The Solution of Shannon Game

Frederic Maire, 2004

Recall that in Shannon game, the player SHORT tries to paint a path between the endpoints of an (unplayable) edge $e = xy$, securing edges by painting them, whereas the player CUT tries to disconnect x and y by deleting edges that have not been painted.

Shannon game is a solved game. The elegant solution that Lehman proposed [ref] forty years ago revolves around the notion of spanning tree that we will explain shortly. Although Lehman presented his solution in the more abstract framework of matroid theory (Shannon game can be extended to matroids), we will illustrate the essential ideas of the solution on graphs to avoid the heavy formalism of matroid theory. We refer the reader to the references for a complete mathematical presentation.

An introductory example

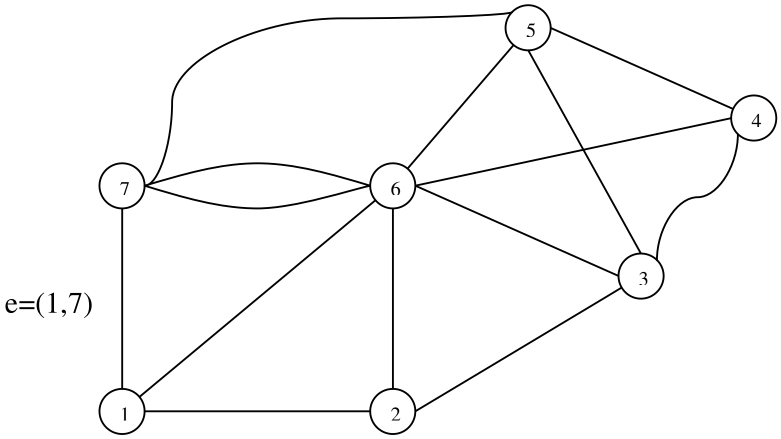

Figure 1 An introductory example

In the graph of Figure 1, SHORT's objective is to connect the endpoints of the edge $e = (1,7)$. That is, SHORT tries to paint a path from the vertex labelled 1 to the vertex labelled 7. CUT tries to prevent this connection from happening by deleting unpainted edges.

Although the graph of Figure 1 is small, a principled strategy is not obvious. On the contrary, the graph of Figure 2 is a trivial win for SHORT, even if SHORT plays second. If CUT deletes the dashed edge, SHORT will play the dotted edge. Reciprocally, if CUT deletes the dotted edge, SHORT will play the dashed edge.

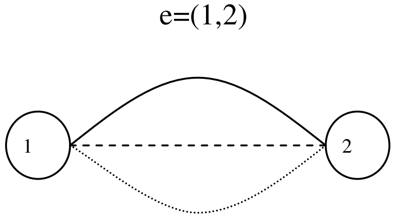

Figure 2 An easy win for SHORT

In the graph of Figure 2, SHORT has two ways to connect the endpoints of *e*. We will see that more generally the winning games for SHORT are characterized by the existence of two structured sets of edges that form an intertwined network connecting the endpoints of the distinguished edge *e*.

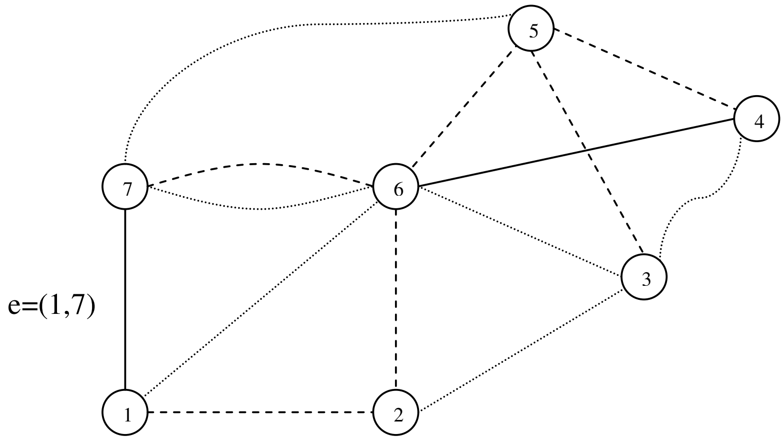

Figure 3 Two co-spanning trees spanning the edge e

In Figure 3, the dotted edges form a tree. The dashed edges form a second tree. Both trees are said *to span the edge e with respect to the circuit).* That is, the addition of the edge *e* to any of the trees creates a circuit. Moreover, each dotted edge is spanned by the dashed tree and each dashed edge is spanned by the dotted tree. Because of this, they are called co-spanning trees. A beautiful theorem of Lehman tells us that a game can be won by SHORT playing second if and only if there exist two co-spanning trees spanning the edge *e*. We can see that this is indeed the case for the graph of Figure 2.

The winning strategy for SHORT consists in restoring this co-spanning trees structure after each move by CUT. In fact, the edges outside the co-spanning trees, like edge (4,6) of Figure 3, can be ignored by SHORT. Suppose that CUT plays edge (5,6). The resulting dashed subgraph is displayed in Figure 4.

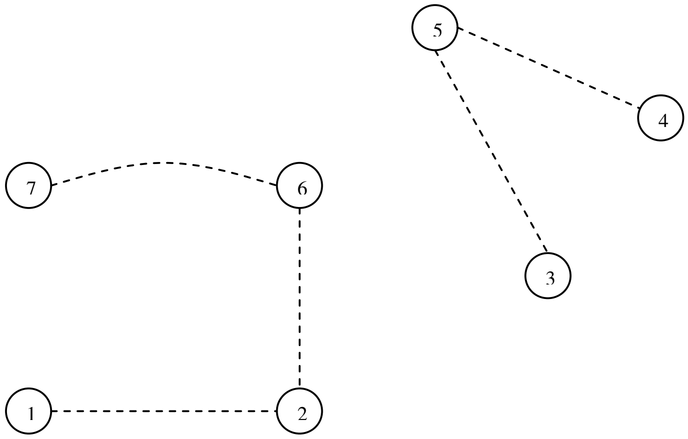

Figure 4 The two components {1,2,6,7} and {3,4,5} of the dashed subgraph

As edge $(5,6)$ is spanned by the dotted tree, there must be a dotted path going from vertex 5 to vertex 6. Along this path, there must be a dotted edge that bridges the two dashed components of Figure 4. This bridging edge is the edge that SHORT must play. Here this edge is edge $(5,7)$. Instead of painting the edge, SHORT can equivalently contract the edge. The resulting graph (Figure 5) contains again two co-spanning trees spanning the edge $e=(1,7)$.

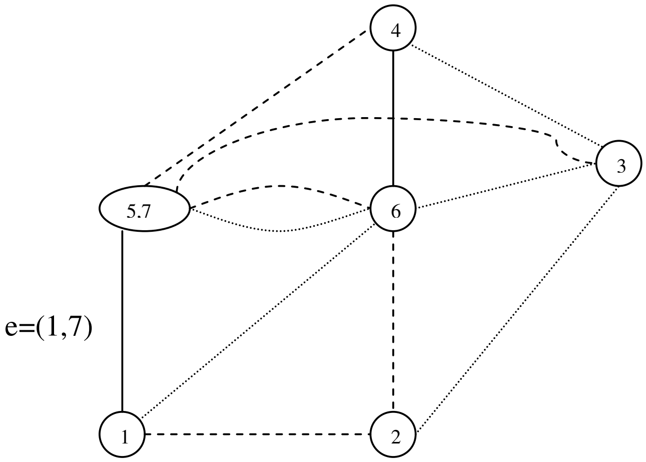

Figure 5 After SHORT plays (5,7) the two co-spanning trees structure is restored

A number of questions need to be answered to build a complete artificial player; what happens when there are no co-spanning trees? When two co-spanning trees exist, how can they be determined? Is there a similar strategy for CUT?

These questions will be answered in the following sections.

The duality of SHORT and CUT games

Surprisingly the SHORT and CUT games can be unified. Both players aim at securing a specific structured edge set with a distinguished edge e; SHORT tries to secure an edge set that with e make up a circuit, whereas CUT tries to secure an edge set that with e make up a cutset (definition below). The strategies for SHORT and CUT are dual strategies because of the duality between cutsets and circuits in graphs. Cutsets and circuits satisfy a number of common properties that have been used in matroid theory to define a set of axioms for a common abstraction. The players' strategies are identical when expressed in this common abstraction framework.

After defining cutsets and circuits, we will show how a pair of maximally distant pair of trees induces a partition of the edge-set of the graph that completely determines the outcome of the Shannon game.

Cutsets and circuits

The set of edges that connect a proper subset A of vertices to the rest of the graph (that is to the vertices in $V \setminus A$) is called a *cut* and will be denoted $\omega(A)$. Notice that $\omega(V \setminus A) = \omega(A)$, and if $\omega(A)$ is not empty, then the removal of $\omega(A)$ increases the number of connected components in the graph.

A *cutset* $\omega(A)$ is a minimal cut. That is $\omega(A)$ should not contain a smaller cut inclusion-wise. For example, the set of dashed edges {(2,3),(3,6),(4,6), (5,6) (5,7)} in Figure 6 is a cutset.

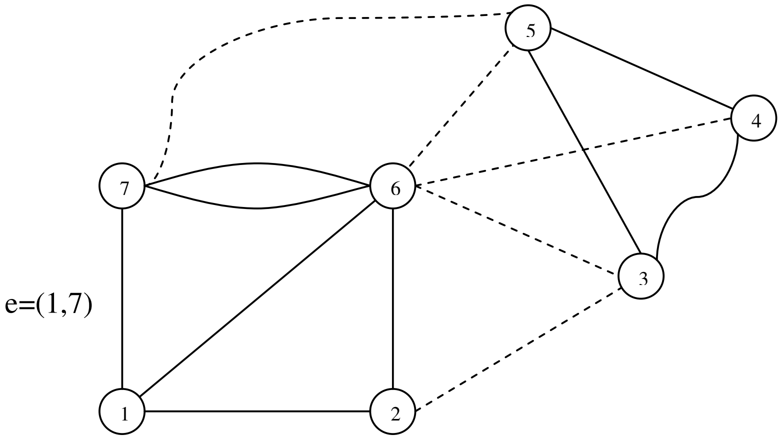

Figure 6 The dashed edges form the cutset $\omega(\{3,4,5\}) = \{(2,3),(3,6),(4,6),(5,6)\ (5,7)\}$

Whereas the cut $\omega(\{5\}) = \{(3,5),(4,5)\}$ in Figure 4 is not a cutset because it is not minimal as it contains a smaller cut, namely $\omega(\{3\}) = \{(3,5)\}$.

Circuits are minimal cycles and are the duals of cutsets. If S is a cutset and C is a circuit, then $|S \cap C|$, the cardinal of the intersection of S and C, is an even number.

Among the interesting properties that cutsets and circuits share is the swapping property;

Theorem 1 If C_1 and C_2 are circuits, if e is an edge of $C_1 \cap C_2$, and if e_1 is an edge of $C_1 \setminus C_2$, then there exists a circuit C such that $e_1 \in C$ and $C \subseteq C_1 \cup C_2 \setminus e$.

A similar result holds for cutsets;

Theorem 2 If S_1 and S_2 are cutsets, if e is an edge of $S_1 \cap S_2$, and if e_1 is an edge of $S_1 \setminus S_2$, then there exists a cutset S such that $e_1 \in S$ and $S \subseteq S_1 \cup S_2 \setminus e$.

By definition, a co-tree T^* is the complement of a tree T. A co-tree T^* is said to *span an edge e with respect the cutsets,* if the addition of e to T^* creates a cutset. Another beautiful theorem of Lehman tells us that a game can be won by CUT playing second if and only if there exist two co-spanning co-trees spanning the edge *e*. This statement is the dual of the statement we made for SHORT. The reduction illustrated in the introductory example concerned co-spanning trees.

The classification of Shannon games (with respect to the outcome of the game) is in an intimate way related to the *principal partition* (explained next) of the edges of a graph.

Principal partition

The edge set of a graph can be partitioned into three disjoint subsets $\{P, P^*, R\}$. This canonical partition is called the *principal partition.* The part P, called the *principal minor*, corresponds to the subset of edges that are spanned by at least one pair of disjoint co-spanning trees. The part P^* , called the *principal co-minor*, corresponds to the subset of edges that are spanned by at least one pair of disjoint co-spanning co-trees. The part R is the subset of the edges that are neither P nor in P^*. That is, $R = E - (P \cup P^*)$. Fortunately, the parts P and P^* can be determined easily once a pair of *maximally distant spanning trees* are known. A pair of spanning trees is maximally distant if the trees have the least possible number of edges in common.

Next, we sketch an algorithm that computes a pair of two maximally distant spanning trees and gives the principal partition as a by-product.

Maximally distant trees

To construct a pair of maximally distant trees, we can start with any pair of trees $\{T_B, T_W\}$, where T_B is a black tree (its edges are painted black), and T_W is a white tree (edged painted in white). If the trees T_B and T_W have no edges in common, they are obviously maximally distant. If there exists a *common chord* (an edge that does not belong to $T_B \cup T_W$), we can try swap this edge with a common edge of T_B and T_W to make the trees more distant.

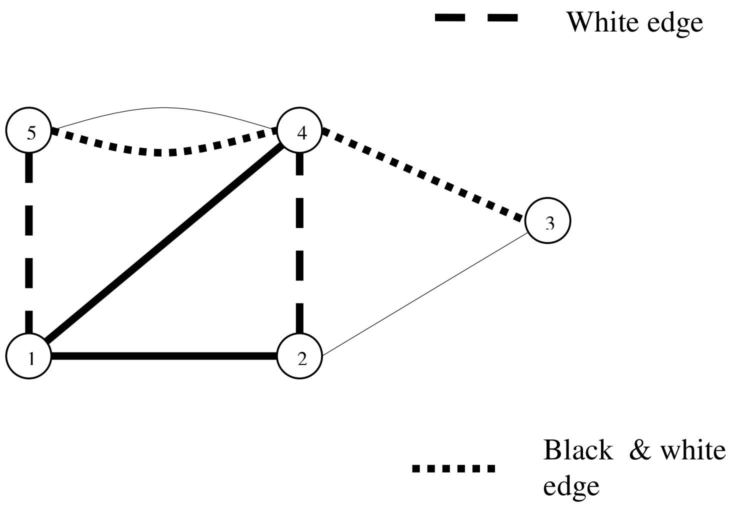

Figure 4 A pair of spanning trees not maximally distant.

For example, the edge (2,3) of Figure 4 can be used to make the Black tree and the White tree more distant. The black fundamental circuit of edge (2,3) is {(2,1),(1,4),(4,3)}. By swapping the common edge (4,3) with the edge (2,3), we obtain a new Black tree that has one less edge in common with the White tree (see Figure 5).

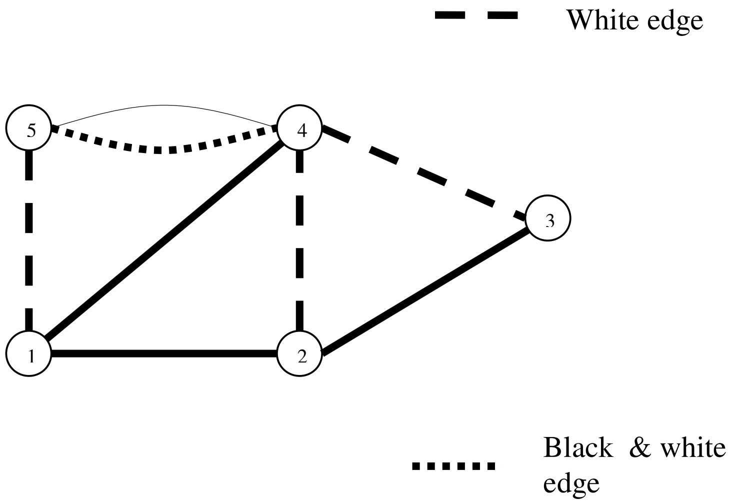

Figure 5 After replacing the common edge (3,4) with the common chord (2,3).

Even when the two fundamental circuits (one Black and one White) of a common chord do not contain a common edge, the trees might be made more distant by an augmenting swapping sequence.

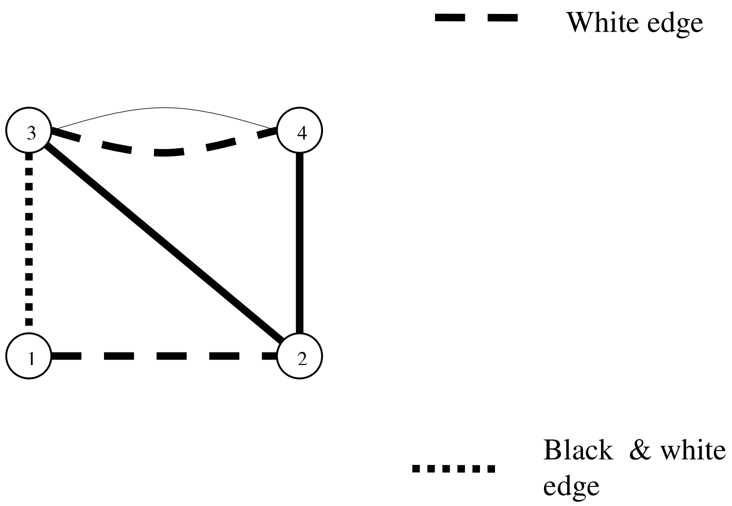

Figure 6 Swapping common chord (3,4) with any edge of its two fundamental circuits would not immediately make the Black and White trees more distant.

The White fundamental circuit for the common chord (3,4) is the White edge (3,4). The Black fundamental circuit is {(3,2),(2,4)}. None of the two fundamental circuits contains a common edge. Swapping the common chord (3,4) with any edge of the two fundamental circuits would not make the trees more distant directly. However, swapping the common chord (3,4) with (2,4), then swapping (2,4) with (1,3) would increase the distance between the two trees as illustrated in Figure 7.

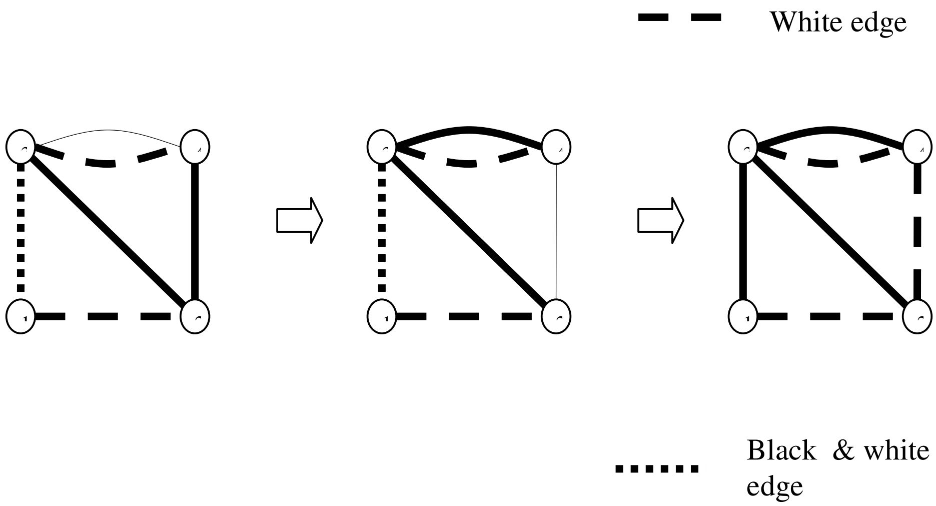

Figure 7 Augmenting swapping sequence; swap of common chord (3,4) with (2,4), followed by swap of (2,4) and (1,3).

Let call P the set of edges that can be reached by a swapping sequence with respect to the Black and White trees. The swapping sequences we consider always start with a common chord and alternate with fundamental circuits of the two trees. We can obtain a pair of maximally distant trees by searching for swapping sequences finishing at a common edge. When no such a sequence can be found, a theorem (unfortunately with a very long proof!) guarantees that the trees are maximally distant. Moreover, it has also been proved that the set of edges P is exactly the set of edges for which the Shannon game can be won for SHORT by starting second. The set P^* is obtained similarly by considering swapping sequences that use the fundamental cutsets instead of the fundamental circuits. The three possible outcomes of the Shannon games between perfect players are summarized in the following theorem.

Theorem 3 Let $\{P, P^*, R\}$ be the principal partition of a graph and e be the unplayable edge of a Shannon game.

- SHORT wins, even if playing second, if and only if $e \in P$.
- CUT wins, even if playing second, if and only if $e \in P^*$.
- Whoever plays first wins, if and only if $e \in R$

To implement a perfect SHORT player program, one has to write a set of functions to compute a pair of maximally distant spanning trees T_B and T_W using augmenting swapping sequences. A by-product of the search for the augmenting swapping sequences is P. Doing the same with cutsets yields P^*.

Once these preliminary computations are done, a game can be played. If $e \in P$ and e belongs neither to T_B nor to T_W, the artificial player can use directly the strategy sketched in the introductory example. The case where $e \in P$ and e belongs to T_B or to T_W requires some pre-processing; the artificial player has to first find a swapping sequence to modify one of the spanning tree so that e is no longer an edge of any of the pair of spanning trees. If $e \in R$, it means that no swapping sequence that would modify the pair of spanning trees so that e is outside the spanning trees can be found. Without loss of generality, assume that the edge e belongs to T_B. SHORT can imagine that e has been played by CUT and that the edge to span is not e, but a virtual common chord edge. This new view of the game reduces it to the first case (edge in P) that we know already how to play. A perfect CUT player would consider cutsets instead of circuits and use P^* instead of P.

References

"Hybrid Graph Theory and Network Analysis" by Novak and Gibbons, Cambridge University Press, 1999