**Algorithm 23.3 Iterated optimization for influence diagrams with acyclic relevance graphs**

```
Procedure Iterated-Optimization-for-IDs (
    𝓘,   // Influence diagram
    𝒢    // Acyclic relevance graph for 𝓘
  )
    Let D_1, ..., D_k be an ordering of 𝒟 that is a topological ordering for 𝒢
    Let σ^0 be some fully mixed strategy for 𝓘
    for i = 1, ..., k
      Choose δ_{D_i} to be locally optimal for σ^{i-1}
      σ^i ← (σ^{i-1}_{-D_i}, δ_{D_i})
    return σ^k
```

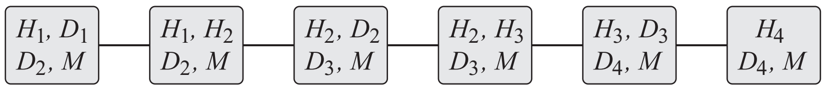


**Figure 23.9 Clique tree for the imperfect-recall influence diagram of figure 23.5.** Although the network has many cascaded decisions, our ability to "forget" previous decisions allows us to solve the problem using a bounded tree-width clique tree.

are already stable, and so will never need to change; for $j > i$, the decision rules for $D_j$ are irrelevant, so that changing them will not require revisiting $D_i$.

One subtlety with this argument relates, once again, to the issue of probability-zero events. If our arbitrary starting strategy $\sigma$ assigns probability zero to a certain decision $d \in Val(D)$ (in some setting), then the local optimization of another decision rule $D'$ might end up selecting a suboptimal decision for the zero probability cases. If subsequently, when optimizing the decision rule for $D$, we ascribe nonzero probability to $D = d$, our overall strategy will not be optimal. To avoid this problem, we can use as our starting point any fully mixed strategy $\sigma$. One obvious choice is simply the strategy that, at each decision $D$ and for each assignment to $\mathrm{Pa}_D$, selects uniformly at random between all of the possible values of $D$.

The overall algorithm is shown in algorithm 23.3.

**Theorem 23.5** *Applying* Iterated-Optimization-for-IDs *on an influence diagram $\mathcal{I}$ whose relevance graph is acyclic, returns a globally optimal strategy for $\mathcal{I}$.*

The proof is not difficult and is left as an exercise (exercise 23.8).

Thus, this algorithm, by iteratively optimizing individual decision rules, finds a globally optimal solution. The algorithm applies to any influence diagram whose relevance graph is acyclic, and hence to any influence diagrams satisfying the perfect recall assumption. Hence, it is at least as general as the variable elimination algorithm of section 23.3. However, as we saw, some influence diagrams that violate the perfect recall assumption have acyclic relevance graphs nonetheless; this algorithm also applies to such cases.

**Example 23.22** *Consider again the influence diagram of example 23.11. Despite the lack of perfect recall in this*