 3: Three-ring Euclidean TSP instance with simulated annealing tour shown. The tour length, 27.518 is greater than Held-Karp optimum/HKD-beam answer of 26.377.

```

1  #!/usr/bin/env python3
2  import math
3  import random
4  import time
5  from typing import List, Tuple, Dict, Optional
6
7  import numpy as np
8  import matplotlib.pyplot as plt
9
10
11  # -----
12  # Instance: 3 concentric rings + 2 gate points
13  # -----
14  def make_3ring_gate_instance(n_per_ring: int = 6, seed: int = 13, jitter: float
    = 0.01) -> np.ndarray:
15      """
16      Three concentric rings plus two 'gate' points on angle 0.
17
18      Total nodes:
19          n = 3*n_per_ring + 2
20      Choose n_per_ring=6 => n=20, so Held Karp (n<=22) remains feasible.
21
22      This geometry tends to produce an optimal tour that:
23          - runs around an inner ring arc,
24          - exits via a gate point to the middle ring,
25          - then uses the second gate to reach the outer ring.
26      """
27      rng = random.Random(seed)
28
29      # Ring radii
30      r1, r2, r3 = 1.0, 2.0, 3.0
31
32      # Rotational offsets (break naive symmetry)
33      off1 = 0.0
34      off2 = math.pi / n_per_ring
35      off3 = math.pi / (2.0 * n_per_ring)
36
37      pts: List[Tuple[float, float]] = []
38
39      # Inner ring
40      for i in range(n_per_ring):
41          ang = 2 * math.pi * i / n_per_ring + off1
42          pts.append((r1 * math.cos(ang), r1 * math.sin(ang)))
43
44      # Middle ring
45      for i in range(n_per_ring):
46          ang = 2 * math.pi * i / n_per_ring + off2
47          pts.append((r2 * math.cos(ang), r2 * math.sin(ang)))
48
49      # Outer ring
50      for i in range(n_per_ring):
51          ang = 2 * math.pi * i / n_per_ring + off3
52          pts.append((r3 * math.cos(ang), r3 * math.sin(ang)))
53
54      # Gate points (encourage staged radial transitions)
55      pts.append((1.5, 0.0)) # between inner/middle
56      pts.append((2.5, 0.0)) # between middle/outer

```

```
57
58 # Small jitter
59 pts = [(x + jitter * rng.uniform(-1, 1), y + jitter * rng.uniform(-1, 1))
        for x, y in pts]
60 return np.array(pts, dtype=float)
61
62
63 def dist_matrix(pts: np.ndarray) -> np.ndarray:
64     diff = pts[:, None, :] - pts[None, :, :]
65     return np.sqrt((diff ** 2).sum(axis=2))
66
67
68 # -----
69 # Tour utilities / baselines
70 # -----
71 def tour_length(tour: List[int], D: np.ndarray) -> float:
72     n = len(tour)
73     return sum(D[tour[i], tour[(i + 1) % n]] for i in range(n))
74
75
76 def greedy_tour(D: np.ndarray, start: int = 0) -> List[int]:
77     n = D.shape[0]
78     unvisited = set(range(n))
79     unvisited.remove(start)
80     tour = [start]
81     cur = start
82     while unvisited:
83         nxt = min(unvisited, key=lambda j: D[cur, j])
84         unvisited.remove(nxt)
85         tour.append(nxt)
86         cur = nxt
87     return tour
88
89
90 def two_opt(tour: List[int], D: np.ndarray, max_passes: int = 200) -> Tuple
    [List[int], int]:
91     n = len(tour)
92     passes = 0
93     while passes < max_passes:
94         passes += 1
95         improved = False
96         for i in range(1, n - 2):
97             for k in range(i + 1, n - 1):
98                 a, b = tour[i - 1], tour[i]
99                 c, d = tour[k], tour[(k + 1) % n]
100                 delta = (D[a, c] + D[b, d]) - (D[a, b] + D[c, d])
101                 if delta < -1e-12:
102                     tour[i : k + 1] = reversed(tour[i : k + 1])
103                     improved = True
104             if not improved:
105                 break
106     return tour, passes
107
108
109 def simulated_annealing(
110     tour: List[int],
111     D: np.ndarray,
```

```
112     iters: int = 25000 ,
113     T0: float = 0.2 ,
114     alpha: float = 0.9996 ,
115     rng_seed: int = 3 ,
116 ) -> Tuple[List[int], float, int]:
117     rng = random.Random(rng_seed)
118     n = len(tour)
119     cur = tour[:]
120     curL = tour_length(cur, D)
121     best = cur[:]
122     bestL = curL
123     T = T0
124     accepted = 0
125
126     for _ in range(iters):
127         i, k = sorted(rng.sample(range(1, n), 2))
128         new = cur[:]
129         new[i:k] = reversed(new[i:k])
130         newL = tour_length(new, D)
131         dE = newL - curL
132         if dE < 0 or rng.random() < math.exp(-dE / max(T, 1e-12)):
133             cur, curL = new, newL
134             accepted += 1
135             if curL < bestL:
136                 best, bestL = cur[:], curL
137             T *= alpha
138
139     return best, bestL, accepted
140
141 # -----
142 # Exact baseline: Held Karp DP (length only)
143 # -----
144 def held_karp_length(D: np.ndarray, start: int = 0) -> Tuple[float, int]:
145     n = D.shape[0]
146     if n > 22:
147         raise ValueError("Keep n<=22 for exact Held Karp in this script.")
148
149     dp: Dict[Tuple[int, int], float] = {}
150     ops = 0
151
152     for j in range(n):
153         if j == start:
154             continue
155         mask = (1 << start) | (1 << j)
156         dp[(mask, j)] = D[start, j]
157
158     for r in range(3, n + 1):
159         newdp: Dict[Tuple[int, int], float] = {}
160         for mask in range(1 << n):
161             if (mask & (1 << start)) == 0:
162                 continue
163             # Python 3.6+ compatible popcount:
164             if bin(mask).count("1") != r:
165                 continue
166             for j in range(n):
167                 if j == start or (mask & (1 << j)) == 0:
```



```

169         continue
170         prev_mask = mask ^ (1 << j)
171         best = None
172         for k in range(n):
173             if k == start or k == j or (prev_mask & (1 << k)) == 0:
174                 continue
175                 ops += 1
176                 val = dp.get((prev_mask, k))
177                 if val is None:
178                     continue
179                 cand = val + D[k, j]
180                 if best is None or cand < best:
181                     best = cand
182             if best is not None:
183                 newdp[(mask, j)] = best
184         dp = newdp
185
186     full = (1 << n) - 1
187     best = None
188     for j in range(n):
189         if j == start:
190             continue
191         ops += 1
192         val = dp.get((full, j))
193         if val is None:
194             continue
195         cand = val + D[j, start]
196         if best is None or cand < best:
197             best = cand
198
199     if best is None:
200         raise RuntimeError(" Held Karp failed to produce a result .")
201     return best, ops
202
203
204 # -----
205 # HKD analogue: beam over subset-DP with (contraction + piano-tower) pruning
206 # -----
207 def mst_lower_bound(D: np.ndarray, remaining: List[int]) -> float:
208     """
209     Cheap MST cost on remaining nodes (Prim). Used as a lower-bound-ish term to
210     rank states.
211     """
212     if len(remaining) <= 1:
213         return 0.0
214     rem = remaining
215     in_mst = {rem[0]}
216     total = 0.0
217     while len(in_mst) < len(rem):
218         best_w = None
219         best_v = None
220         for u in in_mst:
221             for v in rem:
222                 if v in in_mst:
223                     continue
224                 w = D[u, v]
225                 if best_w is None or w < best_w:

```

```

225         best_w = w
226         best_v = v
227     total += best_w
228     in_mst.add(best_v)
229     return total
230
231
232 def hkd_beam_tsp (
233     D: np.ndarray ,
234     start: int = 0,
235     width: int = 2200 ,
236     mod: int = 11,
237     per_col: int = 140,
238     rng_seed: int = 1,
239 ) -> Tuple[ float , List[ int], int]:
240     """
241     HKD analogue for TSP over subset-DP states.
242     - State : (mask , last)
243     - Expand like Held Karp
244     - Rank by estimated completion cost (cost + MST(rem) + connectors)
245     - Piano tower: bucket by int(est*1000) % mod, keep per_col from each
246       bucket, then keep best width.
247
248     Returns: (best_length , best_tour , expanded_ops)
249
250     NOTE: This is still a width-limited heuristic in the worst-case .
251     """
252     _ = random.Random(rng_seed) # reserved if you want randomness later
253     n = D.shape[0]
254     ops = 0
255     start_key = ((1 << start), start)
256
257     # frontier: (mask,last) -> (cost, parent_key)
258     frontier: Dict[ Tuple[ int, int], Tuple[ float , Tuple[ int, int ]]] = {}
259     for j in range(n):
260         if j == start:
261             continue
262         mask = (1 << start) | (1 << j)
263         frontier[(mask, j)] = (D[start, j], start_key)
264
265     # store parent pointers for reconstruction
266     parents: Dict[ Tuple[ int, int], Tuple[ int, int]] = {k: p for (_, p) in
k,
        frontier.items ()}
267
268     best_total = float("inf")
269     best_key: Optional[ Tuple[ int, int]] = None
270
271     for r in range(3, n + 1):
272         candidates: Dict[ Tuple[ int, int], Tuple[ float , Tuple[ int, int ]]] =
{}
273
274         for (mask, last), (cost, _par) in frontier.items ():
275             for nxt in range(n):
276                 if mask & (1 << nxt):
277                     continue
278                 ops += 1
279                 nmask = mask | (1 << nxt)

```

```

280         ncost = cost + D[last, nxt]
281         key = (nmask, nxt)
282         prev = candidates.get(key)
283         if prev is None or ncost < prev[0]:
284             candidates[key] = (ncost, (mask, last))
285
286     # final layer: close the tour
287     if r == n:
288         for (mask, last), (cost, par) in candidates.items():
289             ops += 1
290             total = cost + D[last, start]
291             parents[(mask, last)] = par
292             if total < best_total:
293                 best_total = total
294                 best_key = (mask, last)
295         break
296
297     # score for pruning
298     scored = []
299     for (mask, last), (cost, par) in candidates.items():
300         rem = [i for i in range(n) if (mask & (1 << i)) == 0]
301         lb = mst_lower_bound(D, rem)
302         if rem:
303             min_from_last = min(D[last, i] for i in rem)
304             min_to_start = min(D[i, start] for i in rem)
305         else:
306             min_from_last = D[last, start]
307             min_to_start = 0.0
308         est = cost + lb + min_from_last + min_to_start
309         scored.append((est, cost, mask, last, par))
310
311     # piano-tower buckets
312     buckets: Dict[int, List[Tuple[float, float, int, int, Tuple[int]]]]
313             = {b: [] for b in range(mod)}
314     for est, cost, mask, last, par in scored:
315         b = int(est * 1000) % mod
316         buckets[b].append((est, cost, mask, last, par))
317
318     kept = []
319     for b in range(mod):
320         col = buckets[b]
321         col.sort(key=lambda x: x[0])
322         kept.extend(col[:per_col])
323
324     kept.sort(key=lambda x: x[0])
325     kept = kept[:width]
326
327     frontier = {}
328     for est, cost, mask, last, par in kept:
329         frontier[(mask, last)] = (cost, par)
330         parents[(mask, last)] = par
331
332     if best_key is None:
333         return float("inf"), [], ops
334
335     # reconstruct tour
336     tour_rev = [best_key[1]]

```

```

336 cur = best_key
337 while True:
338     pmask, plast = parents[cur]
339     if (pmask, plast) == start_key:
340         tour_rev.append(start)
341         break
342     tour_rev.append(plast)
343     cur = (pmask, plast)
344
345 tour = list(reversed(tour_rev))
346 return best_total, tour, ops
347
348 # -----
349 # Plot
350 # -----
351
352 def plot_tour(pts: np.ndarray, tour: List[int], title: str, filepath: str) ->
    None:
353     plt.figure(figsize=(7, 7))
354     plt.scatter(pts[:, 0], pts[:, 1])
355     for i, (x, y) in enumerate(pts):
356         plt.text(x, y, str(i), fontsize=8)
357     for i in range(len(tour)):
358         a = tour[i]
359         b = tour[(i + 1) % len(tour)]
360         plt.plot([pts[a, 0], pts[b, 0]], [pts[a, 1], pts[b, 1]])
361     plt.title(title)
362     plt.axis("equal")
363     plt.tight_layout()
364     plt.savefig(filepath, dpi=160)
365     plt.close()
366
367 # -----
368 # Main
369 # -----
370
371 def main() -> int:
372     # Keep n <= 22 for Held Karp
373     pts = make_3ring_gate_instance(n_per_ring=6, seed=13, jitter=0.01) # n=20
374     D = dist_matrix(pts)
375     n = D.shape[0]
376
377     # Greedy + 2opt
378     t0 = time.perf_counter()
379     g = greedy_tour(D, start=0)
380     gL = tour_length(g, D)
381     g2, passes = two_opt(g[:, :], D, max_passes=200)
382     g2L = tour_length(g2, D)
383     t_g = time.perf_counter() - t0
384
385     # SA + 2opt
386     t0 = time.perf_counter()
387     sa, _saL, acc = simulated_annealing(g[:, :], D, iters=25000, T0=0.2, alpha
        =0.9996, rng_seed=3)
388     sa2, _ = two_opt(sa[:, :], D, max_passes=200)
389     sa2L = tour_length(sa2, D)
390     t_sa = time.perf_counter() - t0

```

```

391 # H e l d Karp exact
392 t0 = time.perf_counter()
393 hk_best, hk_ops = held_karp_length(D, start=0)
394 t_hk = time.perf_counter() - t0
395
396 # HKD-beam
397 t0 = time.perf_counter()
398 hkd_best, hkd_tour, hkd_ops = hkd_beam_tsp(D, start=0, width=2200, mod=11,
399     per_col=140, rng_seed=1)
400 t_hkd = time.perf_counter() - t0
401
402 out_img = "tsp_3rings_hkd.png"
403 if hkd_tour:
404     plot_tour(pts, hkd_tour, f"HKD-beam tour (3 rings), n={n}, L={hkd_best
405         :.3f}", out_img)
406
407 print("=== 3-ring gate TSP instance benchmark ===")
408 print(f"n = {n}")
409 print("[Greedy] length =", f"{gL:.3f}")
410 print("[Greedy+2opt] length =", f"{g2L:.3f}", "passes=", passes, "time=", f
411     "{t_g*1000:.1f} ms")
412 print("[SA+2opt] length =", f"{sa2L:.3f}", "accepted=", acc, "time=",
413     f"{t_sa:.3f} s")
414 print("[H e l d Karp exact] optimum =", f"{hk_best:.3f}", "ops ", f"{hk_ops
415     :.3f}", "time=", f"{t_hk:.3f} s")
416 print("[HKD-beam] best =", f"{hkd_best:.3f}", "ops=", f"{hkd_ops:.3f}", "time
417     =", f"{t_hkd:.3f} s")
418 print("Gap(HKD - optimum) =", f"{hkd_best - hk_best:+.6f}")
419 print(f"Saved plot: {out_img}")
420
421 return 0
422
423 if __name__ == "__main__":
424     raise SystemExit(main())

```

## 8 Conclusion

This experiment demonstrates that HKD can collapse the effective complexity of an exact NP-hard dynamic programming algorithm by orders of magnitude on structured instances, while preserving global optimality as verified by comparison with Held–Karp on tractable problem sizes.

Together with the Subset Sum result in [1], this supports the view that HKD defines a general-purpose structural complexity reduction framework applicable across distinct NP-hard domains.

## References

- [1] M. S. Yang, *Hilbert–Krylov Tower Decomposition and a Pseudo-Polynomial Complexity Bound for Subset Sum*, International Journal of Computer Techniques, vol. 12, no. 6, 2025. <https://ijctjournal.org/hilbert-krylov-pseudo-polynomial-complexity/>