

NAG Toolbox

nag_mip_tsp_simann (h03bb)

1 Purpose

nag_mip_tsp_simann (h03bb) calculates an approximate solution to a symmetric travelling salesman problem using simulated annealing via a configuration free interface.

2 Syntax

```
[path, cost, tmode, alg_stats, state, ifail] = nag_mip_tsp_simann(dm, bound,
targc, state, 'nc', nc)
```

```
[path, cost, tmode, alg_stats, state, ifail] = h03bb(dm, bound, targc, state,
'nc', nc)
```

3 Description

nag_mip_tsp_simann (h03bb) provides a probabilistic strategy for the calculation of a near optimal path through a symmetric and fully connected distance matrix; that is, a matrix for which element (i, j) is the pairwise distance (also called the cost, or weight) between nodes (cities) i and j . This problem is better known as the Travelling Salesman Problem (TSP), and symmetric means that the distance to travel between two cities is independent of which is the destination city.

In the classical TSP, which this function addresses, a salesman wishes to visit a given set of cities once only by starting and finishing in a home city and travelling the minimum total distance possible. It is one of the most intensively studied problems in computational mathematics and, as a result, has developed some fairly sophisticated techniques for getting near-optimal solutions for large numbers of cities. nag_mip_tsp_simann (h03bb) adopts a very simple approach to try to find a reasonable solution, for moderately large problems. The function uses simulated annealing: a stochastic mechanical process in which the heating and controlled cooling of a material is used to optimally refine its molecular structure.

The material in the TSP is the distance matrix and a given state is represented by the order in which each city is visited—the path. This system can move from one state to a neighbouring state by selecting two cities on the current path at random and switching their places; the order of the cities in the path between the switched cities is then reversed. The cost of a state is the total cost of traversing its path; the resulting difference in cost between the current state and this new proposed state is called the delta; a negative delta indicates the proposal creates a more optimal path and a positive delta a less optimal path. The random selection of cities to switch uses random number generators (RNGs) from Chapter G05; it is thus necessary to initialize a state array for the RNG of choice (by a call to nag_rand_init_repeat (g05kf) or nag_rand_init_nonrepeat (g05kg)) prior to calling nag_mip_tsp_simann (h03bb).

The simulation itself is executed in two stages. In the first stage, a series of sample searches through the distance matrix is conducted where each proposed new state is accepted, regardless of the change in cost (delta) incurred by applying the switches, and statistics on the set of deltas are recorded. These metrics are updated after each such sample search; the number of these searches and the number of switches applied in each search is dependent on the number of cities. The final collated set of metrics for the deltas obtained by the first stage are used as control parameters for the second stage. If no single improvement in cost is found during the first stage, the algorithm is terminated.

In the second stage, as before, neighbouring states are proposed. If the resulting delta is negative or causes no change the proposal is accepted and the path updated; otherwise moves are accepted based on a probabilistic criterion, a modified version of the Metropolis–Hastings algorithm.

The acceptance of some positive deltas (increased cost) reduces the probability of a solution getting trapped at a non-optimal solution where any single switch causes an increase in cost. Initially the

acceptance criteria allow for relatively large positive deltas, but as the number of proposed changes increases, the criteria become more stringent, allowing fewer positive deltas of smaller size to be accepted; this process is, within the realm of the simulated annealing algorithm, referred to as ‘cooling’. Further exploration of the system is initially encouraged by accepting non-optimal routes, but is increasingly discouraged as the process continues.

The second stage will terminate when:

- a solution is obtained that is deemed acceptable (as defined by supplied values);
- the algorithm will accept no further positive deltas and a set of proposed changes have resulted in no improvements (has cooled);
- a number of consecutive sets of proposed changes has resulted in no improvement.

4 References

Applegate D L, Bixby R E, Chvátal V and Cook W J (2006) *The Traveling Salesman Problem: A Computational Study* Princeton University Press

Cook W J (2012) *In Pursuit of the Traveling Salesman* Princeton University Press

Johnson D S and McGeoch L A The traveling salesman problem: A case study in local optimization *Local search in combinatorial optimization* (1997) 215–310

Press W H, Teukolsky S A, Vetterling W T and Flannery B P (2007) *Numerical Recipes The Art of Scientific Computing (3rd Edition)*

Rego C, Gamboa D, Glover F and Osterman C (2011) Traveling salesman problem heuristics: leading methods, implementations and latest advances *European Journal of Operational Research* **211** (3) 427–441

Reinelt G (1994) *The Travelling Salesman. Computational Solutions for TSP Applications, Volume 840 of Lecture Notes in Computer Science* Springer–Verlag, Berlin Heidelberg New York

5 Parameters

5.1 Compulsory Input Parameters

1: **dm(nc,nc)** – REAL (KIND=nag_wp) array

The distance matrix; each $\mathbf{dm}(i,j)$ is the effective cost or weight between nodes i and j . Only the strictly upper half of the matrix is referenced.

Constraint: $\mathbf{dm}(i,j) \geq 0.0$, for $j = 2, 3, \dots, \mathbf{nc}$ and $i = 1, 2, \dots, j - 1$.

2: **bound** – REAL (KIND=nag_wp)

A lower bound on the solution. If the optimum is unknown set **bound** to zero or a negative value; the function will then calculate the minimum spanning tree for **dm** and use this as a lower bound (returned in **alg_stats**(6)). If an optimal value for the cost is known then this should be used for the lower bound. A detailed discussion of relaxations for lower bounds, including the minimal spanning tree, can be found in Reinelt (1994).

3: **targc** – REAL (KIND=nag_wp)

A measure of how close an approximation needs to be to the lower bound. The function terminates when a cost is found less than or equal to **bound** + **targc**. This argument is useful when an optimal value for the cost is known and supplied in **bound**. It may be sufficient to obtain a path that is close enough (in terms of cost) to the optimal path; this allows the algorithm to terminate at that point and avoid further computation in attempting to find a better path.

If **targc** < 0, **targc** = 0 is assumed.

4: **state**(:) – INTEGER array

Note: the actual argument supplied **must** be the array **state** supplied to the initialization routines `nag_rand_init_repeat (g05kf)` or `nag_rand_init_nonrepeat (g05kg)`.

A valid RNG state initialized by `nag_rand_init_repeat (g05kf)` or `nag_rand_init_nonrepeat (g05kg)`. Since the algorithm used is stochastic, a random number generator is employed; if the generator is initialized to a non-repeatable sequence (`nag_rand_init_nonrepeat (g05kg)`) then different solution paths will be taken on successive runs, returning possibly different final approximate solutions.

5.2 Optional Input Parameters

1: **nc** – INTEGER

Default: the first dimension of the array **dm** and the second dimension of the array **dm**. (An error is raised if these dimensions are not equal.)

The number of cities. In the trivial cases **nc** = 1, 2 or 3, the function returns the optimal solution immediately with **tmode** = 0 (provided the relevant distance matrix entries are not negative).

Constraint: **nc** ≥ 1.

5.3 Output Parameters

1: **path**(**nc**) – INTEGER array

The best path discovered by the simulation. That is, **path** contains the city indices in path order. If **ifail** ≠ 0 on exit, **path** contains the indices 1 to **nc**.

2: **cost** – REAL (KIND=nag_wp)

The cost or weight of **path**. If **ifail** ≠ 0 on exit, **cost** contains the largest model real number (see `nag_machine_model_maxexp (x02bl)`).

3: **tmode** – INTEGER

The termination mode of the function (if **ifail** ≠ 0 on exit, **tmode** is set to -1):

tmode = 0

Optimal solution found, **cost** = **bound**.

tmode = 1

System temperature cooled. The algorithm returns a **path** and associated **cost** that does not attain, nor lie within **target** of, the **bound**. This could be a sufficiently good approximation to the optimal **path**, particularly when **bound** + **target** lies below the optimal **cost**.

tmode = 2

Halted by **cost** falling within the desired **target** range of the **bound**.

tmode = 3

System stalled following lack of improvement.

tmode = 4

Initial search failed to find a single improvement (the solution could be optimal).

4: **alg_stats**(6) – REAL (KIND=nag_wp) array

An array of metrics collected during the initial search. These could be used as a basis for future optimization. If **ifail** ≠ 0 on exit, the elements of **alg_stats** are set to zero; the first five elements are also set to zero in the trivial cases **nc** = 1, 2 or 3.

alg_stats(1)

Mean delta.

alg_stats(2)

Standard deviation of deltas.

alg_stats(3)

Cost at end of initial search phase.

alg_stats(4)

Best cost encountered during search phase.

alg_stats(5)

Initial system temperature. At the end of stage 1 of the algorithm, this is a function of the mean and variance of the deltas, and of the distance from best cost to the lower bound. It is a measure of the initial acceptance criteria for stage 2. The larger this value, the more iterations it will take to geometrically reduce it during stage 2 until the system is cooled (below a threshold value).

alg_stats(6)

The lower bound used, which will be that computed internally when **bound** ≤ 0 on input. Subsequent calls with different random states can set **bound** to the value returned in **alg_stats(6)** to avoid recomputation of the minimal spanning tree.

5: **state(:)** – INTEGER array

Contains updated information on the state of the generator.

6: **ifail** – INTEGER

ifail = 0 unless the function detects an error (see Section 5).

6 Error Indicators and Warnings

Errors or warnings detected by the function:

ifail = 1

Constraint: **nc** ≥ 1 .

ifail = 2

On entry, the strictly upper triangle of **dm** had a negative element.

ifail = 9

On entry, **state** vector has been corrupted or not initialized.

ifail = -99

An unexpected error has been triggered by this routine. Please contact NAG.

ifail = -399

Your licence key may have expired or may not have been installed correctly.

ifail = -999

Dynamic memory allocation failed.

7 Accuracy

The function will not perform well when the average change in cost caused by switching two cities is small relative to the cost; this can happen when many of the values in the distance matrix are relatively close to each other.

The quality of results from this function can vary quite markedly when different initial random states are used. It is therefore advisable to compute a number of approximations using different initial random states. The best cost and path can then be taken from the set of approximations obtained. If no change in results is obtained after 10 such trials then it is unlikely that any further improvement can be made by this function.

8 Further Comments

Memory is internally allocated for $3 \times \mathbf{nc} - 2$ integers and $\mathbf{nc} - 1$ real values.

In the case of two cities that are not connected, a suitably large number should be used as the distance (cost) between them so as to deter solution paths which directly connect the two cities. Solutions which contain an artificial link (i.e., a connection with a large distance between them to indicate no actual link) may be patched, using the shortest path algorithm `nag_mip_shortestpath` (h03ad).

If a city is to be visited more than once (or more than twice for the home city) then the distance matrix should contain multiple entries for that city (on rows and columns i_1, i_2, \dots) with zero entries for distances to itself and identical distances to other cities.

9 Example

An approximation to the best path through 21 cities in the United Kingdom and Ireland, beginning and ending in Oxford, is sought. A lower bound is calculated internally.

9.1 Program Text

```
function h03bb_example

fprintf('h03bb example results\n\n');

% This example demonstrates the use of h03bb to find an approximation
% to the shortest return path from Oxford to 20 other cities.

% The cities in the salesman path
nc = 21;
homecity = 'Oxford';
cities = char(homecity,      'Dundee',      'Cardiff',      'Edinburgh', ...
              'Swansea',    'Perth',      'Stirling',    'Bangor', ...
              'Plymouth',  'Holyhead',  'Exeter',     'Glasgow', ...
              'Newport',   'Inverness', 'St. Davids',  'Aberdeen', ...
              'St. Asaph', 'Cambridge', 'Aberystwyth', 'Birmingham', ...
              'Dublin');

% Distance matrix data for cities
dm = zeros(nc,nc);
dm(1,2:11) = [23961  7112 21331  9050 22548 20667 13227 11617 14292  9455];
dm(2,3:11) = [      25998  4724 27936  2014  3997 20826 30488 21891 28327];
dm(3,4:11) = [          23108  2871 24325 22444 15004  8664 16359  6503];
dm(4,5:11) = [                25203  3444  3379 18093 27755 19158 25593];
dm(5,6:11) = [                    26434 24553 15169 10773 16033  8612];
dm(6,7:11) = [                        2668 19496 29159 20562 26997];
dm(7,8:11) = [                            17550 27212 18615 25051];
dm(8,9:11) = [                                19516  1895 17354];
dm(9,10:11) = [                                    20649  3135];
dm(10,11) = [                                        18537];

dm(1:11,12:21) = ...
    [19634  6394 29483  14068  28136  11052  7228  13771  4752 24111;
     5403  25281  9312  31882  4751  18651  24909  25448  20113 25289;
    21411  1263 31260  7889  29913 12829 12517  8941  7038 26178;
     3598 22547 10592 29149  8868  15918 21956 22715 17380 23484;
    23519  3372 33368  5988  32022 13917 14626  6916  9147 25852;
     4074 23951  7766 30553  6075  17322 23580 24119 18784 23960;
     2127 22005  9586 28606  8239  15375 21634 22172 16837 22013;
    16200 14308 26049 15136 24703  2447 14727  8446  9140 11714;
    25990  7981 35839 15655 34493 17409 17103 15937 11618 30467];
```

```

17383 15491 7232 16033 25886 3630 15910 9343 10323 9866;
23819 5810 33668 13484 32321 15237 14931 13766 9446 28296];

dm(12,13:21) = [21026 10985 27628 9638 14397 20655 21193 15858 20188];
dm(13,14:21) = [      30598 8276 29252 12168 11856 9064 6377 25227];
dm(14,15:21) = [      37538 9425 24307 30565 31103 25769 30945];
dm(15,16:21) = [      35803 14744 19628 6869 14149 26227];
dm(16,17:21) = [      22962 29220 29758 24423 29599];
dm(17,18:21) = [      12712 8242 7126 13457];
dm(18,19:21) = [      15366 6300 25639];
dm(19,20:21) = [      9465 18936];
dm(20,21:21) = [      20048];

% Calculate a lower bound internally and try to find lowest cost path.
bound = -1;
targc = -1;

% Initialize the random number state array
genid = nag_int(2);
subid = nag_int(53);
seed = nag_int([304950,889934,209094,23423990]);
[state,ifail] = g05kf( ...
                    genid,subid,seed);

% Find low cost return path through all cities
[path, cost, tmode, alg_stats, state, ifail] = ...
    h03bb( ...
        dm, bound, targc, state);

fprintf('Initial search end cost: %12.2f\n', alg_stats(3));
fprintf('Search best cost      : %12.2f\n', alg_stats(4));
fprintf('Initial temperature   : %12.2f\n', alg_stats(5));
fprintf('Lower bound              : %12.2f\n', alg_stats(6));
fprintf('Termination mode         : %12d', tmode);
fprintf('\nFinal cost           : %12.2f\n', cost);
fprintf('\nFinal Path:\n');
fprintf(' %s --> %s\n', homecity, cities(path(2),:));
l = numel(homecity);
sp(1,1:l+1) = ' ';
for i=3:nc-1
    fprintf(' %s --> %s\n',sp,cities(path(i),:));
end
fprintf(' %s --> %s\n',sp,homecity);

```

9.2 Program Results

h03bb example results

```

Initial search end cost:    432459.00
Search best cost          :    237068.00
Initial temperature       :    598481.00
Lower bound               :    106350.00
Termination mode         :                3
Final cost                :    131580.00

```

Final Path:

```

Oxford --> Cambridge
        --> Birmingham
        --> Glasgow
        --> Stirling
        --> Edinburgh
        --> Perth
        --> Dundee
        --> Aberdeen
        --> Inverness
        --> Holyhead
        --> Dublin
        --> Bangor
        --> St. Asaph
        --> Aberystwyth

```

--> St. Davids
--> Swansea
--> Cardiff
--> Newport
--> Exeter
--> Oxford
