Dynamic programming

Lecture 4: Examples

Giovanni Righini

Doctoral course, 2023

UNIVERSITÀ DEGLI STUDI DI MILANO

Motivation

Motivation: optimization of the operations within an Automated Storage and Retrieval System (AS/RS).

In AS/RS components and products are stored in identical (standardized) boxes (items).

Items are stored in suitable locations along an aisle and they are transported by a crane equipped with q shuttles, that moves along a rail.

The crane is initially empty at an idle point along the rail: it acts as an I/O location (*origin*).

Hence, the aisle can be represented by two lines with a common origin.

The AS/RS

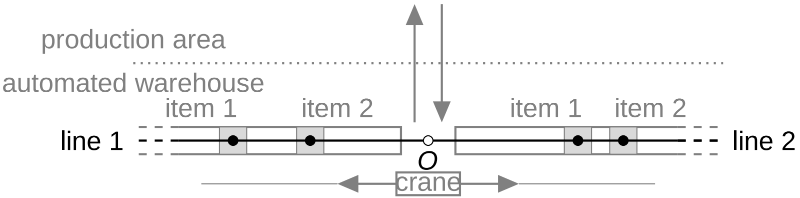

Figura: Gray: real elements (crane, rail, items, in/out site). Black: their abstract representation (origin, lines and points).

In general, each required item can be picked up on either line, because multiple identical copies of the same materials or components can be stored in the AS/RS.

Without loss of generality, each required item can be picked up from two distinct locations, one on each line.

The problem

In each trip

- the crane starts from the origin,
- it moves along one of the two lines,
- it collects at most q required items,
- it returns to the origin,
- it unloads the collected items.

Hence, the total distance travelled in a trip is twice the distance between the origin and the farthest collected item.

Objective: minimize the total distance travelled by a crane of capacity q to collect a set of required items.

Notation

$N = \{1, 2, \ldots, n\}$ is the set of items to be picked-up.

The origin is indicated by O.

Each line $\ell \in \{1, 2\}$ holds one copy of each item $i \in N$ at a given distance $d_\ell(i)$ from O.

Multiple items can be stored at a same location.

The crane capacity is a given positive integer q.

Trips and leading items

Definition (trip). A trip T is a subset of N of cardinality at most q.

The cost of a trip T on line $\ell \in \{1, 2\}$ is

$$C_\ell(T) = \max_{i \in T}\{d_\ell(i)\},$$

i.e. half the distance travelled by the crane.

Definition (leading item). A leading item of a trip T on line ℓ is a (not necessarily unique) item $j \in T$ that is farthest from O:

$$d_\ell(j) = \max_{i \in T}\{d_\ell(i)\}.$$

Hence $C_\ell(T) = d_\ell(j)$, where j is a leading item of T.

Properties

The problem requires to find a pair $(\mathcal{T}^1, \mathcal{T}^2)$ of sets of trips such that

- the trips in $\mathcal{T}^1 \cup \mathcal{T}^2$ partition N;
- the total cost $C(\mathcal{T}) = \sum_{\ell=1}^{2} \sum_{T \in \mathcal{T}^\ell} C_\ell(T)$ is minimum.

Determining a feasible solution consists of

- deciding a line assignment, i.e. determining the line where each item must be picked-up;
- grouping the items assigned to the same line into trips.

Grouping items on each line

We assume a line assignment $\mathcal{A}$ is given.

Two independent instances of the single-line problem must be solved.

Single-line problem.

Data:

- a set N of n locations to be visited on a line,
- the distance $d(i)$ from the origin for each item $i \in N$,
- the crane capacity q.

Constraints: find a partition of N into a set $\mathcal{T}$ of trips such that $|T| \leq q \; \forall T \in \mathcal{T}$.

Objective: minimize the total cost of the trips, $C(\mathcal{T})$.

Compact and complete trips

Definition (compact trips). A set $\mathcal{T}^\ell$ of trips on line $\ell \in \{1, 2\}$ is compact $\Leftrightarrow \forall T_1, T_2 \in \mathcal{T}^\ell$, either $d_\ell(i) \geq d_\ell(j)$ $\forall i \in T_1, \forall j \in T_2$ or $d_\ell(i) \leq d_\ell(j)$ $\forall i \in T_1, \forall j \in T_2$.

A solution $(\mathcal{T}^1, \mathcal{T}^2)$ is compact $\Leftrightarrow \mathcal{T}^\ell$ is compact $\forall \ell \in \{1, 2\}$.

Definition (complete trips). A set $\mathcal{T}^\ell$ of m_ℓ trips on line $\ell \in \{1, 2\}$ is complete $\Leftrightarrow$ its farthest $m_\ell - 1$ trips are made of q items each.

A solution $(\mathcal{T}^1, \mathcal{T}^2)$ is complete $\Leftrightarrow \mathcal{T}^\ell$ is complete $\forall \ell \in \{1, 2\}$.

Greedy solution of the single-line problem: optimal

The single-line sub-problem can be solved to optimality by a greedy algorithm (Brucker et al., 1998) in $O(n \log n)$:

- sort the items on each line by non-increasing distance from O;
- group them in subsets of q, starting from the farthest ones.

Observation. By construction, the solution computed by the greedy algorithm of Brucker et al. is compact and complete.

Proposition. For any given line assignment $\mathcal{A}$, any solution that is compact and complete is optimal for $\mathcal{A}$.

Proof. Trivial, by contradiction (omitted here).

Greedy solution of the line assignment sub-problem: sub-optimal

Example with $n = 6$, $q = 3$.

ℓ	A	B	C	D	E	F
1	**13**	12	11	2	1	∞
2	∞	**11**	10	∞	∞	1

Greedy solution.

Best assignment for each item.

Two trips.

Cost = $13 + 11 = 24$.

ℓ	A	B	C	D	E	F
1	**13**	12	11	**2**	1	∞
2	∞	11	10	∞	∞	**1**

Optimal solution.

Sub-optimal assignments: B, C.

Three trips.

Cost = $13 + 2 + 1 = 16$.

Implicit enumeration

The implicit complete enumeration of solutions can obtained through the implicit complete enumeration of line assignments.

There are 2^n possible line assignments.

The algorithm restricts the search to non-dominated line assignments.

Dominance between line assignments is the basis for a dynamic programming algorithm.

Graphical representation of line assignments

Item $\leftrightarrow$ edge between two points, i.e. its locations on the two lines.

Definition (intersecting edges).

Two distinct edges $i \in N$ and $j \in N$ intersect if and only if

$$(d_{\ell'}(i) \leq d_{\ell'}(j)) \ \wedge \ (d_{\ell''}(i) \geq d_{\ell''}(j)),$$

where $\ell' \neq \ell''$.

Definition (disjoint edges).

Two distinct edges are disjoint $\Leftrightarrow$ they do not intersect.

Graphical representation of line assignments

Assigning an item to a line corresponds to orienting its edge.

Definition (horizontal edges).

Edge $i \in N$ is horizontal $\Leftrightarrow d_1(i) = d_2(i)$.

Definition (orientation).

In any given line assignment

- a horizontal edge is ℓ-oriented $\Leftrightarrow$ item i is assigned to line ℓ;
- a non-horizontal edge $i \in N$ is upward- (downward-) oriented $\Leftrightarrow$ item i is assigned to the line where it is farther from (closer to) the origin.

Graphical representation

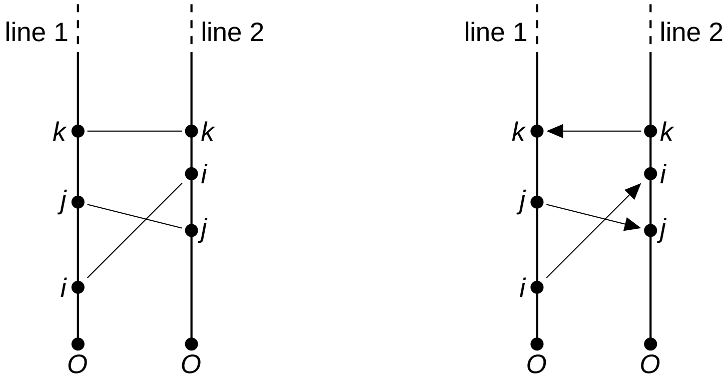

Replacement and swap

Proposition (replacement).
Consider a line $\ell \in \{1, 2\}$ and a compact and complete set $\mathcal{T}^\ell$ of trips on it; let C be its cost.
If $i \in \mathcal{T}^\ell$ is replaced by $j \notin \mathcal{T}^\ell$ with $d_\ell(j) \leq d_\ell(i)$, then the cost of any compact and complete set of trips collecting the items in $T^\ell \setminus \{i\} \cup \{j\}$ on ℓ is not larger than C.

Proof. Trivial, by contradiction (see Barbato et al.).

Proposition (swap).
Consider a solution with a compact and complete set $\mathcal{T}^\ell$ of trips on each line $\ell \in \{1, 2\}$ and let C be its cost.
If $i \in \mathcal{T}^1$ and $j \in \mathcal{T}^2$ with $d_1(j) \leq d_1(i)$ and $d_2(i) \leq d_2(j)$ are swapped, then the cost of the optimal solution corresponding to the new line assignment is not larger than C.

Proof. A swap is equivalent to two replacements.

Ordering

We introduce a lexicographic ordering of line assignments that allows to define an asymmetric and transitive dominance relation between them.

For this purpose, we introduce three quantities α, β and γ associated with line assignments, in order to break ties.

$L(i, \mathcal{A}) \in \{1, 2\}$ is the line to which item $i \in N$ is assigned in line assignment $\mathcal{A}$.

Definition. For any given line assignment $\mathcal{A}$:

$$\alpha(\mathcal{A}) = \sum_{i \in N} d_{L(i,\mathcal{A})}(i),$$

$$\beta(\mathcal{A}) = \sum_{i \in N} d_{L(i,\mathcal{A})}(i) i,$$

$$\gamma(\mathcal{A}) = \sum_{i \in N} L(i, \mathcal{A}) i.$$

An example

$$\alpha(\mathcal{A}) = \sum_{i \in N} d_{L(i,\mathcal{A})}(i) = 14$$

$$\beta(\mathcal{A}) = \sum_{i \in N} d_{L(i,\mathcal{A})}(i) i = 27$$

$$\gamma(\mathcal{A}) = \sum_{i \in N} L(i,\mathcal{A}) i = 8$$

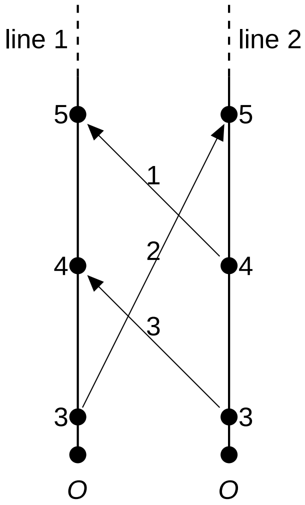

The edge reversal property

Property (edge reversal). Let $\{\ell', \ell''\} = \{1, 2\}$. Consider two distinct intersecting edges $i \in N$ and $j \in N$.

1. Non-coincident edges: $d_{\ell'}(i) \geq d_{\ell'}(j)$, $d_{\ell''}(i) \leq d_{\ell''}(j)$ and at least one of the two inequalities is strict. Assume $L(i, \mathcal{A}) = \ell'$ and $L(j, \mathcal{A}) = \ell''$. Let $\overline{\mathcal{A}}$ be obtained from $\mathcal{A}$ by reversing the orientation of both edges. Then, $C(\overline{\mathcal{A}}) \leq C(\mathcal{A})$.
   Furthermore, $\alpha(\overline{\mathcal{A}}) < \alpha(\mathcal{A})$.
2. Coincident non-horizontal edges: $d_{\ell'}(i) = d_{\ell'}(j) > d_{\ell''}(i) = d_{\ell''}(j)$ and $i > j$. Assume $L(i, \mathcal{A}) = \ell'$ and $L(j, \mathcal{A}) = \ell''$. Let $\overline{\mathcal{A}}$ be obtained from $\mathcal{A}$ by reversing the orientation of both edges. Then, $C(\overline{\mathcal{A}}) = C(\mathcal{A})$.
   Furthermore, $\alpha(\overline{\mathcal{A}}) = \alpha(\mathcal{A})$ and $\beta(\overline{\mathcal{A}}) < \beta(\mathcal{A})$.
3. Coincident horizontal edges: $d_1(i) = d_1(j) = d_2(i) = d_2(j)$ and $i > j$. Assume $L(i, \mathcal{A}) = 2$ and $L(j, \mathcal{A}) = 1$. Let $\overline{\mathcal{A}}$ be obtained from $\mathcal{A}$ by reversing the orientation of both edges. Then, $C(\overline{\mathcal{A}}) = C(\mathcal{A})$.
   Furthermore $\alpha(\overline{\mathcal{A}}) = \alpha(\mathcal{A})$, $\beta(\overline{\mathcal{A}}) = \beta(\mathcal{A})$ and $\gamma(\overline{\mathcal{A}}) < \gamma(\mathcal{A})$.

The proof directly follows from the definitions.

Dominance property

Finally, we prove that an optimal solution is certainly found even if the search is restricted to line assignments that are non-dominated.

Definition. Given two line assignments $\mathcal{A}$ and $\mathcal{A}'$, $\mathcal{A}$ *dominates* $\mathcal{A}'$ if and only if $\mathcal{A}$ is obtained from $\mathcal{A}'$ through a sequence of edge reversal operations.

Property. The dominance relation is asymmetric and transitive.

Proof. If $\mathcal{A}$ dominates $\mathcal{A}'$ then the triple $(\alpha(\mathcal{A}), \beta(\mathcal{A}), \gamma(\mathcal{A}))$ is lexicographically smaller than $(\alpha(\mathcal{A}'), \beta(\mathcal{A}'), \gamma(\mathcal{A}'))$. Since the lexicographic order is asymmetric and transitive so is the dominance relation.

Corollary. There exists at least one optimal solution whose corresponding line assignment is non-dominated.

Primary edges

The implicit complete enumeration of non-dominated line assignments must be done efficiently.

For this purpose, we introduce the definitions of primary edges and primary set and we prove that enumerating primary sets is equivalent to enumerating non-dominated line assignments.

While a line assignment requires to specify the orientation of each of the n edges, a primary set is described by a selection of a (typically small) subset of the n edges.

We prove that the selection of the primary edges is enough to determine the orientation of all edges in a non-dominated line assignment.

In turn, the definition of primary edges relies upon the definition of implications between edges.

Implication between edges

Definition (implication between non-horizontal edges).
Given two distinct edges $i \in N$ and $j \in N$, with $d_{\ell'}(i) < d_{\ell''}(i)$, i implies j if and only if the following three conditions are satisfied:

1. $d_{\ell'}(j) \geq d_{\ell'}(i)$,
2. $d_{\ell''}(j) \leq d_{\ell''}(i)$,
3. at least one of the two inequalities above is strict or $j < i$.

Definition (implication between horizontal edges).
Given two distinct edges $i \in N$ and $j \in N$, with $d_1(i) = d_2(i)$, i implies j if and only if the following three conditions are satisfied:

1. $d_1(j) = d_1(i)$,
2. $d_2(j) = d_2(i)$,
3. $j < i$.

Observation. For any two intersecting edges i and j, either i implies j or j implies i or both.
For any two disjoint edges, none of them implies the other.

Implication between edges

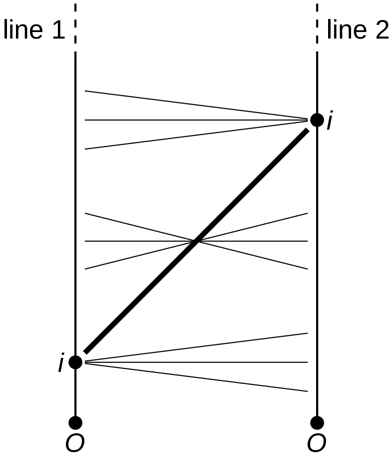

Edge *i* implies all the other edges represented in the figure.

Primary edges

Definition (primary edges).

Given a line assignment,

- a non-horizontal edge is primary if and only if
  1. it is upward-oriented
  2. it is not implied by any other upward-oriented edge;
- a horizontal edge is primary if and only if these three statements hold:
  1. it is 2-oriented,
  2. it is not implied by any upward-oriented edge,
  3. it is not implied by any 2-oriented horizontal edge.

Property.

In any line assignment all primary edges are disjoint.

Primary edges are disjoint: proof

Proof. By contradiction, assume two primary edges i and j intersect. Three cases may occur.

1. None of the edges is horizontal.
   Since they intersect, at least one of them implies the other.
   Since they are primary, they are both upward-oriented.
   Then each one is primary and is intersected by an upward-oriented edge.
   This contradicts the definition of primary edges.
2. One of the edges (w.l.o.g. j) is horizontal.
   By definition of edge implication, i implies j.
   Since i is primary, it is upward-oriented.
   Since j is primary, it is 2-oriented.
   Then j is horizontal and it is implied by an upward-oriented edge.
   This contradicts the assumption that j is primary.
3. Both edges are horizontal (w.l.o.g. i implies j).
   Since they are primary, they are 2-oriented.
   Then, j is horizontal and implied by a 2-oriented horizontal edge.
   This contradicts the assumption that j is primary.

Orientation implied by non-horizontal primary edges

By combining the definition of non-dominated line assignments with the definition of primary edges, we prove that the selection of the primary edges completely determines a non-dominated line assignment.

Property (Non-horizontal primary edge). If a non-horizontal edge $i \in N$ is primary in a non-dominated line assignment $\mathcal{A}$, then

1. edge i is upward-oriented;
2. each edge j implying i is downward-oriented;
3. each edge j implied by i is oriented to $L(i, \mathcal{A})$.

Proof. Statements 1 and 2 are implied by i being primary.
Let ℓ be the line different from $L(i, \mathcal{A})$.
If i implies j and i is upward-oriented to line $L(i, \mathcal{A})$, then $d_\ell(j) \geq d_\ell(i)$ and if the two edges coincide, then $j < i$ (by def. of implication).
By contradiction, if j is oriented to ℓ, then $\mathcal{A}$ is dominated.

Orientation implied by horizontal primary edges

Property (Horizontal primary edge). If a horizontal edge $i \in N$ is primary in a non-dominated line assignment $\mathcal{A}$, then

1. edge i is 2-oriented;
2. all non-horizontal edges implying i are downward-oriented;
3. all horizontal edges implying i are 1-oriented;
4. all horizontal edges implied by i are 2-oriented.

Proof. Statements 1, 2 and 3 are implied by i being primary.
By definition of implication, any edge j implied by i must be horizontal, coincident with i and such that $j < i$.
If j is 1-oriented, then $\mathcal{A}$ is dominated.

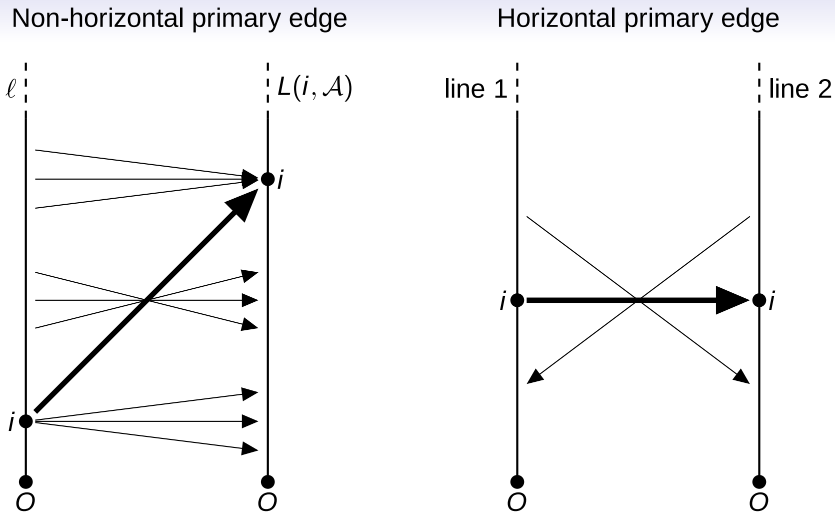

The selection of edge i as a primary edge induces the orientation of edge i itself and all edges intersecting it.

Partial order of the edges

Definition (partial order).

For each pair of distinct edges $i \in N$ and $j \in N$, *i precedes j* (indicated by $i \prec j$) if and only if $d_\ell(i) < d_\ell(j)\ \forall \ell = 1, 2$.

Observation.

For each pair of disjoint edges $i, j \in N$, either $i \prec j$ or $j \prec i$.
For each pair of intersecting edges $i, j \in N$, neither $i \prec j$ nor $j \prec i$.

Non-primary edges

We now state some properties on edge orientations implied by the assumption that an edge is *not* primary in a non-dominated line assignment.

Property (non-primary edge).

Consider two edges $i, j \in N$ with $i \prec j$ that are consecutive primary edges in a non-dominated line assignment $\mathcal{A}$, *i.e.*, there exists no primary edge $k \in N$ with $i \prec k \prec j$ in $\mathcal{A}$. Then,

- every non-horizontal edge $k \in N$ s.t. $i \prec k \prec j$ is downward-oriented;
- every horizontal edge $k \in N$ s.t. $i \prec k \prec j$ is 1-oriented.

The proof is by contradiction.

Proof (case 1)

Case 1: Assume $\exists$ a non-horizontal edge k between i and j that is upward-oriented in $\mathcal{A}$.
Since i and j are consecutive primary edges, k is not primary.
Hence, at least one condition for k to be primary is violated.
Since k is non-horizontal and upward-oriented, there must exist an upward-oriented edge e implying k.
If e and k are assigned to different lines, then $\mathcal{A}$ is dominated (by edge reversal).
Hence, e and k are assigned to the same line.
Let $\ell' = L(k, \mathcal{A}) = L(e, \mathcal{A})$ and ℓ'' the other line.
Then $d_{\ell'}(e) \geq d_{\ell'}(k) > d_{\ell'}(i)$ and $d_{\ell''}(e) \leq d_{\ell''}(k) < d_{\ell''}(j)$.
Since i and j are primary, then e cannot intersect any of them.
Therefore, e is an upward-oriented edge such that $i \prec e \prec j$.
This would force to repeat the same argument indefinitely, which is impossible since the number of edges between i and j is finite.

Proof (case 2)

Case 2: Assume $\exists$ a horizontal edge k between i and j that is 2-oriented in $\mathcal{A}$.
Since i and j are consecutive primary edges, k cannot be primary.
Hence, at least one condition for k to be primary is violated.
So, either k is implied by an upward-oriented edge or k is implied by a 2-oriented horizontal edge.
In the former case, the proof for Case 1 applies.
In the latter case, $\exists$ a 2-oriented horizontal edge e coinciding with k with $e > k$.
Then, $\exists$ a 2-oriented horizontal edge $e > k$ with $i \prec e \prec j$.
This would force to repeat the same argument indefinitely, which is impossible since the number of edges between i and j is finite.

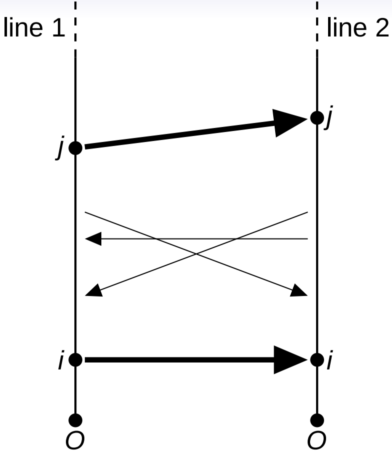

The orientation of all edges between two consecutive primary edges i and j is determined.

Primary set

As a consequence of the properties above, if the search is restricted to non-dominated line assignments, once the primary edges have been selected, the orientation of all the other edges follows.

Definition (primary set).

The primary set of a solution is the set of its primary edges.

For each subset $\mathcal{P}$ of disjoint edges, there exists a unique non-dominated line assignment $\mathcal{A}(\mathcal{P})$ having $\mathcal{P}$ as its primary set.

The enumeration of non-dominated line assignments is achieved by the enumeration of subsets of disjoint edges.

Partial line assignments

Partial line assignments correspond to partial primary sets.

We extend the dominance properties between line assignments to dominance properties betweeen partial line assignments.

The goal is to design a dynamic programming algorithm that implicitly enumerates all primary sets by iteratively adding primary edges to partial primary sets in all possible ways.

For this purpose, we exploit

- the partial order defined above
- the properties of primary edges

For each given partial primary set, it is possible to partition the edges into two subsets, such that

- the edge orientation in one of them is completely determined
- the edge orientation in the other is completely free.

Partial primary sets and edge partitions

Definition (edge partition).

For each edge $i \in N$ we define three subsets of items in which N is partitioned:

- $N^-(i) = \{j \in N : j \prec i\}$;
- $N^+(i) = \{j \in N : i \prec j\}$;
- $N^\pm(i) = N \backslash (N^-(i) \cup N^+(i))$.

Proposition (item positions).

For any non-dominated line assignment $\mathcal{A}$ in which edge $i \in N$ is a primary edge,

1. $d_{L(j,\mathcal{A})}(j) < d_{L(j,\mathcal{A})}(i) \ \ \forall j \in N^-(i)$;
2. $d_{L(j,\mathcal{A})}(j) > d_{L(j,\mathcal{A})}(i) \ \ \forall j \in N^+(i)$;
3. $d_{L(j,\mathcal{A})}(j) \leq d_{L(j,\mathcal{A})}(i) \ \ \forall j \in N^\pm(i)$.

Proof

Statements 1 and 2 directly follow from the definitions.

Statement 3 follows from properties of primary edges and implications: at least one of two intersecting edges must imply the other.

If edge i is non-horizontal and primary, then it is upward-oriented. If edge i implies $j \in N^{\pm}(i)$ and i is primary, then $L(j, \mathcal{A}) = L(i, \mathcal{A})$ and hence, owing to the implication, $d_{L(j,\mathcal{A})}(j) \leq d_{L(j,\mathcal{A})}(i)$.

If edge $j \in N^{\pm}(i)$ implies i and i is non-horizontal and primary, then j is downward-oriented. Owing to the implication, $d_{L(j,\mathcal{A})}(j) \leq d_{L(j,\mathcal{A})}(i)$.

If edge i is horizontal and $j \in N^{\pm}(i)$ is non-horizontal, then j implies i. Hence, if i is primary then j is downward-oriented and, owing to the implication, $d_{L(j,\mathcal{A})}(j) \leq d_{L(j,\mathcal{A})}(i)$.

If edge i is horizontal and $j \in N^{\pm}(i)$ is horizontal, then j coincides with i and therefore $d_{L(j,\mathcal{A})}(j) = d_{L(j,\mathcal{A})}(i)$.

Independent subsets

For any line assignment $\mathcal{A}$,

$$N_\ell(\mathcal{A}) = \{j \in N : L(j, \mathcal{A}) = \ell\} \;\; \forall \ell = 1, 2.$$

For any given primary item $i \in N$ in $\mathcal{A}$, consider the partition of $N_\ell(\mathcal{A})$ into two subsets:

$$S_\ell(i, \mathcal{A}) = \{j \in N_\ell(\mathcal{A}) : d_\ell(j) > d_\ell(i)\}$$

$$R_\ell(i, \mathcal{A}) = \{j \in N_\ell(\mathcal{A}) : d_\ell(j) \leq d_\ell(i)\}.$$

Property (independent subsets). If $\mathcal{A}$ is non-dominated,

- the elements in $S_\ell(i, \mathcal{A})$ are determined only by the orientation of the edges in $N^+(i)$;
- the elements in $R_\ell(i, \mathcal{A})$ are determined only by the orientation of the edges in $N^-(i) \cup N^\pm(i)$.

Proof. This property immediately follows from the previous one and the definitions of $S_\ell(i, \mathcal{A})$ and $R_\ell(i, \mathcal{A})$.

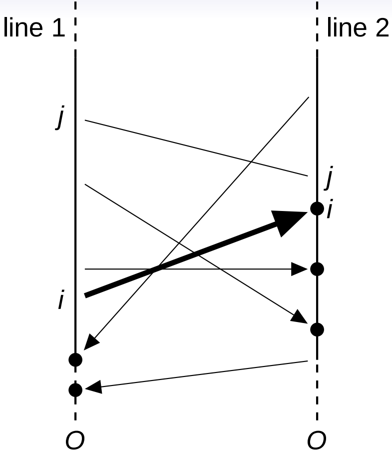

When a partial primary set is defined up to edge *i*, all locations closer to *O* than the endpoints of edge *i* are determined (black dots), while all locations farther from *O* than the endpoints of edge *i* are undetermined (e.g. edge *j*).

Partial line assignment

Definition (partial line assignment).

For any given edge $i \in N$, a partial line assignment $\mathcal{A}_i$ is an assignment to the lines of all items in $N^-(i) \cup N^\pm(i)$ so that i is the last (farthest) primary edge.

The corresponding partial primary set $\mathcal{P}_i$ is the set of primary items of $\mathcal{A}_i$.

When a partial primary set is defined up to edge i, the orientation of all edges in $R_\ell(i, \mathcal{A})$ is defined $\forall \ell$, while that of all edges in $S_\ell(i, \mathcal{A})$ is unconstrained $\forall \ell$.

The argument cannot be reversed to iteratively construct partial line assignments from the farthest edges to O.

Consecutive primary edges

Let us introduce

- a dummy edge 0 preceding all edges in N,
- a dummy edge $n+1$ preceded by all edges in N.

Let $i \in \{0, \dots, n\}$ and $j \in \{1, \dots, n+1\}$ be consecutive primary edges in a non-dominated line assignment $\mathcal{A}$, with $i \prec j$.

Let $N_\ell^{ij}(\mathcal{A})$ be the set of edges in $(N^-(j) \cup N^\pm(j)) \cap N^+(i)$ that are oriented to line ℓ in $\mathcal{A}$, that is,

$$N_\ell^{ij}(\mathcal{A}) = S_\ell(i, \mathcal{A}) \cap R_\ell(j, \mathcal{A}) \;\; \forall i \prec j.$$

$N_\ell^{0j}(\mathcal{A}) = R_\ell(j, \mathcal{A}) \; \forall j = 1, \dots, n$
$N_\ell^{i,n+1}(\mathcal{A}) = S_\ell(i, \mathcal{A}) \; \forall i = 1, \dots, n$
$N_\ell^{0,n+1}(\mathcal{A}) = N_\ell(\mathcal{A})$.

Consecutive primary edges

Property (assignments between consecutive primary edges).

In any non-dominated line assignment $\mathcal{A}$ in which $i \in N \cup \{0\}$ and $j \in N \cup \{n+1\}$ with $i \prec j$ are consecutive primary edges, the elements in $N_\ell^{ij}(\mathcal{A})$ are determined only by the primary edges i and j.

Proof.

The set $N_\ell^{ij}(\mathcal{A})$ is the intersection between $S_\ell(i, \mathcal{A})$ and $R_\ell(j, \mathcal{A})$.
$S_\ell(i, \mathcal{A})$ only depends on the orientation of the edges in $N^+(i)$.
$R_\ell(j, \mathcal{A})$ only depends on the orientation of the edges in $N^-(j) \cup N^\pm(j)$.
Since i and j are consecutive primary edges in $\mathcal{A}$, no other primary edges can be contained in $(N^-(j) \cup N^\pm(j)) \cap N^+(i)$.

Therefore, the elements of each set $N_\ell^{ij}(\mathcal{A})$ can be pre-computed, independently of $\mathcal{A}$.

Leading items in partial assignments

The cost of a set of trips on a line depends on their leading items.

In turn, the leading items are determined by the greedy algorithm starting from the farthest ones.

Unfortunately, the construction of partial primary sets must proceed starting from the closest items.

Hence, it is not possible to determine the cost implied by the oriented edges in a partial primary set, because it is not known which items among them are leading in their trips.

However, the number of possibilities is only q.

For any given partial primary set, the total cost corresponding to the leading items in the subset of oriented edges may have q distinct values.

Therefore a dynamic programming algorithm may associate q distinct states with each partial primary set.

Residual items in partial line assignments

Given a non-dominated line assignment $\mathcal{A}$ and a primary edge $i \in N$, we define the number of residual items on each line at edge i as

$$r_\ell(i, \mathcal{A}) = |S_\ell(i, \mathcal{A})| \bmod q \quad \forall \ell = 1, 2.$$

Let $\mathcal{T}(\mathcal{A})$ be the compact and complete solution obtained from line assignment $\mathcal{A}$ through the greedy algorithm.

Its cost is the sum of two contributions for each line ℓ:

- the cost of the trips with their leading item in $S_\ell(i, \mathcal{A})$
- the cost of the trips with their leading item in $R_\ell(i, \mathcal{A})$.

Assume that $S_\ell(i, \mathcal{A})$ and $R_\ell(i, \mathcal{A})$ are represented as vectors indexed from 1, sorted by non-increasing distances from O on line ℓ.

Denoting by $S_\ell(i, \mathcal{A})[t]$ and $R_\ell(i, \mathcal{A})[t]$ the t-th entry of such vectors the sets of leading items in $S_\ell(i, \mathcal{A})$ and $R_\ell(i, \mathcal{A})$ are

$$\mathcal{L}^S_\ell(i, \mathcal{A}) = \{S_\ell(i, \mathcal{A})[t] : t \bmod q = 1\}$$

$$\mathcal{L}^R_\ell(i, \mathcal{A}) = \{R_\ell(i, \mathcal{A})[t] : (t + r_\ell) \bmod q = 1\}.$$

Cost of partial line assignments

The first cost term is

$$C_\ell(S_\ell(i,\mathcal{A})) = \sum_{k \in \mathcal{L}_\ell^S(i,\mathcal{A})} d_\ell(k).$$

This sum includes the cost terms given by the edges in $N^+(i)$ and it does not depend on the orientation of the edges in $N^-(i) \cup N^\pm(i)$.

The second cost term is

$$C_\ell(R_\ell(i,\mathcal{A}), r_\ell) = \sum_{k \in \mathcal{L}_\ell^R(i,\mathcal{A})} d_\ell(k).$$

This sum includes the cost terms given by the edges in $N^-(i) \cup N^\pm(i)$ and it does not depend on the orientation of the edges in $N^+(i)$, but only on the number of residual items $r_\ell(i,\mathcal{A})$.

Cost of partial line assignments

Setting

$$C^{+}(i,\mathcal{A}) = \sum_{\ell=1}^{2} C_\ell(S_\ell(i,\mathcal{A}))$$

$$C^{-}(i,\mathcal{A},r_1,r_2) = \sum_{\ell=1}^{2} C_\ell(R_\ell(i,\mathcal{A}),r_\ell)$$

the cost $C(\mathcal{A})$ of a solution $\mathcal{T}(\mathcal{A})$ is

$$C(\mathcal{A}) = C^{+}(i,\mathcal{A}) + C^{-}(i,\mathcal{A},r_1,r_2).$$

Extension.

A line assignment $\mathcal{A}$ *extends* $\mathcal{A}_i$ if and only if the items in $N^{-}(i) \cup N^{\pm}(i)$ are assigned to the same lines in both $\mathcal{A}$ and $\mathcal{A}_i$.

Partial line assignments and dominance

Property (dominance between partial line assignments).

For a given $i \in N$, let $\mathcal{A}_i'$ and $\mathcal{A}_i''$ be two partial line assignments.
Let $r_\ell = r_\ell(i, \mathcal{A}_i') = r_\ell(i, \mathcal{A}_i'')$ for $\ell = 1, 2$.
Let $\mathcal{A}''$ be a complete line assignment extending $\mathcal{A}_i''$.
Then, if $C^-(i, \mathcal{A}_i', r_1, r_2) < (=) C^-(i, \mathcal{A}_i'', r_1, r_2)$, then $\exists$ a complete line assignment $\mathcal{A}'$ extending $\mathcal{A}_i'$ such that $C(\mathcal{A}') < (=) C(\mathcal{A}'')$.

Proof.

Construct $\mathcal{A}'$ by orienting the edges in $N^-(i) \cup N^\pm(i)$ as in $\mathcal{A}_i'$ and the edges in $N^+(i)$ as in $\mathcal{A}''$.
Then, $\mathcal{A}'$ extends $\mathcal{A}_i'$.
By construction, $S_\ell(i, \mathcal{A}') = S_\ell(i, \mathcal{A}'')$.
Hence $|S_\ell(i, \mathcal{A}'')| \mod q = r_\ell \ \forall \ell \Rightarrow |S_\ell(i, \mathcal{A}')| \mod q = r_\ell \ \forall \ell$.
Then, $\mathcal{A}'$ and $\mathcal{A}''$ have the same leading items in $N^+(i)$:
$\mathcal{L}_\ell^S(i, \mathcal{A}'') = \mathcal{L}_\ell^S(i, \mathcal{A}')$ for $\ell = 1, 2$.
Then $C^+(i, \mathcal{A}'') = C^+(i, \mathcal{A}')$.

Implicit enumeration by dynamic programming

An optimal solution can be found by (implictly) enumerating all primary sets.

The algorithm iteratively builds primary sets by adding a primary item at each iteration. At intermediate iterations this generates partial line assignments.

Dominated partial line assignments can be discarded without losing the optimality guarantee.

Bellman optimality principle (basis for a dynamic programming algorithm): a sub-policy characterized by $\mathcal{P}'_i$ dominates a sub-policy characterized by $\mathcal{P}''_i$ for a given pair of residual values (r_1, r_2) if and only if $C^-(i, \mathcal{A}'_i, r_1, r_2) < C^-(i, \mathcal{A}''_i, r_1, r_2)$.

In case of tie, any arbitrary criterion can be used to discard one of the two equivalent partial assignments (e.g. the lexicographic ordering of assignments).

States

A state of the dynamic programming algorithm is a triple $\{i, r_1, r_2\}$, with $i \in N \cup \{0, n+1\}$ and r_1 and r_2 satisfying

$$(r_1 + r_2) \bmod q = \rho_i,$$

where $\rho_i = |N^+(i)| \bmod q$.

Hence, there are q distinct states for each item $i \in N \cup \{0\}$ and the following relations allow to obtain r_1 from r_2 and vice versa:

$$r_1 = (\rho_i - r_2) \bmod q$$
$$r_2 = (\rho_i - r_1) \bmod q.$$

A state $\{i, r_1, r_2\}$ with $i \in N$ corresponds to a partial line assignment $\mathcal{A}_i$ where i is the farthest primary item.

Initial and final states

Triples of the form $\{0, r_1, r_2\}$ are initial states: in this case r_1 and r_2 indicate the number of residual items on each line before orienting any edge.

These q initial states are given by the q possible values of (r_1, r_2) pairs such that $(r_1 + r_2) \mod q = n \mod q$, that is, $\{(0, \rho), \ldots, (\rho, 0), (\rho + 1, q - 1), \ldots, (q - 1, \rho + 1)\}$, where $\rho = n \mod q$.

The final state$\{n + 1, 0, 0\}$ corresponds to a full line assignment without residual items; hence, it must be reached to guarantee that on each line the leading item of the farthest trip is a farthest item.

Extension rule

States are iteratively extended according to the partial order of the items, with the addition of items 0 and $n+1$.

Extending a state from a predecessor state $\{i, r_1, r_2\}$ to a successor state $\{j, r'_1, r'_2\}$ means extending the partial line assignment $\mathcal{A}_i$ corresponding to $\{i, r_1, r_2\}$ to the partial line assignment $\mathcal{A}_j$ such that

- edge j is primary in $\mathcal{A}_j$
- no primary edge exists between the primary edges i and j.

$$r_\ell = (r'_\ell + |N^{ij}_\ell|) \bmod q \quad \forall \ell \in \{1, 2\}.$$

So, a one-to-one correspondence is established between the q states of edge i and the q states of edge j for each (i, j) pair such that $i \prec j$.

For each state $\{j, r'_1, r'_2\}$, $Pred(j, r'_1, r'_2)$ is the set of its predecessor states:

$$Pred(j, r'_1, r'_2) = \{\{i, r_1, r_2\} : (i \prec j) \wedge (r_\ell = r'_\ell + |N^{ij}_\ell| \bmod q \;\; \forall \ell \in \{1, 2\})\}.$$

Cost extension

Each state $\{i, r_1, r_2\}$ has an associated cost $C(i, r_1, r_2)$, that is the minimum cost of a partial solution corresponding to the state.

$$C(0, r_1, r_2) = 0 \ \forall (r_1, r_2) : (r_1 + r_2) \bmod q = \rho.$$

$$C(j, r'_1, r'_2) = \min_{\{i, r_1, r_2\} \in Pred(j, r'_1, r'_2)} \{C(i, r_1, r_2) + \Delta(i, j, r'_1, r'_2)\}.$$

The cost increase $\Delta(i, j, r'_1, r'_2)$ is the sum of the distances of the leading items in N_1^{ij} and N_2^{ij}:

$$\Delta(i, j, r'_1, r'_2) = \sum_{\ell=1}^{2} \sum_{k \in \mathcal{L}_\ell^{ij}} d_\ell(k).$$

These depend on the values of r'_1 and r'_2.

Cost extension

For each line $\ell \in \{1, 2\}$, consider an array made of the edges in N_ℓ^{ij} indexed from 1 and sorted by non-increasing distance from O.

Let $N_\ell^{ij}[t]$ be the edge in position t in the array.

Then the set of leading items in N_ℓ^{ij} is

$$\mathcal{L}_\ell^{ij} = \{k = N_\ell^{ij}[t] : (t + r'_\ell) \bmod q = 1\}.$$

Complexity

States: There are $O(nq)$ states.

Extensions: the number of (i, j) pairs such that $i \prec j$ is $O(n^2)$.

Predecessors: q extensions are done for each (i, j) pair.
Each state of $j \in N \cup \{n+1\}$ has a unique predecessor among the states of $i \prec j$.

Pre-computing all N_ℓ^{ij} takes at most $O(n^3)$.

Therefore

- the asymptotic worst-case time complexity to pre-compute N_ℓ^{ij} is $O(n^3)$;
- the asymptotic worst-case time complexity of the algorithm is $O(n^2 q)$, which is bounded by $O(n^3)$.

Main ideas

Define local moves that can improve solutions.
Define as dominated the solutions that can be improved.
Restrict the enumeration to non-dominated solutions.

Consider the locally optimal decisions of a greedy algorithm.
Consider the effects of locally non-optimal decisions in non-dominated solutions.

Enumerate the sequences of locally non-optimal decisions.

For each possible pair of locally non-optimal decisions, define its cost.
Accept the enumeration of all possible cases if needed.

Nodes: locally non-optimal decisions.
Arcs: pairs of consecutive locally non-optimal decisions.