

Branch-and-Cut is our swiss army knife

Matteo Fischetti, University of Padova



A simple idea (?)

- Mixed-Integer Programs (MIPs) can be solved by two alternative techniques:
 - **Cutting planes** (notably, Mixed-Integer Gomory cuts)
 - **Branch and Bound**
- Pros and cons are complementary, so ...
... why not **merging** them?
- This idea was around already in the 1980's
- BUT: how to actually **implement** it?

Why bothering about implementations?

- Implementation is **not** just coding!



```
selfsplit.c — src
x main.c x ic_separation.c x locbra.c x selfsplit.c
300 ~
301 > ssdata.phase = 1;
302 > CPXmipopt(env, lp);
303 > nodelim = nodelim - CPXgetnodecnt(env, lp); if ( nodelim < 0 ) nodelim = 0;
304 ~
305 > // 2nd phase: kill nodes belonging to the other workers ~
306 > CPXsetintparam(env, CPX_PARAM_THREADS, 1);
307 > CPXsetintparam(env, CPX_PARAM_NODELIM, INT_MAX);
308 > ssdata.phase = 2;
309 > CPXmipopt(env, lp);
310 ~
311 > // create the synchronization file (if does not exist already) ~
312 > ssdata.counter = 0;
313 > if ( strcmp(sync_file, "NULL") == 0 ) ssdata.counter = -1; // means no synchronization file
314 > if ( strcmp(sync_file, "NONE") == 0 ) ssdata.counter = -1; // means no synchronization file
315 > if ( (ssdata.counter >= 0) && (!SS_file_exists(sync_file)) ) // synchronization file does not exist
316 > {
317 >     double zstar; if ( CPXgetobjval(env, lp, &zstar) ) zstar = CPX_INFBOUND;
318 >     double *xbest = (double *) calloc(ssdata.ncols, sizeof(double));
319 >     if ( zstar < CPX_INFBOUND/2.0 ) CPXgetx(env, lp, xbest, 0, ssdata.ncols-1);
320 >     SS_write_sol(&ssdata, zstar, xbest);
321 >     free(xbest);
322 > }
323 ~
324 > // 3rd phase: complete the optimization with the original parameters ~
325 > if ( ssdata.verbose >= 1 ) printf("\n## SelfSplit: unit %d out of %d starts the final runt at epoch %ld\n\n", ssdata.unit, ssdata.num_units, time(NULL));
326 > CPXsetintparam(env, CPX_PARAM_NODELIM, nodelim);
327 > CPXsetintparam(env, CPX_PARAM_BBINTERVAL, bbininterval);
328 > CPXsetintparam(env, CPX_PARAM_NOSEL, nodelim);
```

- Needed if we **#orms** want to have an impact in practical applications
- ... but often omitted in papers as “of no interest for a typical reader”
- Ask yourself: would Artificial Intelligence (notably: deep learning) be so successful without gradient-descent **algorithms** served with their efficient **#backpropagation implementations**?

Algorithms as theorems

Theorem 2 Assume w.l.o.g. that $\text{rank}(A) = n$. Given a vertex x^* of P , let the system $Ax \geq b$ be partitioned into $Bx \geq b_B$ and $Nx \geq b_N$, where $Bx^* = b_B$ and B is an $n \times n$ nonsingular matrix. Let (u_B, v_B) and (u_N, v_N) denote the Farkas multipliers associated with the rows of B and N , respectively. For a given disjunction (2) with $\eta^* = \pi x^* - \pi_0 \in [0, 1]$, let $u_0^* = 1 - \eta^*$, $v_0^* = \eta^*$, $u_N^* = v_N^* = 0$, $u_B^* = [\pi B^{-1}]_+$ and $v_B^* = [-\pi B^{-1}]_+$, while γ^* and γ_0^* are defined through (4) and (6), respectively. Then $(\gamma^*, \gamma_0^*, u^*, v^*, u_0^*, v_0^*)$ is an optimal CGLP solution w.r.t. the trivial normalization (10).

Proof We first prove feasibility. Consistency between (4) and (5) requires $u^*A - u_0^*\pi = v^*A + v_0^*\pi$, i.e., $u_B^* - v_B^* = (u_0^* + v_0^*)\pi B^{-1} = \pi B^{-1}$, which follows directly from the definition of u_B^* and v_B^* . Analogously, consistency between (6) and (7) requires $(u_B^* - v_B^*)b_B = (u_0^* + v_0^*)\pi_0 + v_0^*$, i.e., $\pi B^{-1}b_B = \pi_0 + v_0^*$. This latter equation is indeed satisfied because $B^{-1}b_B = x^*$ and $v_0^* = \eta^* = \pi x^* - \pi_0$. As to optimality, we observe that $u_0^* + v_0^* = 1$ holds by definition. Because of (4) and (6), $\gamma x^* - \gamma_0 = u^*(Ax^* - b) - u_0^*(\pi x^* - \pi_0) = u_B^*(Bx^* - b_B) + u_N^*(Nx^* - b_N) - u_0^*\eta^* = 0 + 0 - (1 - \eta^*)\eta^*$, hence the cut violation attains bound UB3 of Lemma 1. \square

Algorithms without implementation

Theorem 2 Assume w.l.o.g. that $\text{rank}(A) = n$. Given a vertex x^* of P , let the system $Ax \geq b$ be partitioned into $Bx \geq b_B$ and $Nx \geq b_N$, where $Bx^* = b_B$ and B is an $n \times n$ nonsingular matrix. Let (u_B, v_B) and (u_N, v_N) denote the Farkas multipliers associated with the rows of B and N , respectively. For a given disjunction (2) with $\eta^* = \pi x^* - \pi_0 \in [0, 1]$, let $u_0^* = 1 - \eta^*$, $v_0^* = \eta^*$, $u_N^* = v_N^* = 0$, $u_B^* = [\pi B^{-1}]_+$ and $v_B^* = [-\pi B^{-1}]_+$, while γ^* and γ_0^* are defined through (4) and (6), respectively. Then $(\gamma^*, \gamma_0^*, u^*, v^*, u_0^*, v_0^*)$ is an optimal CGLP solution w.r.t. the trivial normalization (10).

Proof: omitted as of no interest to the typical MP reader.

Describing an **Algorithm** without **Implementation** is like stating a **Theorem** without **Proof**

#just_a_computational_conjecture

Branch & Cut™

A BRANCH-AND-CUT ALGORITHM
FOR THE RESOLUTION OF LARGE-SCALE
SYMMETRIC TRAVELING SALESMAN PROBLEMS *

MANFRED PADBERG† AND GIOVANNI RINALDI‡

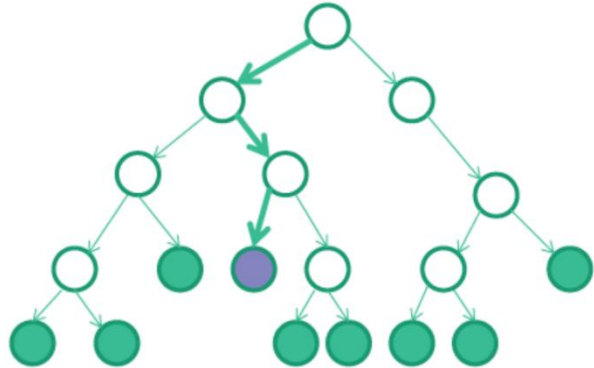
- A “trademark” of Manfred Padberg and Giovanni Rinaldi
- Proposed in the 1990’s for the TSP (and soon extended)
- Comes as an **algorithm** entangled with its **implementation**


Theorem. *Using cuts within an enumerative scheme is good.*

Proof. Assume w.l.o.g. a good LP solver. Then apply B&Bound but

- make use of families of (problem dependent) globally-valid inequalities
- perform efficient exact/heuristic cut separation on the fly
- use a data-structure (cut pool) to effectively share cuts among nodes
- price variables in a dynamic way (well before branch-and-price!)
- alternate row and column generation in a sound way ...
- suspend a node if “unattractive”
- ...

Modern B&C implementation


- Modern B&C solvers such as Cplex, Gurobi, Express, SCIP etc. can be fully **customized** by using **callback functions**
 - Callback functions are just **entry points** in the B&C code where an advanced user (you!) can add his/her customizations
 - Most-used callbacks (using old-style Cplex's jargon)
 - **Lazy constraint**: add “lazy constr.s” that should be part of the original model
 - **User cut**: add additional contr.s that hopefully help enforcing feasibility/integrality
 - Heuristic: try to improve the incumbent (primal solution) as soon as possible
 - Branch: modify the branching strategy
 - ...
- 



Lazy constraint callback

CPX_CALLBACKCONTEXT_CANDIDATE

- Automatically invoked when a solution is going to update the **incumbent** (meaning it is **integer** and **feasible** w.r.t. current model, e.g., because it comes from an internal primal heuristic)
- This is the **last checkpoint** where we can discard a solution for whatever reason (e.g., because it violates a constraint that is not part of the current model)
- To avoid be bothered by this solution again and again, we can/should return a **violated constraint (cut)** that is added (globally or locally) to the current model
- Cut generation is often **simplified** by the fact that the solution to be cut is known to be **integer** (e.g., SECs for TSP)



Usercut callback

CPX_CALLBACKCONTEXT_RELAXATION

- Automatically invoked at every B&B node when the current solution is **noninteger** (e.g., just before branching)
- A **violated cut** can possibly be returned, to be added (locally or globally) to the current model → often leads to an improved convergence to integer solutions
- If no cut is returned, **branching** occurs as usual
- Cut generation **can be hard** as the point is noninteger (heuristic approaches can be used)
- User cuts are **not mandatory** for B&C correctness → being too clever on them can actually **slow-down** the solver because of the overhead in generating and using them (larger/denser LPs etc.)



Other callbacks

- **Branch callback:** invoked at the end of each node (even when the LP solution is integer and apparently does not require any cut/branching) and used to impose/customize branching
- **Heuristic callback:** used to build new (possibly problem-specific) feasible integer solutions to be **posted**, i.e., passed to the solver which will use them (at the appropriate time) to possibly update the incumbent
- etc. etc.

Application: non-convex MIQP

(based on joint work with Michele Monaci, Univ. Bologna)

- **Goal: implement a Mixed-Integer (non-convex) Quadratic solver**
- Two approaches:
 1. start with a continuous QP solver and add enumeration on top of it
→ implement B&B to handle integer var.s
 2. start with a MILP solvers (B&C) and customize it to handle the non-convex quadratic terms → add **McCormick & spatial branching**
 - PROS: ...
 - CONS: ...

MIQP as a MILP with bilinear eq.s

- The fully-general MIQP of interest reads

$$\begin{aligned} (MIQP) \quad & \min a_0^T x + x^T Q^0 x \\ & a_k^T x + x^T Q^k x @ b, \quad k = 1, \dots, m \\ & \ell_j \leq x_j \leq u_j, \quad j = 1, \dots, n \\ & x_j \text{ integer}, \quad j \in \mathcal{I}, \\ & x_j \text{ continuous}, \quad j \in \mathcal{C}, \end{aligned}$$

and can be restated as

$$\begin{aligned} (MIBLP) \quad & \min_x c^T x \\ & Ax = b \\ & \ell_j \leq x_j \leq u_j, \quad j = 1, \dots, n \\ & x_j \text{ integer}, \quad j \in \mathcal{I} \\ & x_j \text{ continuous}, \quad j \in \mathcal{C} \\ & x_{r_k} = x_{p_k} x_{q_k}, \quad k = 1, \dots, K, \end{aligned}$$

McCormick inequalities

- To simplify notation, rewrite the generic bilinear eq. $x_{r_k} = x_{p_k} x_{q_k}$ as:

$$z = x y$$

$$\ell_x \leq x \leq u_x$$

$$\ell_y \leq y \leq u_y$$

- Obviously

$(x - \ell_x)(y - \ell_y) \geq 0$	\rightarrow	mc1) $z \geq \ell_y x + \ell_x y - \ell_x \ell_y$
$(x - u_x)(y - u_y) \geq 0$	\rightarrow	mc2) $z \geq u_y x + u_x y - u_x u_y$
$(x - \ell_x)(y - u_y) \leq 0$		mc3) $z \leq u_y x + \ell_x y - \ell_x u_y$
$(x - u_x)(y - \ell_y) \leq 0$		mc4) $z \leq \ell_y x + u_x y - u_x \ell_y$

(just replace xy by z in the products on the left)

- Note: $mc1)$ and $mc2)$ can be improved in case $x=y \rightarrow$ **gradients cuts**

$$z \geq x_0^2 + 2x_0(x - x_0), \quad \text{for each } x_0 \in \mathbb{R}$$

Spatial branching

- McCormick inequalities are not perfect
→ they are **tight** only when x and/or y are at their lower/upper bound
- $$\begin{aligned}(x - \ell_x)(y - \ell_y) &\geq 0 \\(x - u_x)(y - u_y) &\geq 0 \\(x - \ell_x)(y - u_y) &\leq 0 \\(x - u_x)(y - \ell_y) &\leq 0\end{aligned}$$

→ at some B&C nodes, it may happen that the current (fractional or integer) solution satisfies **all** MC inequalities but some bilinear eq.s $z = xy$ are still violated (we call this **#bilinear_infeasibility**)

→ we need a **bilinear-specific branching** (the usual MILP branching on integrality does not work if all var.s are integer already)

- Spatial branching**: if $z^* = x^* y^*$ is a violated bilinear eq., branch on **$(x \leq x^*)$ OR $(x \geq x^*)$**
to make the upper (resp. lower) bound on x tight on the left (resp. right) child node – thus improving the corresponding MC inequality

Vanilla B&C implementation


- **Lazy constraint callback:** separation of MC inequalities
- **Usercut callback:** not needed (and sometimes detrimental)
- **Branch callback:** spatial branching on the “most violated” $z = xy$
- **Precision:** LP precision higher (more restrictive) than bilinear tolerance
- **MILP heuristics** (kindly provided by the MILP solver): active at their default level
- **MIQP-specific heuristics:** not implemented
- Implemented but not used in the vanilla version:
 - additional bilinear-specific cuts → Balas’ **Intersection Cuts (ICs)**
 - **semi-spatial** branching (branch threshold $x^* + \delta \rightarrow x^*$ violates the x -bound in one of the two children, MC only needed in the other one)

Does it work?

- Comparison with the SCIP 5.0 MIQP solver using CPLEX 12.8 as **LP solver** + internal nonlinear solver
- Preliminary test on the quadratic **MINLPlib** (700+ instances) ...
... but some instances removed as root LP was **unbounded**
→ they need bound tightening by preprocessing (TODO)
- Results of our B&C callback-based vanilla implementation using CPLEX 12.8 as **MILP solver**; 1-thread runs (parallel runs not allowed in SCIP); only instances solved by both codes in the 1-hour time limit.
 - Overall, we are as fast as SCIP (but the latter solves more instances within the time limit → SCIP qualifies as a more robust solver).
 - We are 2 to 10 times faster than SCIP when the optimal/best-known solution from MINLPlib is used as a warm-start for both codes → evidently, we miss a sound bilinear-specific heuristic (TODO)

More detailed comparison

- **SCIP** vs **noic** (our “vanilla” version with no ICs and classical spatial branching)



- Results with incumbent warm-start (only instances solved by both codes)

