FAIR ALLOCATION METHOD FOR STEINER TREE NETWORKS

Reynard M. Hilman
Department of Mathematics

Communications networks are costly to build and so the users of the network must contribute to the development and maintenance of the network. It would be good to allocate costs so that no group of users is asked to contribute more than they did if they were to build a network to meet their own needs. Such an allocation is called group rational. Unfortunately, group rational allocations are not always possible. This paper provides a computationally efficient algorithm for finding a group rational allocation in a special class of network configurations.

Reynard Hilman
Prof. David Housman
Summer 2000

Allocation Method for Steiner Tree Networks

A cost game is defined as a set N = {1, 2, … n} of players and cost c(S) associated with each subset S of players. Example:
N = {1, 2}
c(1) = 2 → cost of player 1 alone
c(2) = 3 → cost of player 2 alone
c(12) = 4 → cost if player 1 and 2 collaborate

The problem that arises from this cost game is how to allocate the total cost, when all players collaborate, to each player. Let us denote the allocation for a game as $x = (x_1, x_2, \ldots, x_n\}$ where x_i is the allocation for player i.

Core

An allocation is said to be in the core if the sum of the allocation of players in any subset S of N is ~~less~~ than the cost for that subset.

$$\sum_{s \in S} x_s \leq c(S), \quad \forall\, S \subset N$$

There are 2 common methods that are used to decide the allocation:

1. Shapley value which takes the average of the marginal contribution of each player over all possible orders. Marginal contribution for a player i is the ~~amount /~~ cost that a player i adds to a coalition when player i joins the coalition (c(S U {i}) – c(S) for i ∉ S. The Shapley value is not always in the core.
2. Nucleolus which lexicographically minimize the maximum coalitional complaints. Nucleolus is always in the core provided the cord is not empty.

Steiner tree

Given a set of points. Steiner tree is a tree that connects all points in S with possibility of having some additional points, called Steiner points, not in S. Example:

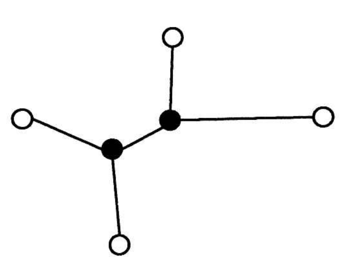

Stiner tree

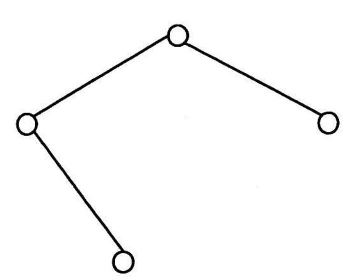

Spanning tree

Rectilinear Steiner tree is a Steiner tree that consist of only horizontal and vertical lines as the connector.

Game theory
Allocation Method for Steiner Tree Networks

Reynard Hilman
Prof. David Housman
Summer 2000

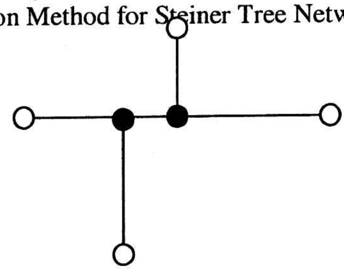

Rectilinear Steiner tree

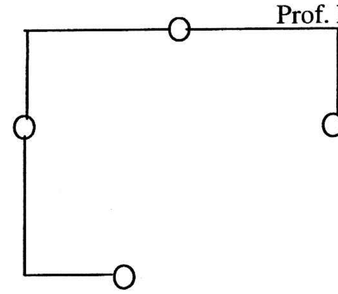

Rectilinear Spanning tree

The problem of finding the minimum rectilinear Steiner tree has been proven to be NP complete problem (*) therefore there is no polynomial time algorithm to find the MRST. However, some good algorithms have been found for some specific cases. In (**) Aho et. al. presented an O(n) dynamic programming algorithm to construct an MRST for points lying on two parallel lines.

Steiner tree on points lying on two parallel lines

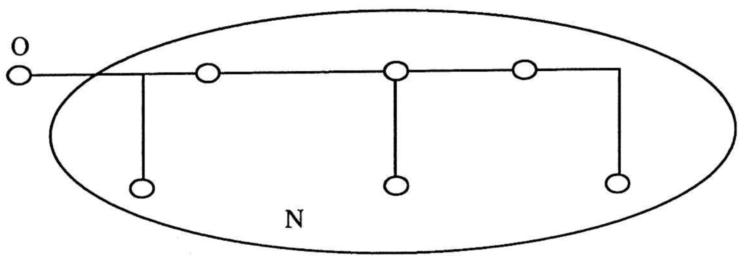

Let's denote a set of points lying on 2 parallel lines as N and the source where all points must be connected to as O. Let's further assume that the source is on the top left.

Notation:

d (i,j) is the rectilinear distance between points i and j
e (i,j) is the edge connecting point i and j
T_S is the minimal rectilinear Steiner tree on points in $S \cup \{O\}$
$|T_S|$ is the total length of the edges in the tree T_S
So T_N is the minimum tree on all points including the source and $|T_N|$ is the total length of the tree.

Since there are many MRST on a subset of points, we use the following tie breaking rules: No vertical lines could be shifted to the left.

1. Shift all vertical lines to the left as far as possible
2. Take the tree with less number of Steiner points. For Example:

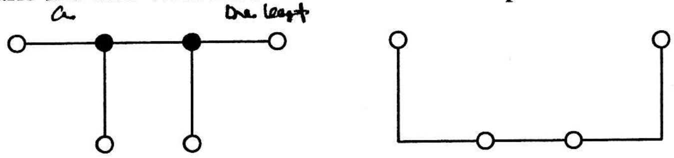

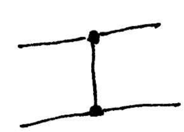

Reynard Hilman
Prof. David Housman
Summer 2000

3. Shift all Steiner points to left as far as possible.

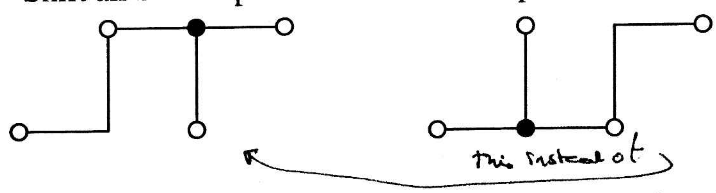

It is clear that there will not be any 3 adjacent Steiner points. Furthermore, from tie breaking rules no 2, we can always eliminate 2 adjacent Steiner points.
A point q is connected to the source through p if p lies on the unique path from the source O to q.
We can break the tree into smaller components. Let p be a point in N and S be a subset of all points that are connected to the source through p.

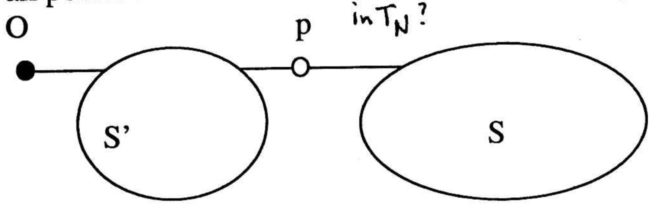

We can break this tree into two components:

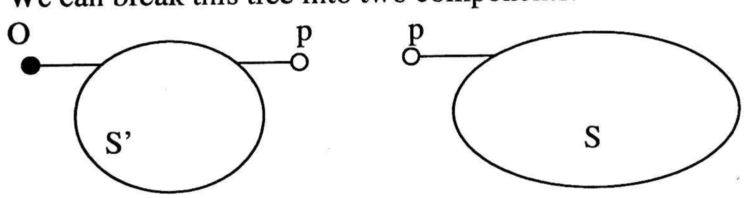

(i) (ii)

Let's denote each component as $T_{p,S}$ where p is the point in N that connect points in S to the source. $|T_{p,S}|$ is the length of the component. From the example above, component (i) will be denoted as $T_{O,S'}$ and component (ii) as $T_{p,S}$.
We can keep breaking each component until we get the smallest components.
Therefore the tree T_N can be broken into 4 basic components:

1.  $T_{p,\{u\}}$

2. 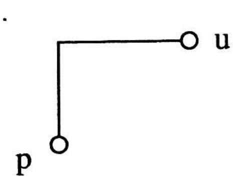 $T_{p,\{u\}}$

3. 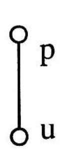 $T_{p,\{u\}}$

... $p \in N$. Let T_p denote the subgraph of T_N generated by p and its immediate neighbors in T_N that are further away from the origin. [ancestor

Since T_N is a tree on $\{0\} \cup N$, for each $q \in N$ there is a unique path connecting 0 to q. ~~The point~~ Let p be the point on this path that is adjacent to q. We call p the parent of q, and we call q a child of p. For $p \in \{0\} \cup N$, let T_p denote the subgraph of T_N generated by p and its children.

be the last point

last point in N on this path

No descendants.

p O

Descendant to the left

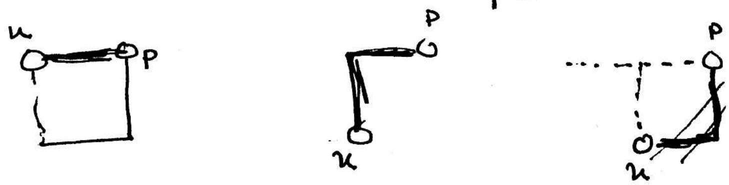

violates tree rule #1 T_N not a tree T_N not minimum length.

Lemma 1. Steiner points in T_N are not adjacent.

Lemma 2. ~~The immediate descendants are vertical or to the right.~~ A child is not found to the left of its parent.

Lemma 3. The subgraph generated by a Steiner point and its neighbors must look like

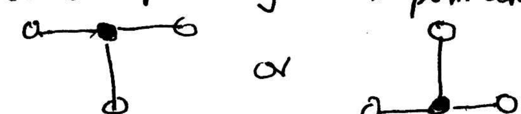

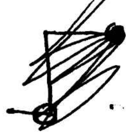

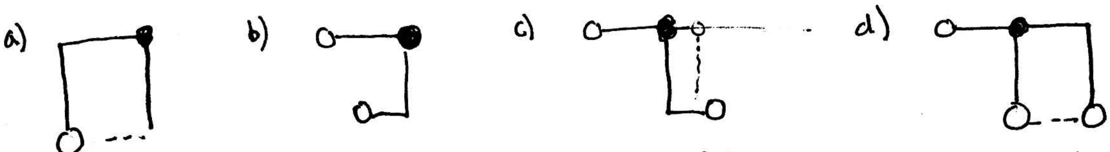

a) not minimum b) violates rule #1 c) not minimum d) not minimum.

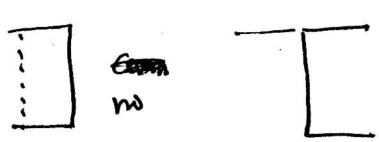

no rule #1

Any vertical line must be incident to a point in $N \cup \{0\}$

Reynard Hilman
Prof. David Housman
Summer 2000

4.

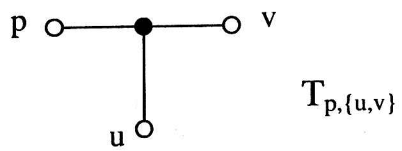

After we get the MRST on a set of points $S \cup \{O\}$ the next thing is how to divide the overall cost of the tree $|T_S|$ to each point in S. Following is a simple allocation that is in the core.

For basic component of type 1,2 and 3 we assign $x_u = |T_{p,\{u\}}|$ So u pays the amount that it needs to connect to p (which is the first point that connects u to the source.

For basic component of type 4

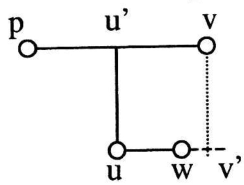

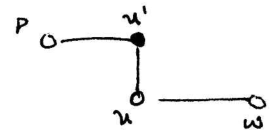

Where w is the first point connected to u that we encounter when we sweep a vertical line from v to the left (notice that w could be v' or u).

$x_v = \min\{ d(v,w), d(v,p) \}$

$x_u = | T_{p,\{u,v\}}| - x_v$

Note that in this case $d(u,w) \leq d(u,u')$ otherwise v can connect to w instead of u'

From this allocation, it is clear for basic component of type 1,2 and 3 that each point has to pay at least the closest distance to any point to it's left. $x_i \leq d(i,j)$ where j is the closest points to the left of i. However, for basic component of type 4 it is not immediately clear that x_u is less than the closest point to its left. Proof that x_u is less than or equal the closest point to its left:

$$\mathbf{x_u \leq \min\{ d(p,u), d(q,u) \}}$$

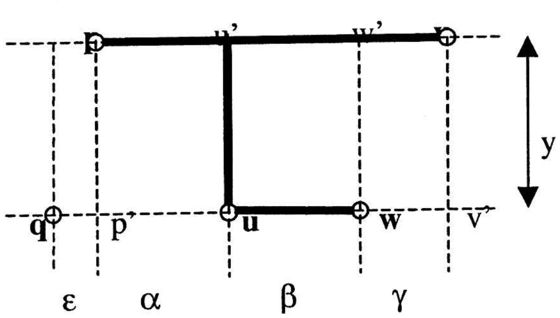

$Y = d(v,v')$

$\varepsilon = d(q,p')$

$\alpha = d(p',u)$

$\beta = d(u,w)$

$\gamma = d(w,v')$

$| T_{p,\{u,v\}}| = \alpha + \beta + \gamma + y$

$x_v = \min\{ (y + \gamma), (\alpha + \beta + \gamma) \}$

Reynard Hilman
Prof. David Housman
Summer 2000

I $x_v = y + \gamma$, $x_u = \alpha + \beta$ Or **II** $x_v = \alpha + \beta + \gamma$, $x_u = y$

$\mathbf{x_u \leq d\ (p,u) = \alpha + y}$

1. if $x_u = y$, it is clear that $\mathbf{x_u \leq d\ (p,u)}$
2. if $x_u = \alpha + \beta$, suppose $x_u > d\ (p,u)$

$$\alpha + \beta > \alpha + y$$
$$\beta > y$$

then we can replace e(u',w') with e(w,w') and get a new tree smaller than T_N (contradiction, T_N is the minimum tree). So, $\mathbf{x_u \leq d\ (p,u)}$

$\mathbf{x_u \leq d\ (q,u) = \varepsilon + \alpha}$

1. if $x_u = y$, suppose $x_u > d\ (q,u)$

$$y > d\ (q,u) = \varepsilon + \alpha$$

then we can replace e(u,u') with e(q,u) and get a new tree smaller than T_N (contradiction, T_N is the minimum tree). So, $\mathbf{x_u \leq d\ (q,u)}$

2. if $x_u = \alpha + \beta$, suppose $x_u > d\ (q,u)$

$$\alpha + \beta > \varepsilon + \alpha$$
$$\beta > \varepsilon$$

then we can replace e(u',w') with e(q,p'), shift e(p,u') down to e(p',u), shift e(u,u') to e(w,w'), and get a smaller three than T_N (contradiction, T_N is the minimum tree). So, $\mathbf{x_u \leq d\ (q,u)}$

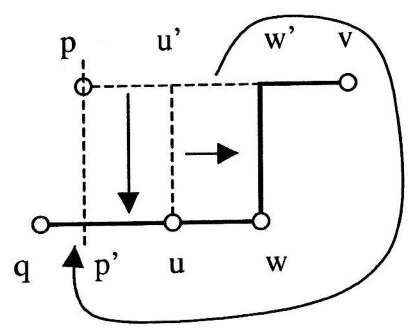

Following is the proof that the allocation is in the core:
Suppose the allocation is not in the core. Let's take a maximum S such that

$$|T_{O,S}| < \sum_{s \in S} x_s$$

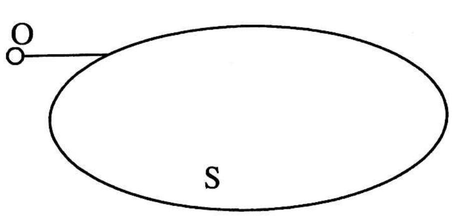

$$\exists\, p \notin S \ni |T_{O,\, S \cup \{p\}}| \geq \sum_{s \in S \cup \{p\}} x_s$$

Reynard Hilman
Prof. David Housman
Summer 2000

Before we move further, let's denote all points not in S that is connected directly (that is, not connected through another points) to any point in S as peripheral points of S. We are interested in the peripheral points of S, for the other points in N \ S can be connected to S through these peripheral points.

Let p' ∈ S, then p can not be any point ∈ N \ S that is connected directly to p' in $T_{O,N}$, such as component 1, 2, 3.

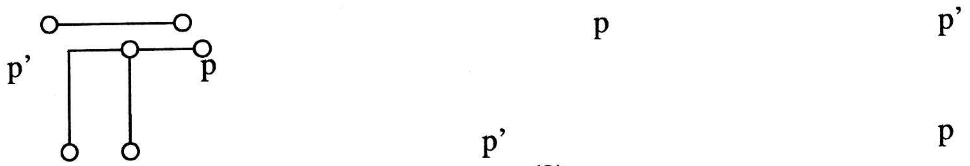

(1) (2) (3)

Because for (1) and (2) $x_p = d(p,p')$

for (3) $x_p \geq d(p,p')$

Therefore $| T_{O,S \cup \{p\}} | < \sum_{s \in S \cup \{p\}} x_s$

The only possible p is a point that is connected to a steiner point in $T_{O,N}$

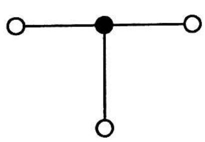

One of these three points must be p and at least one of them must be in S.

Therefore we have 6 cases to look at:

1. O

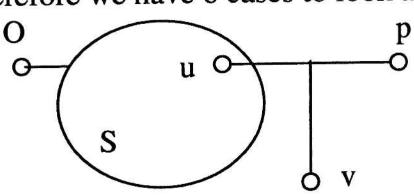

u ∈ S, v ∉ S

Since $| T_{u,\{v,p\}} | = x_v + x_p$

$$| T_{O,S \cup \{v,p\}} | < \sum_{s \in S \cup \{v,p\}} x_s$$

Therefore S is not maximum (contradiction)

2. O

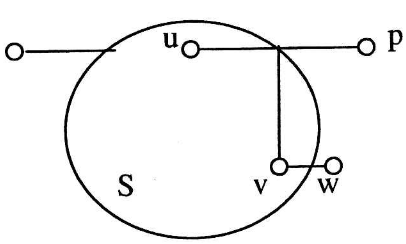

u ∈ S, v ∈ S

$x_p = \min\{\, d(u,p),\ d(v,w)\,\}$
where w is the rightmost point before p connected to v in T_N.
a. If $x_p = d(u,p)$

We still have the inequality:

$$\left| T_{O,S\cup\{p\}} \right| < \sum_{s\in S\cup\{p\}} x_s$$

Therefore S is not maximum (contradiction)
b. If $x_p = d(v,w)$

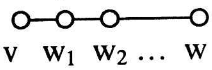

We can connect all points along that are not in S yet vw (R={w1,w2, …,w), this will cost at most ($x_{w1}+x_{w2}+\ldots+x_w$). and connect p to w, which cost x_p. So we still have the inequality:

$$\left| T_{O,S\cup R\cup\{p\}} \right| < \sum_{s\in S\cup R\cup\{p\}} x_s$$

Therefore S is not maximum (contradiction)

3.

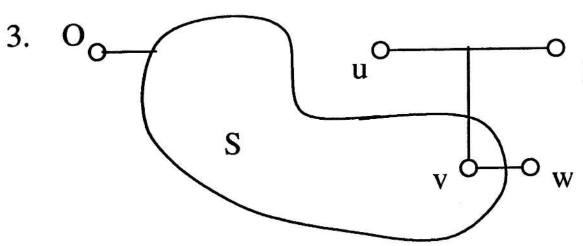

$v \in S,\ u \notin S$
$x_p = \min\{\, d(u,p),\ d(v,w)\,\}$
a. If $x_p = d(v,w)$ same as case 2b, we can connect all points along vw and connect p to w.
b. If $x_p = d(u,p)$.

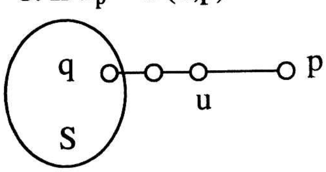

If there is a path in T_N cosisting of basic components of type 1, 2 or 3 from u to any point in S, u can be connected to S by connecting this series of points, then we can connect p to u and maintain the inequality:

$$\left| T_{O,S\cup R\cup\{p\}} \right| < \sum_{s\in S\cup R\cup\{p\}} x_s$$

where R is the set of points along the path from u to q (excluding q).
If there is no such path, it means that in T_N, u is connected to a peripheral point m that connects to S through a steiner point. Notice that anything below m and u is in S. So we only have 2 possibilities.

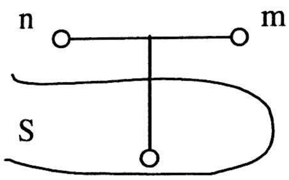

The 2 possibilities are:

3.1. If $n \in S$

$x_m \geq \min \{ d(m,m'), d(m,n) \}$

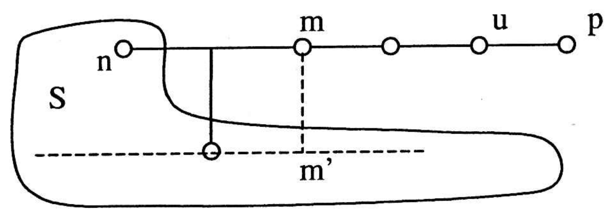

We can connect all points from m to p with cost less than $x_m + \ldots x_p$ and maintain the the inequality:

$$\left| T_{O,S \cup R \cup \{p\}} \right| < \sum_{s \in S \cup R \cup \{p\}} x_s$$

Therefore S is not maximum (contradiction)

3.2. If $n \notin S$

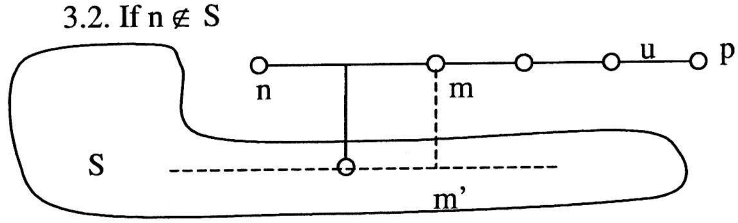

$x_m \geq \min \{ d(m,m'), d(m,n) \}$

If $x_m \geq d(m,m')$, we can connect all points from m to p with cost less than $x_m + \ldots x_p$ and maintain the the inequality:

$$\left| T_{O,S \cup R \cup \{p\}} \right| < \sum_{s \in S \cup R \cup \{p\}} x_s$$

Therefore S is not maximum (contradiction)

If $x_m \geq d(m,n)$ then we have to look at how n is connected to the source in T_N. So we can treat n as another u. Since the number of points is limited, there must be a point q in S such that connecting all series of points from q to p in T_N costs less than the allocation for those points.

$$\left| T_{O,S \cup R \cup \{p\}} \right| < \sum_{s \in S \cup R \cup \{p\}} x_s$$

Where R is the set of points on the path from q to p in T_N.

This also implies that S is not maximum (contradiction)

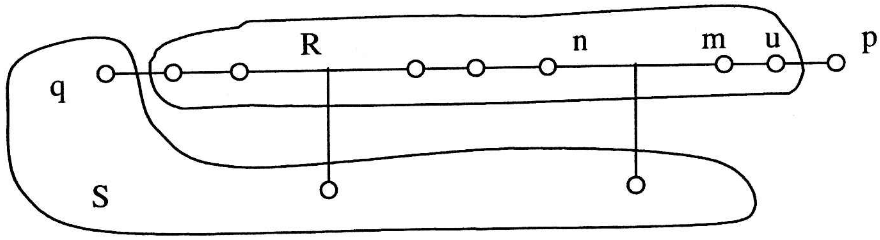

4.

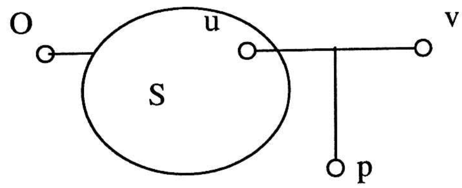

$u \in S, v \notin S$
same as case 1.

5.

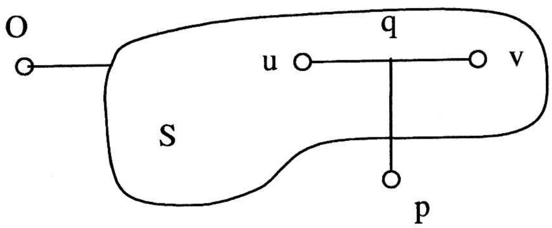

$u \in S, v \in S$
$x_p \bullet d(p,q)$

$$|T_{O,S \cup \{p\}}| < \sum_{s \in S \cup \{p\}} x_s$$

Therefore S is not maximum (contradiction)

6.

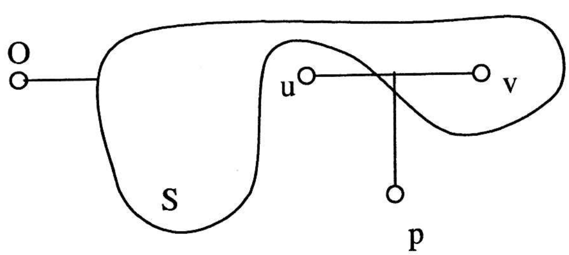

$v \in S, u \notin S$
There must be some path through p or u that connect v to source. Since p is not in S so the path must go thorough u. This implies that u must also be in S (same as case 5).

For the case where p is the point on the left:

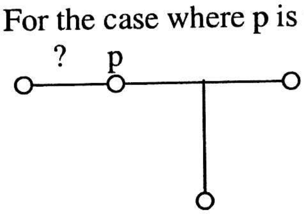

We are interested in looking at how p is connected to the source in T_N. So what we have to look is the cases with p on the right or bottom of the Steiner point. By symetry, this

Reynard Hilman
Prof. David Housman
Summer 2000

proof also applies for the case where p lies on the bottom line (by turing all pictures upside down).

Since we have shown that no such p exists, the assumption that the allocation is not in the core is wrong. Therefore the allocation is in the core.