AUTOMATED REASONING

Agostino Dovier

Università di Udine
CLPLAB

Udine, November 2016

GLOBAL CONSTRAINTS

- Constraints on sets of variables (with some precise meaning) are called global constraints
- Often global constraints can be rewritten as combination of binary constraints. However, propagation on these binary constraints is rather poor wrt the constraint viewed as a whole.
- Therefore, global constraints are often studied independently.
- The most famous is all_different.

- Let $X_1, \ldots, X_k$ be variables with domains $\mathcal{D}_1, \ldots, \mathcal{D}_k$.
- The (k-ary) constraint all_different($X_1, \ldots, X_k$) is defined as follows:

$$\text{all\_different}(X_1, \ldots, X_k) = (\mathcal{D}_1 \times \cdots \times \mathcal{D}_k) \setminus \{(a_1, \ldots, a_k) \in \mathcal{D}_1 \times \cdots \times \mathcal{D}_k \,:\, \exists i \exists j\, 1 \leq i < j \leq k\, (a_i = a_j)\}$$

- A CSP is said *diff-arc consistent* iff every all_different-constraint in it is *hyper arc consistent*.
- Namely, for every $i \in \{1, \ldots, k\}$ and every $a_i \in \mathcal{D}_i$ there are $a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_k$ s.t. $\langle a_1, \ldots, a_k \rangle \in \text{all\_different}(X_1, \ldots, X_k)$

- Let $X_1, \ldots, X_k$ be variables with domains $\mathcal{D}_1, \ldots, \mathcal{D}_k$.
- The (k-ary) constraint all_different($X_1, \ldots, X_k$) is defined as follows:

$$\text{all\_different}(X_1, \ldots, X_k) = (\mathcal{D}_1 \times \cdots \times \mathcal{D}_k) \setminus \{(a_1, \ldots, a_k) \in \mathcal{D}_1 \times \cdots \times \mathcal{D}_k \ : \ \exists i \exists j \, 1 \leq i < j \leq k \, (a_i = a_j)\}$$

- A CSP is said *diff-arc consistent* iff every all_different-constraint in it is *hyper arc consistent*.
- Namely, for every $i \in \{1, \ldots, k\}$ and every $a_i \in \mathcal{D}_i$ there are $a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_k$ s.t. $\langle a_1, \ldots, a_k \rangle \in$ all_different($X_1, \ldots, X_k$)

All Different Constraint

- Let us observe that the two CSP $\langle \text{all\_different}(X_1, \ldots, X_k); \mathcal{D}_\in \rangle$ and $\langle X_1 \neq X_2, X_1 \neq X_3, \ldots, X_1 \neq X_k, X_2 \neq X_3, \ldots, X_{k-1} \neq X_k; \mathcal{D}_\in \rangle$ are equivalent.
- hyper-arc-consistency of all_different$(X_1, \ldots, X_k)$ implies (binary) arc consistency in the second CSP.
- The converse does not hold:
  $\langle \text{all\_different}(X_1, X_2, X_3); \mathcal{D}_1 = \mathcal{D}_2 = \mathcal{D}_3 = \{0, 1\} \rangle$ vs
  $\langle X_1 \neq X_2, X_1 \neq X_3, X_2 \neq X_3; \mathcal{D}_1 = \mathcal{D}_2 = \mathcal{D}_3 = \{0, 1\} \rangle$.
- Let $d_i = |\mathcal{D}_i|$ for $i \in \{1, \ldots, k\}$ and $d = \max_{i=1}^{k}\{d_i\}$.
- A propagation algorithm for hyper-arc-consistency based on the definition has cost $O(\ldots d^{k+1})$.
- Not applicable with large k. But is it an intrinsic problem or just a too naive algorithm?

All Different Constraint

- Let us observe that the two CSP $\langle \text{all\_different}(X_1, \ldots, X_k); \mathcal{D}_\in \rangle$ and $\langle X_1 \neq X_2, X_1 \neq X_3, \ldots, X_1 \neq X_k, X_2 \neq X_3, \ldots, X_{k-1} \neq X_k; \mathcal{D}_\in \rangle$ are equivalent.
- hyper-arc-consistency of $\text{all\_different}(X_1, \ldots, X_k)$ implies (binary) arc consistency in the second CSP.
- The converse does not hold:
  $\langle \text{all\_different}(X_1, X_2, X_3); \mathcal{D}_1 = \mathcal{D}_2 = \mathcal{D}_3 = \{0, 1\} \rangle$ vs
  $\langle X_1 \neq X_2, X_1 \neq X_3, X_2 \neq X_3; \mathcal{D}_1 = \mathcal{D}_2 = \mathcal{D}_3 = \{0, 1\} \rangle$.
- Let $d_i = |\mathcal{D}_i|$ for $i \in \{1, \ldots, k\}$ and $d = \max_{i=1}^{k}\{d_i\}$.
- A propagation algorithm for hyper-arc-consistency based on the definition has cost $O(\ldots d^{k+1})$.
- Not applicable with large k. But is it an intrinsic problem or just a too naive algorithm?

Propagation of all_different

Bipartite Graphs

- An *bipartite graph* is a triple $G = \langle X, Y, E \rangle$ where X and Y are disjoint sets of nodes and $E \subseteq X \times Y$ is a set of edges
- Edges are treated as not directed
- A *matching* $M \subseteq E$ is a set of edges such that there are no pairs of edges that share the same node.
- Given G and M a node is said *matched* if it is in some edge in M; otherwise it is *free*.
- A *path* is a sequence of edges $(x_1, y_1), (y_1, x_2), (x_2, y_2), \ldots$

- Given $G = \langle X, Y, E \rangle$ bipartite graph and $M \subseteq E$ matching, a path in G is *alternating* for M if edges of the path are alternatively in M and not in M.
- An alternating path is *augmenting* for M if it is *acyclic* and starts and ends in free nodes.
- Observe that every augmenting path starting in a node in X ends in Y (or vice versa).
- If $M = \emptyset$, any (set containing a single) edge is an augmenting path for M.

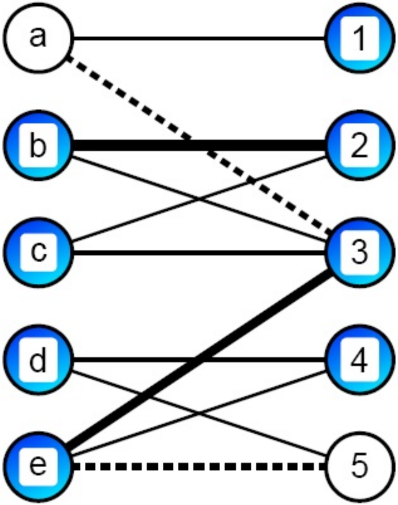
a
1
b
2
c
3
d
4
e
5

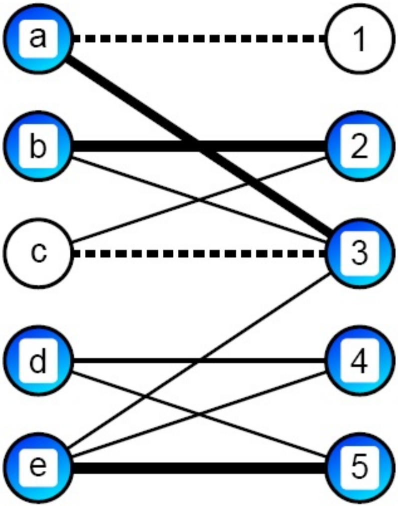
a
1
b
2
c
3
d
4
e
5

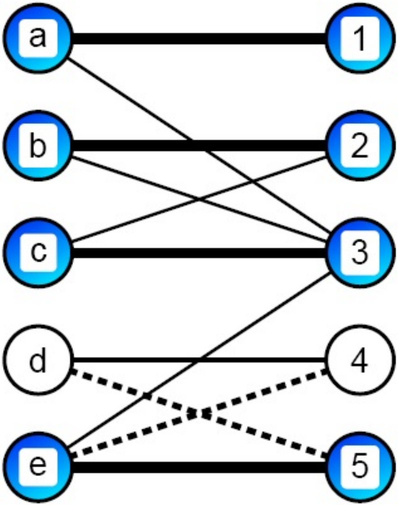
a
1
b
2
c
3
d
4
e
5

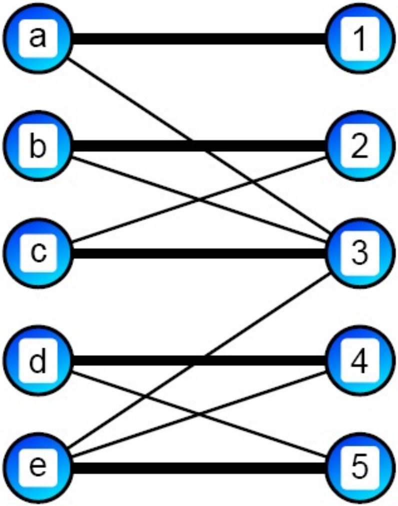
a
1
b
2
c
3
d
4
e
5

PROP. Given a matching M and an augmenting path P for M, we have that: $M' = M \oplus P = (M \setminus P) \cup (P \setminus M)$ is a matching such that $|M'| = |M| + 1$.

Proof.

Let $M = \{(x_1, y_1), \ldots, (x_k, y_k),$ The part of M in the path
$(x_{k+1}, y_{k+1}), \ldots, (x_n, y_n)\}$

Suppose P is the augmenting path: $P =$
$\{(y_0, x_1), \underline{(x_1, y_1)}, (y_1, x_2), \underline{(x_2, y_2)}, \ldots, (y_{k-1}, x_k), \underline{(x_k, y_k)}, (y_k, x_0)\}$
(wlog, $y_0 \in Y$ and $x_0 \in X$ and the path uses the 'first' k edges of the set M).

Let $M' = (M \setminus P) \cup (P \setminus M)$
$= \{(y_0, x_1), (y_1, x_2), \ldots, (y_{k-1}, x_k), (y_k, x_0),$
$(x_{k+1}, y_{k+1}), \ldots, (x_n, y_n)\}$

M' is a matching: $x_1, \ldots, x_n, y_1, \ldots, y_n$ are all different (M is a matching and x_0 and y_0 are different since free by hypothesis).
Its length is $|M| + 1$.

PROP. Given a matching M and an augmenting path P for M, we have that: $M' = M \oplus P = (M \setminus P) \cup (P \setminus M)$ is a matching such that $|M'| = |M| + 1$.

Proof.

Let $M = \{(x_1, y_1), \dots, (x_k, y_k), (x_{k+1}, y_{k+1}), \dots, (x_n, y_n)\}$ — The part of M in the path: $(x_1, y_1), \dots, (x_k, y_k)$

Suppose P is the augmenting path: $P = \{(y_0, x_1), \underline{(x_1, y_1)}, (y_1, x_2), \underline{(x_2, y_2)}, \dots, (y_{k-1}, x_k), \underline{(x_k, y_k)}, (y_k, x_0)\}$
(wlog, $y_0 \in Y$ and $x_0 \in X$ and the path uses the 'first' k edges of the set M).

$$\begin{aligned} \text{Let } M' &= (M \setminus P) \cup (P \setminus M) \\ &= \{(y_0, x_1), (y_1, x_2), \dots, (y_{k-1}, x_k), (y_k, x_0), \\ &\quad (x_{k+1}, y_{k+1}), \dots, (x_n, y_n)\} \end{aligned}$$

M' is a matching: $x_1, \dots, x_n, y_1, \dots, y_n$ are all different (M is a matching and x_0 and y_0 are different since free by hypothesis).
Its length is $|M| + 1$.

- We are interested in matching of maximum size *maximum bipartite matching*.
- Of course the size of the maximum matching is $\leq \min\{|X|, |Y|\}$.
- **Theorem** (Berge–1957). Given a bipartitite graph $G = \langle X, Y, E \rangle$, ($M$ is a maximum matching) if and only if (there are not augmenting paths for M).
- ($\rightarrow$) To show $A \rightarrow B$ we prove $\neg B \rightarrow \neg A$.
  Suppose there is an augmenting path for M.
  Then (previous Prop) we can define a matching M' of size $|M| + 1$, thus M is not of maximum size.

- ($\leftarrow$) To show $A \leftarrow B$ we prove that $\neg A \rightarrow \neg B$. Namely, if M is a matching not of maximum size, then there is an augmenting path for M. Let M be a not maximal matching. Then there is a matching M' of bigger size ($|M'| > |M|$). Now let $U = \langle X, Y, M \oplus M' \rangle$.
  1. Being M and M' matchings, at most one edge of M and one of M' might incide on a node in U. Thus the degree of a node in U is ≤ 2.
  2. The graph U can have cycles. In this case half of the edges are from M and the other from M' (and the length of the cycle is even).
- Now, removing the cycles, the paths use alternatively edges from M and from M'. Since $|M'| > |M|$ there must be at least one path with more edges from M' than from M.
- That path starts and ends with edges in M' starting and ending with two nodes free for M: There is an augmenting path for M. □

MAXIMUM MATCHING

NAIVE ALGO

Let $n = \min\{|X|, |Y|\}$ and $m = |E|$.

```
Max_Matching_Naive(⟨X, Y, E⟩)
  M ← ∅;
  while (there is an augmenting path P for M)
  do
      M ← M ⊕ P;
  return M;
```

The algorithm terminates in $\leq n$ iterations.
Let us analyze briefly the cost of each iteration.

MAXIMUM MATCHING

NAIVE ALGO

Find_Augmenting_Path($\langle X, Y, E\rangle, M$)

$S \leftarrow X$; $A \leftarrow E$;
trovato $\leftarrow$ false;
while (S contains a free node $\wedge \neg$*trovato*)
do
    choose a free node x in S;
    depth first search of an augmenting path for M in $\langle S, Y, A\rangle$;
    let $E(x)$ be the set of edges visited starting from x;
    if (a path is found)
      then
        trovato $\leftarrow$ `true`
      else $S \leftarrow S \setminus \{x\}$; $A \leftarrow A \setminus E(x)$;
return *trovato*/*path*

$O(|E|)$: globally we have $O(nm)$.
Hopcroft-Karp: $O(m\sqrt{n}) = O(n^2\sqrt{n})$.

Applying these results to the propagation of the all_different constraint

Given all_different($X_1, \ldots, X_k$), with domains $\mathcal{D}_1, \ldots, \mathcal{D}_k$, let us define the bipartite graph $GV(C) = \langle X_C, Y_C, E_C \rangle$ as follows:

- $X_C = \{X_1, \ldots, X_k\}$
- $Y_C = \bigcup_{i=1}^{k} \mathcal{D}_i$
- $E_C = \{(X_i, a) \ : \ a \in \mathcal{D}_i\}$

All Different Constraint

Eample: all_different($X_1, \ldots, X_7$)

$X_1 \in 1..2, X_2 \in 2..3, X_3 \in \{1,3\}, X_4 \in \{2,4\},$
$X_5 \in 3..6, X_6 \in 6..7, X_7 \in \{8\}$

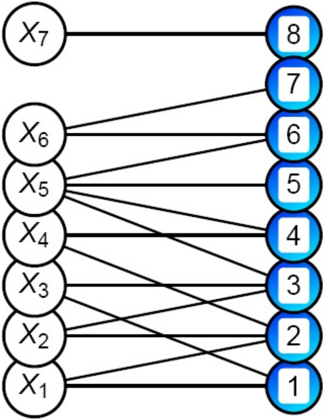

Theorem: A CSP $\mathcal{P} = (\mathcal{C}; \mathcal{D}_{\in})$ is diff-arc consistent if and only if forall all_different-constraint C in $\mathcal{C}$ any edge in $GV(C)$ belongs to (at least) one matching of the same size of the set of variables of C.

Proof:

$(\rightarrow)$ Let C in $\mathcal{C}$ a all_different-constraint. Let $X_1, \ldots, X_k$ be its variables. Let (X_i, a_i) in $GV(C)$ (this implies that $a_i \in \mathcal{D}_i$). Since $\mathcal{P}$ is diff-arc consistent, C is hyper-arc consistent. Thus, there are $a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_k$ such that $X_1 = a_1, \ldots, X_k = a_k$ is a solution of C. This solution is a matching of the desired size to which the edge (X_i, a_i) belongs.

$(\leftarrow)$ For any all_different-constraint C (assume it has k vars) consider an edge (X_i, a_i) belonging to a matching of size k. From that matching we obtain the values for the other variables to guarantee the hyper-arc consistency property. □

Theorem: A CSP $\mathcal{P} = (\mathcal{C}; \mathcal{D}_{\in})$ is diff-arc consistent if and only if forall all_different-constraint C in $\mathcal{C}$ any edge in $GV(C)$ belongs to (at least) one matching of the same size of the set of variables of C.

Proof:

$(\rightarrow)$ Let C in $\mathcal{C}$ a all_different-constraint. Let $X_1, \dots, X_k$ be its variables. Let (X_i, a_i) in $GV(C)$ (this implies that $a_i \in \mathcal{D}_i$). Since $\mathcal{P}$ is diff-arc consistent, C is hyper-arc consistent. Thus, there are $a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_k$ such that $X_1 = a_1, \dots, X_k = a_k$ is a solution of C. This solution is a matching of the desired size to which the edge (X_i, a_i) belongs.

$(\leftarrow)$ For any all_different-constraint C (assume it has k vars) consider an edge (X_i, a_i) belonging to a matching of size k. From that matching we obtain the values for the other variables to guarantee the hyper-arc consistency property. □

- For building $GV(C)$, we introduce k nodes (for the variables), $|\mathcal{D}_1 \cup \cdots \cup \mathcal{D}_k|$ for the domain objects, and $e = d_1 + \cdots + d_k \leq kd$ edges.
- A constraint $C = \text{all\_different}(X_1, \ldots, X_k)$ is hyper-arc consistent if and only if any edge in $GV(C)$ belongs to a matching of size k (maximum).
- We know that we can find ONE maximum matching in time $O(\sqrt{k}e) = O(k^{3/2}d)$.
- We are interested in edges that belong to at least one maximum matching.
- We need a new idea here!

- For building $GV(C)$, we introduce k nodes (for the variables), $|\mathcal{D}_1 \cup \cdots \cup \mathcal{D}_k|$ for the domain objects, and $e = d_1 + \cdots + d_k \leq kd$ edges.
- A constraint $C = \text{all\_different}(X_1, \ldots, X_k)$ is hyper-arc consistent if and only if any edge in $GV(C)$ belongs to a matching of size k (maximum).
- We know that we can find ONE maximum matching in time $O(\sqrt{k}e) = O(k^{3/2}d)$.
- We are interested in edges that belong to at least one maximum matching.
- We need a new idea here!

Teorema: [Berge–1970] Let $G = \langle X, Y, E \rangle$ be a bipartire graph.
An edge e belongs to some but not to all maximum matchings if and only if for an arbitrary maximum matching M, e belongs to:

- an acyclic even length alternating path that starts in a free vertex OR
- an alternating cycle (of course of even length).

Proof. ($\Leftarrow$) Let M be a maximum matching.

1. Let P an alternating even length path with a free extreme (the other is not free). Let $M' = M \oplus P$ another matching with the same size. Half of the edges of the path P are in M, the other half in M'.
2. Let P be an alternating cycle. Similar.

All Different Constraint

An edge e belongs to some but not to all maximum matchings $\Rightarrow$ for an arbitrary maximum matching M, e belongs to: an acyclic even length alternating path that starts in a free vertex OR an alternating cycle (of course of even length).

Let (x_0, x_1) be an edge that belongs to some but not to all maximum matchings.
Let M a generic max matching s.t. $(x_0, x_1) \in M$ and M' a generic max matching s.t. $(x_0, x_1) \notin M$.
By hypothesis, at least one of such M and M' exists.
Now, let $M'' = M \oplus M'$. We have that $(x_0, x_1) \in M''$.
We already know that the degree of each node in M'' is ≤ 2.
Iteratively choose nodes $x_2, x_3, x_4, x_5, \dots, x_m$ such that $(x_i, x_{i+1}) \in M''$ and $x_{i+2} \neq x_i$.
Stop when there are no longer nodes to be chosen or when x_m has been previously chosen.

An edge e belongs to some but not to all maximum matchings $\Rightarrow$ for an arbitrary maximum matching M, e belongs to: an acyclic even length alternating path that starts in a free vertex OR an alternating cycle (of course of even length).

Let (x_0, x_1) be an edge that belongs to some but not to all maximum matchings.
Let M a generic max matching s.t. $(x_0, x_1) \in M$ and M' a generic max matching s.t. $(x_0, x_1) \notin M$.
By hypothesis, at least one of such M and M' exists.
Now, let $M'' = M \oplus M'$. We have that $(x_0, x_1) \in M''$.
We already know that the degree of each node in M'' is ≤ 2.
Iteratively choose nodes $x_2, x_3, x_4, x_5, \ldots, x_m$ such that $(x_i, x_{i+1}) \in M''$ and $x_{i+2} \neq x_i$.
Stop when there are no longer nodes to be chosen or when x_m has been previously chosen.

All Different Constraint

An edge e belongs to some but not to all maximum matchings $\Rightarrow$ for an arbitrary maximum matching M, e belongs to: an acyclic even length alternating path that starts in a free vertex OR an alternating cycle (of course of even length).

Two cases:

- A cycle has been pointed out (i.e., $x_m = x_1$)
- The sequence is not a cycle. In this case, go back to x_0 and choose a new sequence (that, in a sense, lead to x_0) introducing nodes $x_{-1}, x_{-2}, \ldots, x_{-t}$ such that $(x_{i-1}, x_i) \in M''$. Also in this case:
  - A cycle has been generated (but this is not possible ... why?)
  - No cycle has been generated.

An edge e belongs to some but not to all maximum matchings $\Rightarrow$ for an arbitrary maximum matching M, e belongs to: an acyclic even length alternating path that starts in a free vertex OR an alternating cycle (of course of even length).

If there are not cycles, then we have an alternating path both for M and for M'. If the path would be of odd length then it would be augmenting for one of them that would contradict the maximality of M or M'. Then the path is even.

In the case there is a cycle, there are two cases. Either (x_0, x_1) is in the path or not. Assume it is in the path, then there is a node of degree 3 in M'': this is absurdum.
Then (x_0, x_1) belongs to an alternating cycle both for M and for M'.

An edge e belongs to some but not to all maximum matchings $\Rightarrow$ for an arbitrary maximum matching M, e belongs to: an acyclic even length alternating path that starts in a free vertex OR an alternating cycle (of course of even length).

If there are not cycles, then we have an alternating path both for M and for M'. If the path would be of odd length then it would be augmenting for one of them that would contradict the maximality of M or M'. Then the path is even.

In the case there is a cycle, there are two cases. Either (x_0, x_1) is in the path or not. Assume it is in the path, then there is a node of degree 3 in M'': this is absurdum.
Then (x_0, x_1) belongs to an alternating cycle both for M and for M'.

- Find a maximal matching M (time $O(k^{3/2}d)$). If the size is less than k stop with unsat.
- For every free node, find the even alternating paths. All nodes reached and edges vidited are retained.
- All cycles are detected (algorithm for strongly connected components)
- All edges in M outside these visits should be in all matchings.
- All edges outside M outside these visits are removed.
- Global cost: $O(k^{3/2}d)$
- Data structure is retained to speed up further visits after some labelings.

All Different Constraint

Efficient Filtering

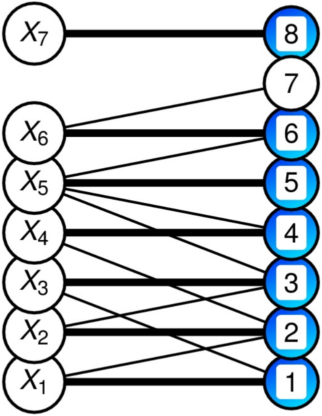

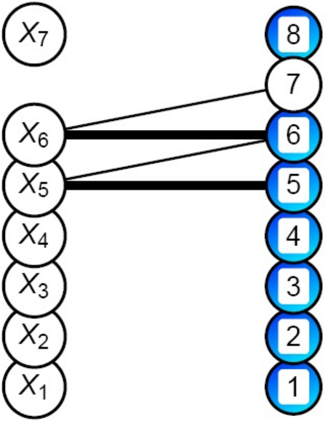

X_7
X_6
X_5
X_4
X_3
X_2
X_1
8
7
6
5
4
3
2
1

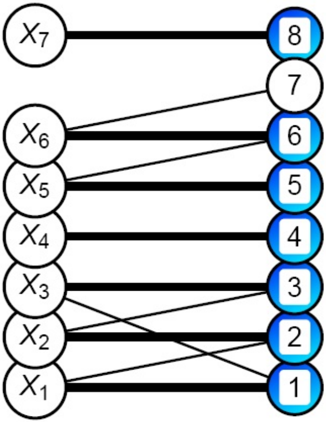

X_7
X_6
X_5
X_4
X_3
X_2
X_1
8
7
6
5
4
3
2
1