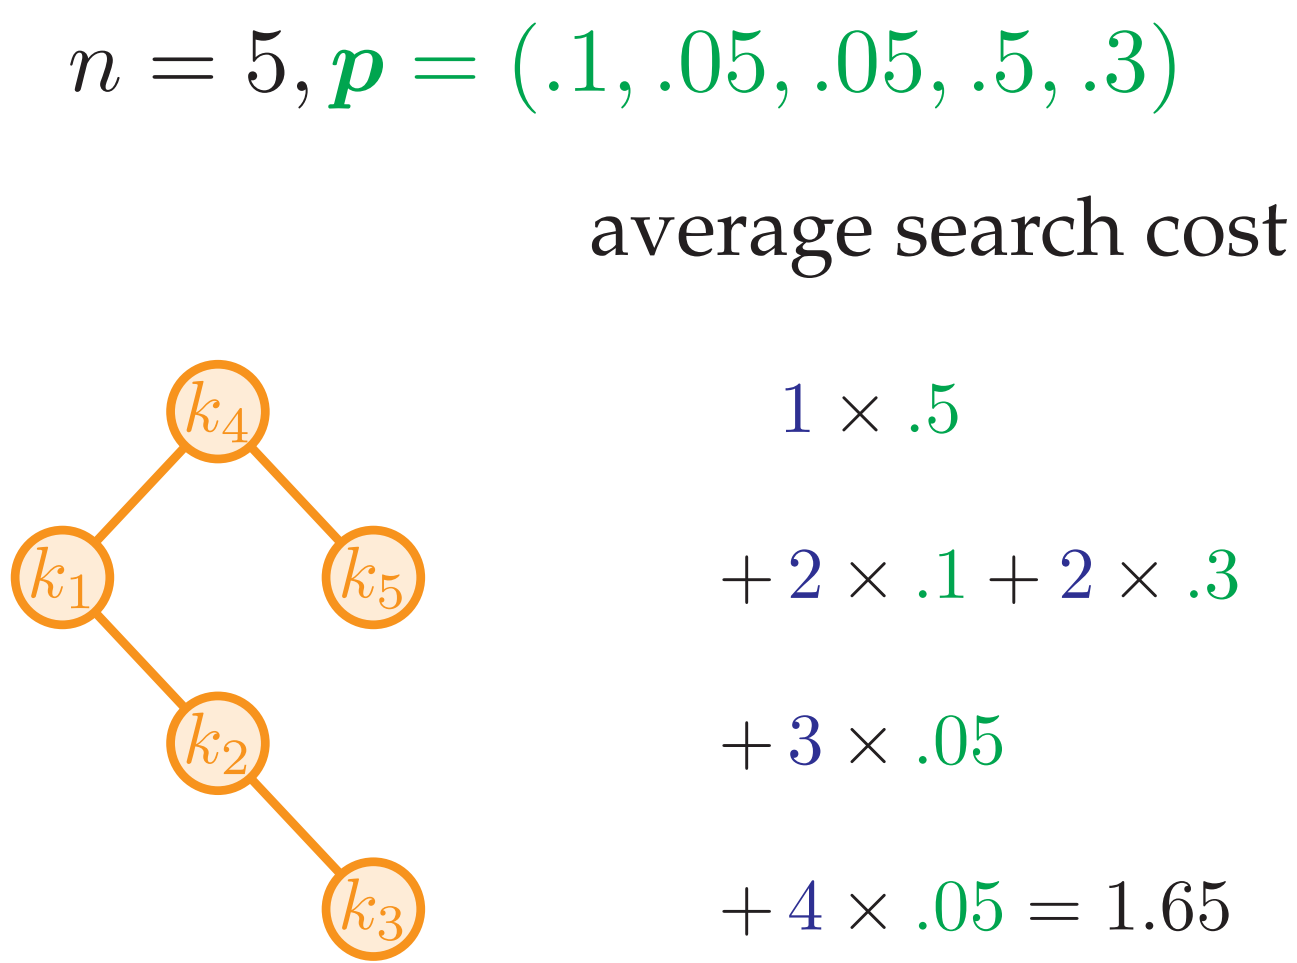


An Example of the Prediction Game

Fix a set of n ordered keys $k_1 < \dots < k_n$. In each of the T trials

1. The learner randomly predicts with a *Binary Search Tree (BST)*
2. The adversary reveals the *search probabilities* for all the keys $p \in [0,1]^n$ s.t. $\sum_i p_i = 1$
3. The learner incurs a loss of *average search cost* of the predicted BST i.e. $\mathbb{E}[\sum_{i=1}^n \text{depth}(k_i) \cdot p_i]$



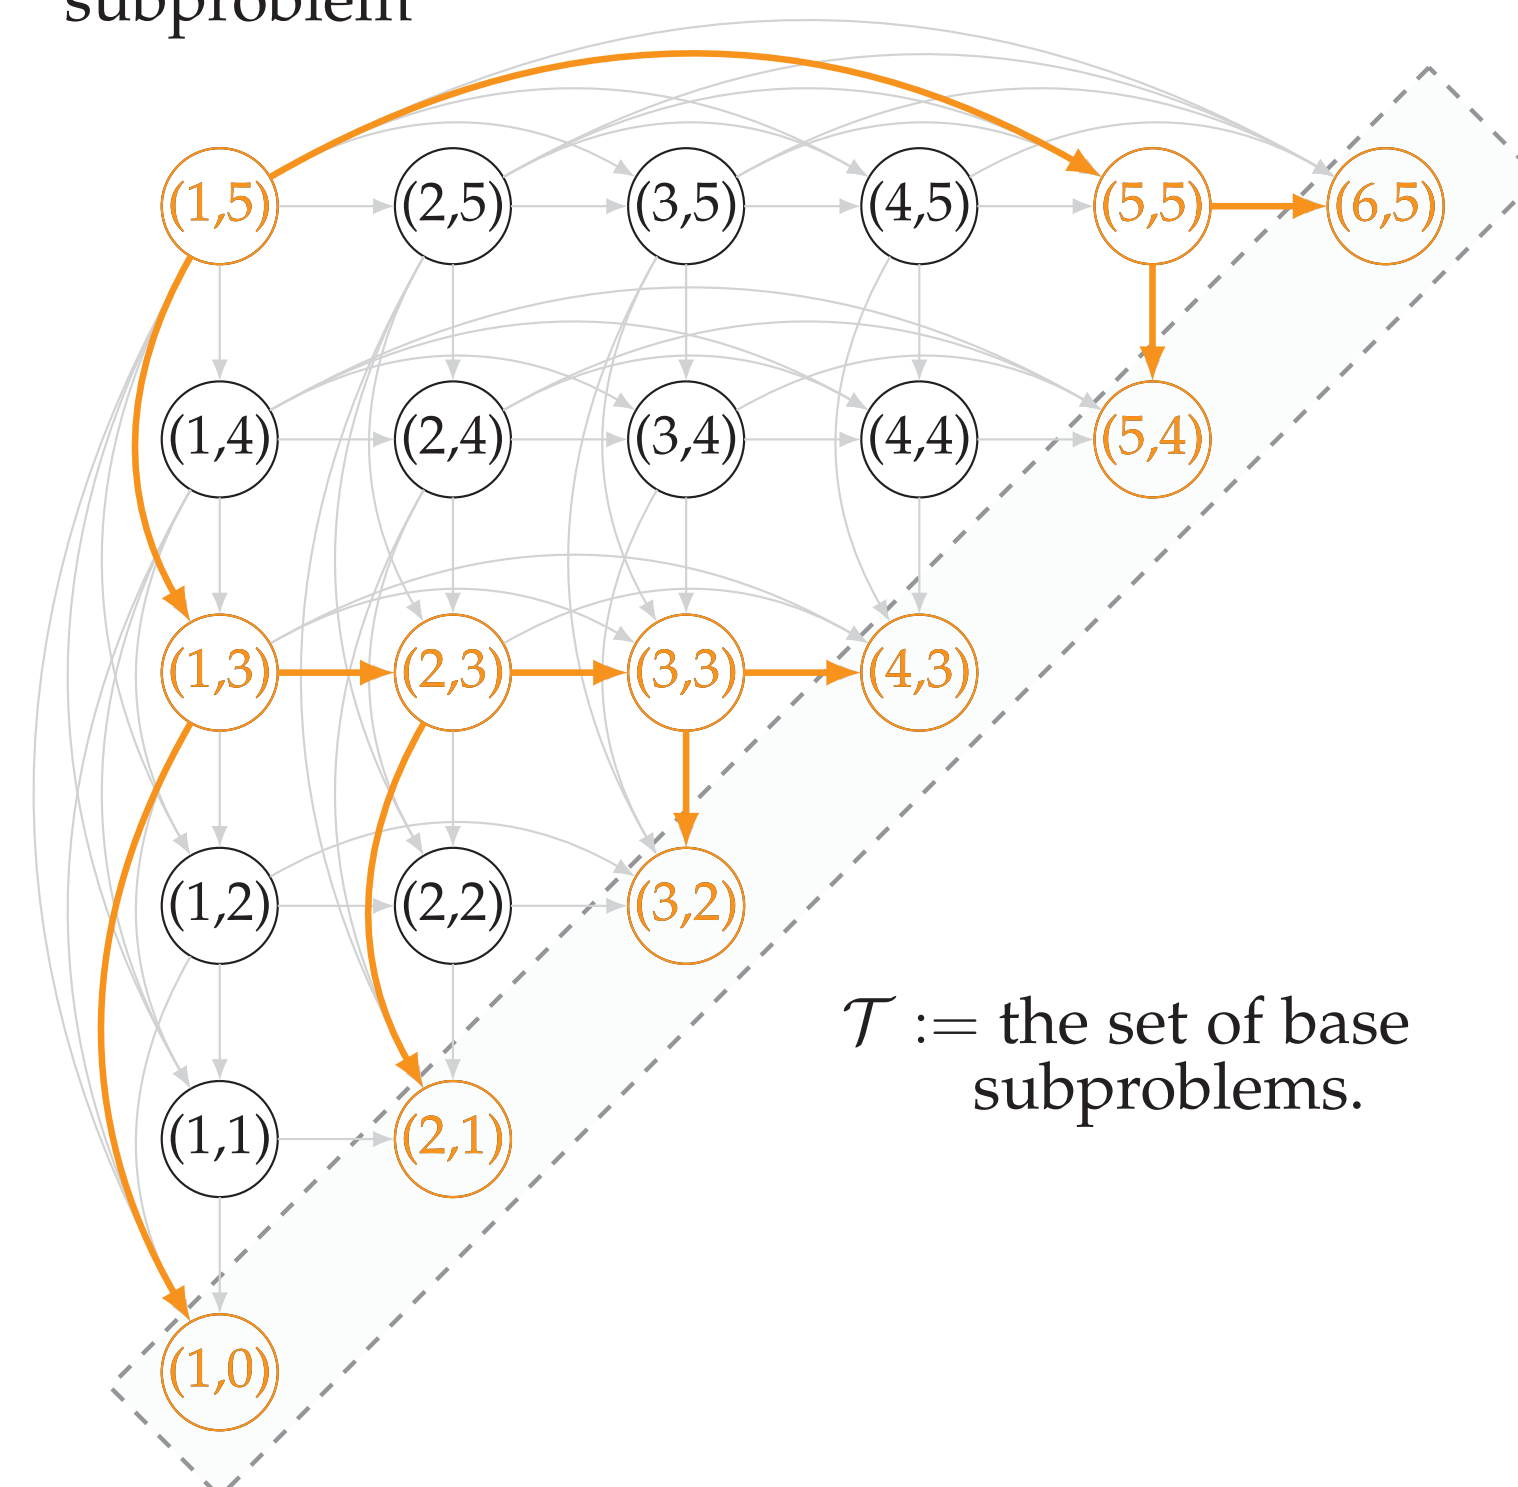
The Dynamic Programming Representation

DAG of Subproblems

- Vertices V encode the subproblems
- Edges E encode the dependency between subproblems: (i, j) to $(i, r-1)$ and $(r+1, j)$ for $i \leq r \leq j$

$$\text{The subproblem with } k_i \dots k_j \quad \text{OPT}(i, j) = \begin{cases} 0 & j = i - 1 \\ \min_{i \leq r \leq j} \{ \text{OPT}(i, r-1) + \text{OPT}(r+1, j) + \sum_{k=i}^j p_k \} & i \leq j \end{cases}$$

s := the final subproblem



Multiedges encode the recursive calls

- A **multiedge** is the pair of edges of form (i, j) to $(i, r-1)$ and $(r+1, j)$
- The associated loss with any outgoing multiedge m from (i, j) is defined as $\ell_m = \sum_{k=i}^j p_k$

Multipaths encode the solutions

- A **multipath** starts with a single multiedge at s , continues with successor multiedges at the internal nodes, and ends at \mathcal{T}

Benefits of multipath representation:

- Every BST can be represented by a multipath and vice versa
- Linear loss over the components (i.e. multiedges)
- Works for a large class of min-sum dynamic programmings (M_v := the set of outgoing multiedges from vertex v)

$$\text{OPT}(v) = \min_{m \in M_v} \left\{ \ell_m + \sum_{u: (v,u) \in m} \text{OPT}(u) \right\} \quad \text{where for all vertices } v \in V \begin{cases} 1. |M_v| = k \text{ for a fixed } k \\ 2. m\text{'s in } M_v \text{ do not overlap} \end{cases}$$

Algorithm #1 – Expanded Hedge (EH)

Maintain a **distribution** W on all multipaths. With each **multiedge** m , associate a weight w_m . Exploiting the linear loss, let W be in *canonical product form* so that the multiplicative update factorizes over multiedges and sampling becomes easy:

Weight Space – Canonical Product Form

1. $W(\pi) = \prod_{m \in \pi} w_m$
2. $\sum_{m \in M_v} w_m = 1$ for all $v \in V - \mathcal{T}$
3. $\sum_{\pi} W(\pi) = 1$

Prediction Sample from W using 1 and 2 above

Multiplicative update for each multipath π results in multiplicative updates per multiedge m :

$$\text{Update} \quad W^{\text{new}}(\pi) \propto W(\pi) \exp(-\eta \pi \cdot \ell) = \prod_{m \in \pi} \underbrace{w_m \exp[-\eta \ell_m]}_{:= \hat{w}_m}$$

Going back to the canonical product form:

Projection – Generalized Weight Pushing

Goal: For each multiedge m , find factor f_m s.t. the new weights $w_m^{\text{new}} := \hat{w}_m f_m$ are in canonical product form.

- Solution:**
1. $Z_v := 1, \forall v \in \mathcal{T}$
 2. $Z_v := \sum_{m \in M_v} \hat{w}_m \prod_{u: (v,u) \in m} Z_u, \forall v \in V - \mathcal{T}$
 3. $f_m := (\prod_i Z_{u_i}) / Z_v$ for each multiedge m from v to u_1 and u_2

Regret Bounds

$$\mathbb{E}[L_{\text{EH}}] - L_{\text{best}} \leq \sqrt{2 \log(4)} n^{\frac{3}{2}} \sqrt{T} + n^2 \log(4)$$

Other Combinatorial Objects

- CH has better bounds, but it is hard to obtain the projections.

Problem	FPL	EH	CH
Optimal Binary Search Trees	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n (\log n)^{\frac{1}{2}} \sqrt{T})$
Matrix-Chain Multiplications	—	$\mathcal{O}(n^{\frac{3}{2}} D^3 \sqrt{T})$	$\mathcal{O}(n (\log n)^{\frac{1}{2}} D^3 \sqrt{T})$
Knapsack	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n (\log n C)^{\frac{1}{2}} \sqrt{T})$
Rod Cutting	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n (\log n)^{\frac{1}{2}} \sqrt{T})$
Weighted Interval Scheduling	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n^{\frac{3}{2}} \sqrt{T})$	$\mathcal{O}(n (\log n)^{\frac{1}{2}} \sqrt{T})$

* D := the maximum dimensionality of the matrices in the matrix-chain multiplication problem


* C := the capacity of the knapsack in the knapsack problem

Algorithm #2 – Component Hedge (CH)

We can also work with **edges** as components. With each edge e , associate a weight w_e . Arbitrarily distribute the loss ℓ_m of each multiedge m over its edges. Having a linear loss, it is sufficient to maintain a **mean vector** w of multipaths:

$$\mathbb{E}[\pi \cdot \ell] = \mathbb{E}[\pi] \cdot \ell = w \cdot \ell$$

w will be in a polytope with a small number of facets:



Weight Space – 2-Flow Polytope

1. $w_{\text{out}}(s) = 2$
2. For any multiedge $m = \{e, e'\}, w_e = w_{e'}$
3. $w_{\text{out}}(v) = 2 \times w_{\text{in}}(v)$, for each vertex $v \in V - \mathcal{T} - \{s\}$

We sample with the same expectation as w :

Prediction Decompose w into a convex combination of multipaths. Decomposition is done iteratively via a greedy approach which removes one multipath at a time.

Multiplicative update per edge:

Update For each edge $e \in E, \hat{w}_e = w_e \exp(-\eta \ell_e)$

Going back to the polytope with relative entropy projection:

Projection – Iterative Bregman Projections

Goal: Find $w^{\text{new}} := \arg \min_{w \in \text{2-flow polytope}} \Delta(w || \hat{w})$.

Solution: Project onto the flow constraint of a particular vertex and then repeatedly cycle through the vertices

Regret Bounds

$$\mathbb{E}[L_{\text{CH}}] - L_{\text{best}} \leq 4\sqrt{2} n (\log n)^{\frac{1}{2}} \sqrt{T} + 16 n \log n$$