Resolving sets for incidence graphs

Robert Bailey
University of Regina

23rd British Combinatorial Conference
Exeter
5th July 2011

Metric dimension

- Let $\Gamma = (V, E)$ be a graph.

Metric dimension

- Let $\Gamma = (V, E)$ be a graph.
- A subset $S = \{v_1, v_2, \ldots, v_k\} \subseteq V$ is a *resolving set* for Γ if, for every vertex w, the list of distances

  $$(\mathsf{d}(w, v_1), \mathsf{d}(w, v_2), \ldots, \mathsf{d}(w, v_k))$$

  is unique.

Metric dimension

- Let $\Gamma = (V, E)$ be a graph.
- A subset $S = \{v_1, v_2, \ldots, v_k\} \subseteq V$ is a *resolving set* for Γ if, for every vertex w, the list of distances

  $$(\mathsf{d}(w, v_1), \mathsf{d}(w, v_2), \ldots, \mathsf{d}(w, v_k))$$

  is unique.
- In other words, for all distinct vertices u, w, there exists $v_i \in S$ such that $\mathsf{d}(u, v_i) \neq \mathsf{d}(w, v_i)$.

Metric dimension

- Let $\Gamma = (V, E)$ be a graph.
- A subset $S = \{v_1, v_2, \ldots, v_k\} \subseteq V$ is a *resolving set* for Γ if, for every vertex w, the list of distances

$$(\mathsf{d}(w, v_1), \mathsf{d}(w, v_2), \ldots, \mathsf{d}(w, v_k))$$

  is unique.
- In other words, for all distinct vertices u, w, there exists $v_i \in S$ such that $\mathsf{d}(u, v_i) \neq \mathsf{d}(w, v_i)$.
- The *metric dimension* of Γ, denoted $\mu(\Gamma)$, is the smallest size of a resolving set for Γ.

Metric dimension

- Let $\Gamma = (V, E)$ be a graph.
- A subset $S = \{v_1, v_2, \ldots, v_k\} \subseteq V$ is a *resolving set* for Γ if, for every vertex w, the list of distances

$$(\mathsf{d}(w, v_1), \mathsf{d}(w, v_2), \ldots, \mathsf{d}(w, v_k))$$

  is unique.
- In other words, for all distinct vertices u, w, there exists $v_i \in S$ such that $\mathsf{d}(u, v_i) \neq \mathsf{d}(w, v_i)$.
- The *metric dimension* of Γ, denoted $\mu(\Gamma)$, is the smallest size of a resolving set for Γ.
- Metric dimension was introduced in the 1970s by Harary and Melter, and (independently) by Slater.

Examples

- Complete graphs: $\mu(K_n) = n - 1$.

Examples

- Complete graphs: $\mu(K_n) = n - 1$.
- Complete bipartite graphs: $\mu(K_{m,n}) = m + n - 2$.

Examples

- Complete graphs: $\mu(K_n) = n - 1$.
- Complete bipartite graphs: $\mu(K_{m,n}) = m + n - 2$.
- Trees: a precise formula due to Slater.

Examples

- Complete graphs: $\mu(K_n) = n - 1$.
- Complete bipartite graphs: $\mu(K_{m,n}) = m + n - 2$.
- Trees: a precise formula due to Slater.
- Petersen graph: $\mu(P) = 3$.

Example: Petersen graph

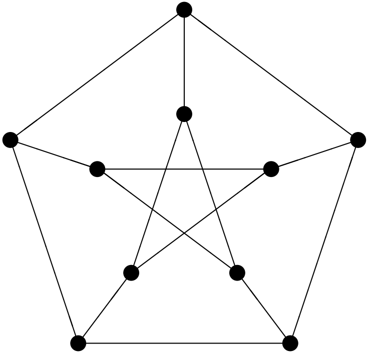

Example: Petersen graph

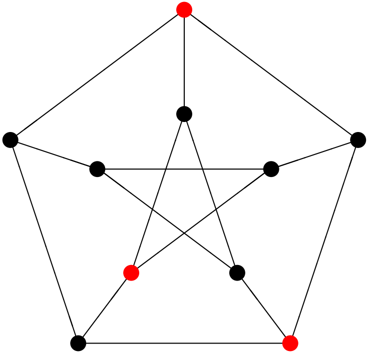

Symmetric designs

- A *symmetric design* with parameters (v, k, λ) is a pair $(X, \mathcal{B})$ such that:

Symmetric designs

- A *symmetric design* with parameters (v, k, λ) is a pair $(X, \mathcal{B})$ such that:
  - X is a set of v *points*

Symmetric designs

- A *symmetric design* with parameters (v, k, λ) is a pair $(X, \mathcal{B})$ such that:
  - X is a set of v *points*
  - $\mathcal{B}$ is a set of k-subsets of X, called *blocks*

Symmetric designs

- A *symmetric design* with parameters (v, k, λ) is a pair $(X, \mathcal{B})$ such that:
  - X is a set of v *points*
  - $\mathcal{B}$ is a set of k-subsets of X, called *blocks*
  - any pair of points lie in exactly λ blocks

Symmetric designs

- A *symmetric design* with parameters (v, k, λ) is a pair $(X, \mathcal{B})$ such that:
  - X is a set of v *points*
  - $\mathcal{B}$ is a set of k-subsets of X, called *blocks*
  - any pair of points lie in exactly λ blocks
  - any pair of blocks agree in exactly λ points.

Symmetric designs

- A *symmetric design* with parameters (v, k, λ) is a pair $(X, \mathcal{B})$ such that:
  - X is a set of v *points*
  - $\mathcal{B}$ is a set of k-subsets of X, called *blocks*
  - any pair of points lie in exactly λ blocks
  - any pair of blocks agree in exactly λ points.
- The last two conditions imply that $|\mathcal{B}| = v$.

Symmetric designs

- A *symmetric design* with parameters (v, k, λ) is a pair $(X, \mathcal{B})$ such that:
  - X is a set of v *points*
  - $\mathcal{B}$ is a set of k-subsets of X, called *blocks*
  - any pair of points lie in exactly λ blocks
  - any pair of blocks agree in exactly λ points.
- The last two conditions imply that $|\mathcal{B}| = v$.
- If $\lambda = 1$, we have a *finite projective plane*. (We call the blocks *lines* in that case.)

Symmetric designs: Example

There is a unique symmetric design with parameters $(7, 3, 1)$:

Symmetric designs: Example

There is a unique symmetric design with parameters $(7, 3, 1)$:

$1\,2\,4$
$2\,3\,5$
$3\,4\,6$
$4\,5\,7$
$5\,6\,1$
$6\,7\,2$
$7\,1\,3$

Symmetric designs: Example

There is a unique symmetric design with parameters $(7, 3, 1)$:

$1\,2\,4$
$2\,3\,5$
$3\,4\,6$
$4\,5\,7$
$5\,6\,1$
$6\,7\,2$
$7\,1\,3$

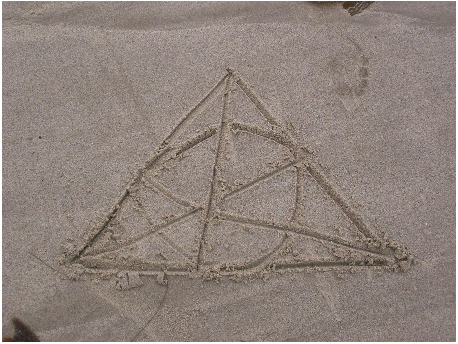

Incidence graphs

- From any symmetric design $(X, \mathcal{B})$, one can construct a bipartite graph as follows:

Incidence graphs

- From any symmetric design $(X, \mathcal{B})$, one can construct a bipartite graph as follows:
  - Vertices: $X \cup \mathcal{B}$

Incidence graphs

- From any symmetric design $(X, \mathcal{B})$, one can construct a bipartite graph as follows:
  - Vertices: $X \cup \mathcal{B}$
  - (p, B) is an edge iff $p \in B$.

Incidence graphs

- From any symmetric design $(X, \mathcal{B})$, one can construct a bipartite graph as follows:
  - Vertices: $X \cup \mathcal{B}$
  - (p, B) is an edge iff $p \in B$.
- This is the *incidence graph* (or *Levi graph*) of the design.

Incidence graphs: Example

The incidence graph of the Fano plane:

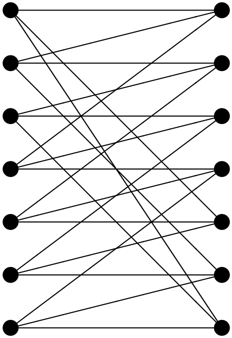

Incidence graphs: Example

The incidence graph of the Fano plane:

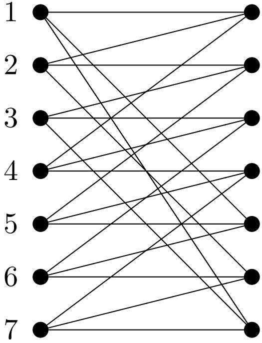

Incidence graphs: Example

The incidence graph of the Fano plane:

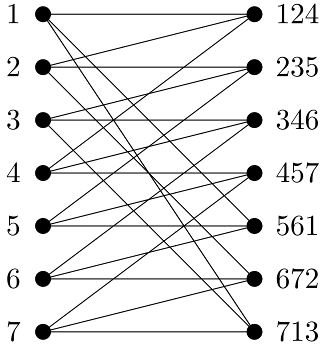

Distance in incidence graphs

- The incidence graph of a symmetric design is a bipartite graph of diameter 3:

Distance in incidence graphs

- The incidence graph of a symmetric design is a bipartite graph of diameter 3:
  - $\mathsf{d}(u, v) = 1$ iff (u, v) are an incident point and block (a *flag*)

Distance in incidence graphs

- The incidence graph of a symmetric design is a bipartite graph of diameter 3:
  - $\mathsf{d}(u, v) = 1$ iff (u, v) are an incident point and block (a *flag*)
  - $\mathsf{d}(u, v) = 2$ iff (u, v) are two distinct points or two distinct blocks

Distance in incidence graphs

- The incidence graph of a symmetric design is a bipartite graph of diameter 3:
  - $\mathsf{d}(u,v) = 1$ iff (u,v) are an incident point and block (a *flag*)
  - $\mathsf{d}(u,v) = 2$ iff (u,v) are two distinct points or two distinct blocks
  - $\mathsf{d}(u,v) = 3$ iff (u,v) are a non-incident point and block (an *antiflag*).

Distance in incidence graphs

- The incidence graph of a symmetric design is a bipartite graph of diameter 3:
  - $\mathsf{d}(u, v) = 1$ iff (u, v) are an incident point and block (a *flag*)
  - $\mathsf{d}(u, v) = 2$ iff (u, v) are two distinct points or two distinct blocks
  - $\mathsf{d}(u, v) = 3$ iff (u, v) are a non-incident point and block (an *antiflag*).
- So distances in the incidence graph can be entirely described using the structure of the design.

Distance in incidence graphs

- The incidence graph of a symmetric design is a bipartite graph of diameter 3:
  - $\mathsf{d}(u,v) = 1$ iff (u,v) are an incident point and block (a *flag*)
  - $\mathsf{d}(u,v) = 2$ iff (u,v) are two distinct points or two distinct blocks
  - $\mathsf{d}(u,v) = 3$ iff (u,v) are a non-incident point and block (an *antiflag*).
- So distances in the incidence graph can be entirely described using the structure of the design.
- From now on, we'll restrict ourselves to $\lambda = 1$, i.e. projective planes.

Resolving sets in incidence graphs

- Let $\Pi = (\mathcal{P}, \mathcal{L})$ be a projective plane of order q, and let Γ_Π denote its incidence graph.

Resolving sets in incidence graphs

- Let $\Pi = (\mathcal{P}, \mathcal{L})$ be a projective plane of order q, and let Γ_Π denote its incidence graph.
- Recall that there are $q^2 + q + 1$ points and $q^2 + q + 1$ lines, and that each line contains $q + 1$ points.

Resolving sets in incidence graphs

- Let $\Pi = (\mathcal{P}, \mathcal{L})$ be a projective plane of order q, and let Γ_Π denote its incidence graph.
- Recall that there are $q^2 + q + 1$ points and $q^2 + q + 1$ lines, and that each line contains $q + 1$ points.
- Any resolving set S for Γ_Π will have the form $S_\mathcal{P} \cup S_\mathcal{L}$.

Resolving sets in incidence graphs

- Let $\Pi = (\mathcal{P}, \mathcal{L})$ be a projective plane of order q, and let Γ_Π denote its incidence graph.
- Recall that there are $q^2 + q + 1$ points and $q^2 + q + 1$ lines, and that each line contains $q + 1$ points.
- Any resolving set S for Γ_Π will have the form $S_\mathcal{P} \cup S_\mathcal{L}$.
- Note that, since any pair of points are at distance 2, if $S \subseteq \mathcal{P}$, we must take at least $|\mathcal{P}| - 1 = q^2 + q$ points.

Resolving sets in incidence graphs

- Let $\Pi = (\mathcal{P}, \mathcal{L})$ be a projective plane of order q, and let Γ_Π denote its incidence graph.
- Recall that there are $q^2 + q + 1$ points and $q^2 + q + 1$ lines, and that each line contains $q + 1$ points.
- Any resolving set S for Γ_Π will have the form $S_\mathcal{P} \cup S_\mathcal{L}$.
- Note that, since any pair of points are at distance 2, if $S \subseteq \mathcal{P}$, we must take at least $|\mathcal{P}| - 1 = q^2 + q$ points.
- So, for a smaller resolving set, we'll need to use both points and lines.

Split resolving sets in bipartite graphs

- Suppose $\Gamma = (V, E)$ is a bipartite graph, with bipartition $V = X \dot{\cup} Y$.

Split resolving sets in bipartite graphs

- Suppose $\Gamma = (V, E)$ is a bipartite graph, with bipartition $V = X \dot{\cup} Y$.
- A set of vertices

$$S_X \cup S_Y = \{x_1, \ldots, x_r\} \cup \{y_1, \ldots, y_s\}$$

(where $x_i \in X$, $y_i \in Y$) is called a *split resolving set* if:
  - for all $x \in X$, $\exists\, y_i, y_j \in S_Y$ with $\mathsf{d}(x, y_i) \neq \mathsf{d}(x, y_j)$;
  - for all $y \in Y$, $\exists\, x_a, x_b \in S_X$ with $\mathsf{d}(y, x_a) \neq \mathsf{d}(y, x_b)$.

Split resolving sets in bipartite graphs

- Suppose $\Gamma = (V, E)$ is a bipartite graph, with bipartition $V = X \dot{\cup} Y$.
- A set of vertices

$$S_X \cup S_Y = \{x_1, \ldots, x_r\} \cup \{y_1, \ldots, y_s\}$$

(where $x_i \in X$, $y_i \in Y$) is called a *split resolving set* if:
  - for all $x \in X$, $\exists\, y_i, y_j \in S_Y$ with $\mathsf{d}(x, y_i) \neq \mathsf{d}(x, y_j)$;
  - for all $y \in Y$, $\exists\, x_a, x_b \in S_X$ with $\mathsf{d}(y, x_a) \neq \mathsf{d}(y, x_b)$.
- If a split resolving set exists in Γ, we denote the smallest possible size by $\mu^*(\Gamma)$. (Otherwise, set $\mu^*(\Gamma) = \infty$.)

Split resolving sets in bipartite graphs

- Suppose $\Gamma = (V, E)$ is a bipartite graph, with bipartition $V = X \dot{\cup} Y$.
- A set of vertices

$$S_X \cup S_Y = \{x_1, \ldots, x_r\} \cup \{y_1, \ldots, y_s\}$$

  (where $x_i \in X$, $y_i \in Y$) is called a *split resolving set* if:
  - for all $x \in X$, $\exists\, y_i, y_j \in S_Y$ with $\mathsf{d}(x, y_i) \neq \mathsf{d}(x, y_j)$;
  - for all $y \in Y$, $\exists\, x_a, x_b \in S_X$ with $\mathsf{d}(y, x_a) \neq \mathsf{d}(y, x_b)$.
- If a split resolving set exists in Γ, we denote the smallest possible size by $\mu^*(\Gamma)$. (Otherwise, set $\mu^*(\Gamma) = \infty$.)
- Clearly a split resolving set is an (ordinary) resolving set, and so $\mu^*(\Gamma) \geq \mu(\Gamma)$.

Split resolving sets in bipartite graphs

- Suppose $\Gamma = (V, E)$ is a bipartite graph, with bipartition $V = X \dot{\cup} Y$.
- A set of vertices

$$S_X \cup S_Y = \{x_1, \ldots, x_r\} \cup \{y_1, \ldots, y_s\}$$

  (where $x_i \in X$, $y_i \in Y$) is called a *split resolving set* if:
    - for all $x \in X$, $\exists\, y_i, y_j \in S_Y$ with $\mathsf{d}(x, y_i) \neq \mathsf{d}(x, y_j)$;
    - for all $y \in Y$, $\exists\, x_a, x_b \in S_X$ with $\mathsf{d}(y, x_a) \neq \mathsf{d}(y, x_b)$.
- If a split resolving set exists in Γ, we denote the smallest possible size by $\mu^*(\Gamma)$. (Otherwise, set $\mu^*(\Gamma) = \infty$.)
- Clearly a split resolving set is an (ordinary) resolving set, and so $\mu^*(\Gamma) \geq \mu(\Gamma)$.
- We call the sets S_X and S_Y *semi-resolving sets* for Γ.

Resolving lines with points

- A set of points $S_{\mathcal{P}}$ is a semi-resolving set for Γ_{Π} if, for any two distinct lines L_1, L_2, there is a point $p \in S_{\mathcal{P}}$ with $\mathsf{d}(p, L_1) \neq \mathsf{d}(p, L_2)$.

Resolving lines with points

- A set of points $S_{\mathcal{P}}$ is a semi-resolving set for Γ_{Π} if, for any two distinct lines L_1, L_2, there is a point $p \in S_{\mathcal{P}}$ with $\mathsf{d}(p, L_1) \neq \mathsf{d}(p, L_2)$.
- There is an equivalent notion for lines.

Resolving lines with points

- A set of points $S_{\mathcal{P}}$ is a semi-resolving set for Γ_{Π} if, for any two distinct lines L_1, L_2, there is a point $p \in S_{\mathcal{P}}$ with $\mathsf{d}(p, L_1) \neq \mathsf{d}(p, L_2)$.
- There is an equivalent notion for lines.
- $S_{\mathcal{P}}$ is a semi-resolving set if and only if P lies on *exactly one* of the lines L_1, L_2.

Resolving lines with points

- A set of points $S_{\mathcal{P}}$ is a semi-resolving set for Γ_{Π} if, for any two distinct lines L_1, L_2, there is a point $p \in S_{\mathcal{P}}$ with $\mathsf{d}(p, L_1) \neq \mathsf{d}(p, L_2)$.
- There is an equivalent notion for lines.
- $S_{\mathcal{P}}$ is a semi-resolving set if and only if P lies on *exactly one* of the lines L_1, L_2.
- How can we construct such a set?

Blocking sets and double blocking sets

- A *blocking set* in a projective plane Π is a collection of points which intersects every line in at least one point.

Blocking sets and double blocking sets

- A *blocking set* in a projective plane Π is a collection of points which intersects every line in at least one point.
- A *t-fold blocking set* in Π is a collection of points which intersects every line in at least t points.

Blocking sets and double blocking sets

- A *blocking set* in a projective plane Π is a collection of points which intersects every line in at least one point.
- A *t-fold blocking set* in Π is a collection of points which intersects every line in at least t points.
- If $t = 2$, we have a *double blocking set*.

Blocking sets and double blocking sets

- A *blocking set* in a projective plane Π is a collection of points which intersects every line in at least one point.
- A t*-fold blocking set* in Π is a collection of points which intersects every line in at least t points.
- If $t = 2$, we have a *double blocking set*.
- **Proposition:** A double blocking set D for Π is a semi-resolving set for Γ_{Π}.

Blocking sets and double blocking sets

- A *blocking set* in a projective plane Π is a collection of points which intersects every line in at least one point.
- A *t-fold blocking set* in Π is a collection of points which intersects every line in at least t points.
- If $t = 2$, we have a *double blocking set*.
- **Proposition:** A double blocking set D for Π is a semi-resolving set for Γ_Π.
- **Proof:**
  - By the definition of a double blocking set, D contains two points $p, q \in L_1$.

Blocking sets and double blocking sets

- A *blocking set* in a projective plane Π is a collection of points which intersects every line in at least one point.
- A *t-fold blocking set* in Π is a collection of points which intersects every line in at least t points.
- If $t = 2$, we have a *double blocking set*.
- **Proposition:** A double blocking set D for Π is a semi-resolving set for Γ_Π.
- **Proof:**
  - By the definition of a double blocking set, D contains two points $p, q \in L_1$.
  - By the definition of a projective plane, at most one of these also lies on L_2.

Blocking sets and double blocking sets

- A *blocking set* in a projective plane Π is a collection of points which intersects every line in at least one point.
- A *t-fold blocking set* in Π is a collection of points which intersects every line in at least t points.
- If $t = 2$, we have a *double blocking set*.
- **Proposition:** A double blocking set D for Π is a semi-resolving set for Γ_{Π}.
- **Proof:**
  - By the definition of a double blocking set, D contains two points $p, q \in L_1$.
  - By the definition of a projective plane, at most one of these also lies on L_2.
- In fact, we can delete a point from a double blocking set and still have a semi-resolving set.

Upper bounds on metric dimension

- It is easy to construct a double blocking set of size $3q$, by taking three non-concurrent lines.

Upper bounds on metric dimension

- It is easy to construct a double blocking set of size $3q$, by taking three non-concurrent lines.
- Thus $\mu(\Gamma_\Pi) \leq \mu^*(\Gamma_\Pi) \leq 6q - 2$.

Upper bounds on metric dimension

- It is easy to construct a double blocking set of size $3q$, by taking three non-concurrent lines.
- Thus $\mu(\Gamma_\Pi) \leq \mu^*(\Gamma_\Pi) \leq 6q - 2$.
- However, better examples are known: if $q = p^{2e}$ and Π is the Desarguesian plane $\mathrm{PG}(2, q)$, there is a double blocking set of size $2q + 2\sqrt{q} + 2$ (the union of two disjoint *Baer subplanes*).

Upper bounds on metric dimension

- It is easy to construct a double blocking set of size $3q$, by taking three non-concurrent lines.
- Thus $\mu(\Gamma_\Pi) \leq \mu^*(\Gamma_\Pi) \leq 6q - 2$.
- However, better examples are known: if $q = p^{2e}$ and Π is the Desarguesian plane $\mathrm{PG}(2, q)$, there is a double blocking set of size $2q + 2\sqrt{q} + 2$ (the union of two disjoint *Baer subplanes*).
- Thus $\mu^*(\Gamma_{\mathrm{PG}(2,q)}) \leq 4q + 4\sqrt{q} + 2$ (if $q = p^{2e}$).

Lower bounds

- Aart Blokhuis (personal communication) has determined that, in a semi-resolving set $S_{\mathcal{P}}$, we must have $|S_{\mathcal{P}}| \gtrsim 2q + \sqrt{2q}$.

Lower bounds

- Aart Blokhuis (personal communication) has determined that, in a semi-resolving set $S_{\mathcal{P}}$, we must have $|S_{\mathcal{P}}| \gtrsim 2q + \sqrt{2q}$.
- Thus the size of a *split* resolving set for Γ_{Π} has size $\gtrsim 4q + 2\sqrt{2q}$.

Lower bounds

- Aart Blokhuis (personal communication) has determined that, in a semi-resolving set $S_{\mathcal{P}}$, we must have $|S_{\mathcal{P}}| \gtrsim 2q + \sqrt{2q}$.
- Thus the size of a *split* resolving set for Γ_{Π} has size $\gtrsim 4q + 2\sqrt{2q}$.
- In particular, in the case of $\mathrm{PG}(2, p^{2e})$, the upper and lower bounds are very similar!

A non-split resolving set

Bill Martin (not-so-personal communication) has come up with a non-split resolving set for Γ_{Π} of size $4q-1$, consisting of $2q$ lines and $2q-1$ points:

A non-split resolving set

Bill Martin (not-so-personal communication) has come up with a non-split resolving set for Γ_Π of size $4q-1$, consisting of $2q$ lines and $2q-1$ points:

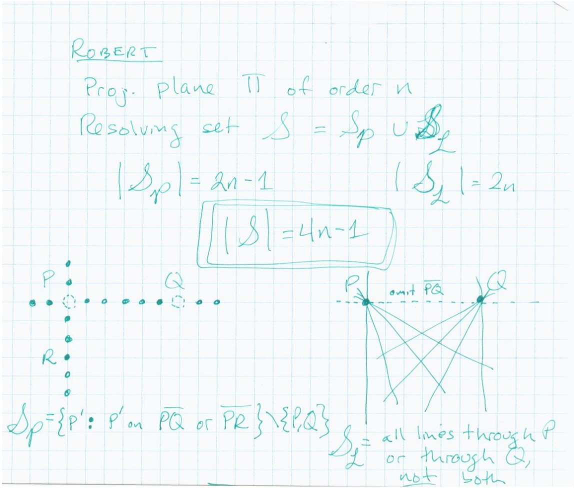

THE END