Polar Cuts from Relaxed Decision Diagrams

Christian Tjandraatmadja, Willem-Jan van Hoeve
Tepper School of Business, Carnegie Mellon University

Carnegie Mellon University

Can we reach strong cuts through a different relaxation than an LP relaxation?

Dynamic programming formulations can be turned into discrete relaxations: **relaxed decision diagrams**.

Bergman et al.

Polar cuts are derived from a **vertex description of a polyhedron**, such as a relaxation or projection of the integer hull.

In previous work, Becker et al. used decision diagrams as an optimization oracle to generate cuts through Lagrangian relaxation.

We revisit cuts from decision diagrams with the old idea of **polarity** and the new idea of **relaxed decision diagrams**...

Decision diagrams

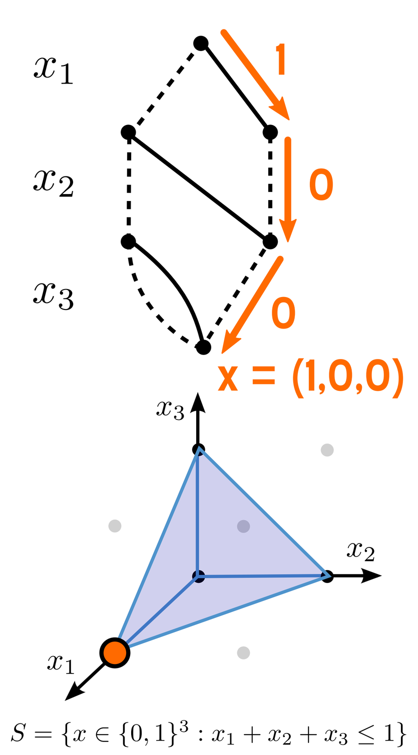

Decision diagrams are graphical representations of sets of points with finite domain.

Main property: One-to-one correspondence between each path from root to terminal and each point of the set.

Relaxed decision diagrams, proposed by Andersen et al., are discrete relaxations (i.e. a superset of the feasible set) that are tractable in this form, constructed by limiting its width.

Polar cuts

(also known as target cuts)

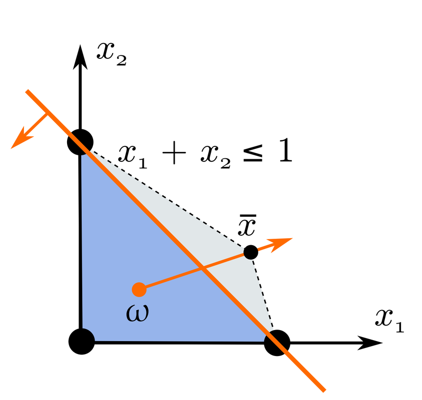

An elegant geometrical interpretation:
Shoot a ray from an interior point ω towards $\bar{x}$ and take a facet that intersects it.

Recently revived by Buchheim et al. under the name target cuts, polar cuts rely on the polar set:

$$S^\circ = \{u \in \mathbb{R}^n : u^\top x \leq 1 \; \forall x \in S\}$$

Proposition: Let P be a full-dimensional polytope. Let ω be an interior point of P, $\bar{x} \notin P$, and

$$u^* = \arg\max_u \{u^\top(\bar{x} - \omega) : u \text{ is an extreme point of } (P - \omega)^\circ\}.$$

Then $u^{*\top} x \leq 1 + u^{*\top}\omega$ is a facet-defining cut for $\bar{x}$ w.r.t. P.

How to generate polar cuts from decision diagrams?

Let D be a decision diagram with vertices V, root s, and terminal t representing the set S.

Define $P^*(D) = \{(u, v) \in \mathbb{R}^n \times \mathbb{R}^{|V|} : v_j \leq v_i - \ell u_k \quad \forall$ arc (i, j) of layer k with label ℓ,
$v_s = 1, \; v_t = 0\}$.

Theorem: $\text{proj}_u(P^*(D)) = S^\circ$.

Polar cuts are very expensive to compute, so a relaxation is necessary. Relaxed decision diagrams allow us to choose a trade-off between strength of the cut and effort to find a cut.

We need $(P - \omega)^\circ$ to generate a polar cut.

Simply replace $v_s = 1$ by $v_s = 1 + u^\top\omega$.

But an extreme point of $P^*(D)$ may not be an extreme point of $\text{proj}_u(P^*(D))$!

Indeed, and this is non-trivial. Experiments indicate that reoptimizing with a perturbed objective over the optimal face results in a cut as strong as a facet-defining cut. (See insights to the right.)

A side corollary...

This theorem is not only useful for cut generation, but it also provides insight on the complexity of decision diagrams.

Corollary: If D is a decision diagram representing S, then the set of points of any face of $\text{conv}(S)$ is representable by a sub-DD of D.

In other words: we cannot expect to construct a compact decision diagram unless every face is also compact!

A proof of concept: the maximum independent set problem

The maximum independent set problem is to find a largest set of pairwise nonadjacent vertices in a graph. We test our cuts on high density (80%) random graphs (average of 10 instances), using a clique cover IP formulation. Greater widths mean a tighter relaxation but more expensive cuts.

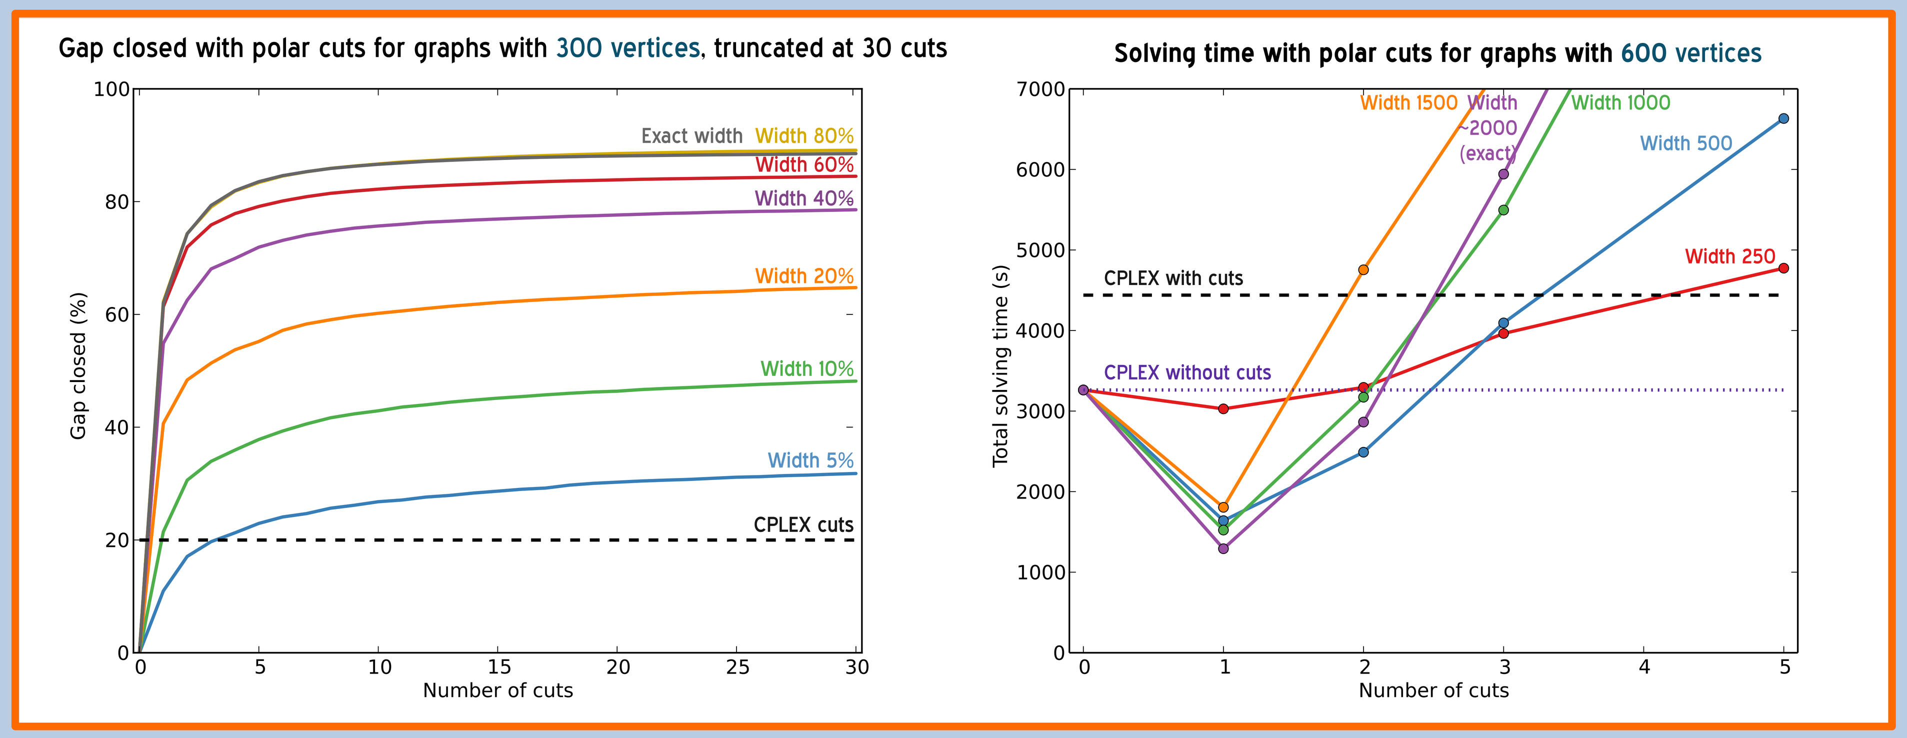

Despite a good gap closed even for weaker relaxations, in practice only the first cut is worth its cost to generate, but it greatly benefits the solver. On medium or small density graphs, polar cuts are not worth generating because they require large decision diagrams for a tight relaxation.

This however serves as a proof of concept: **if a problem has strong relaxed decision diagrams, then very useful cuts can be extracted from them.**

Conclusions and other insights

We develop a method to generate polar cuts directly from a decision diagram representation of the feasible set. Our computational results suggest that this framework can be helpful for solving problems with good dynamic programming formulations, either entirely or as a substructure.

Moreover, this can also be used to analyze facet structure with exact decision diagrams. It allows us to perform shooting experiments: shoot random rays and find out which facets are hit more often.

Insights from other experiments:

How important is "facet-defining"?

Letting $\omega = 0$ often yields cuts that are not facet-defining. If we make them facet-defining, we may double the gap closed, even though the violation is the same.

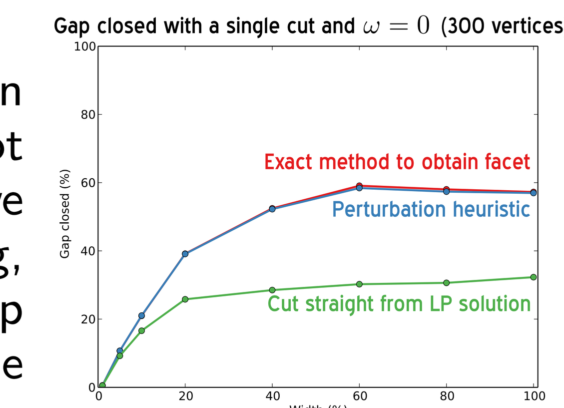

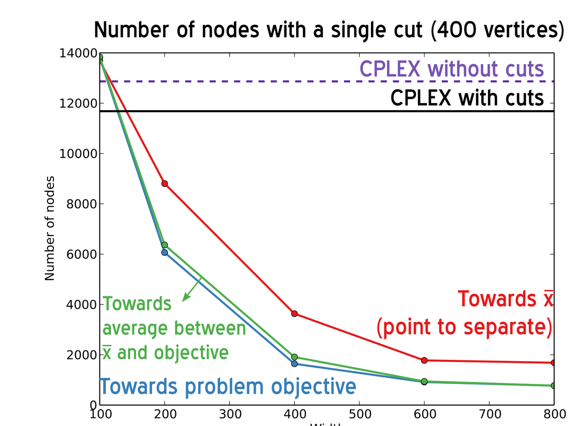

Should we always follow the LP optimum?

Experiments suggest that if consider the objective when aiming, we can obtain better cuts.

This is not the end: **there is a lot to explore!**

- What practical problems can we apply these cuts to?
- Can we identify and exploit substructures with tight decision diagram relaxations?
- What else can we say about decision diagrams with respect to polyhedra?

References

H. R. Andersen, T. Hadzic, J. N. Hooker, and P. Tiedemann. "A constraint store based on multivalued decision diagrams". In: Principles and Practice of Constraint Programming–CP 2007. Springer, 2007, pp. 118–132.

B. Becker, M. Behle, F. Eisenbrand, and R. Wimmer. "BDDs in a branch and cut framework". In: Experimental and Efficient Algorithms. Springer, 2005, pp. 452–463.

D. Bergman, A. A. Ciré, W.-J. van Hoeve, and J. N. Hooker. "Discrete optimization with decision diagrams". To appear.

C. Buchheim, F. Liers, and M. Oswald. "Local cuts revisited". In: Operations Research Letters 36.4 (2008), pp. 430–433.