I

Groups and graphs

Sections 1 and 2 collect together the basic definitions on group actions and graphs, and Section 3 introduces the concept of a graph of groups. Section 4 then describes the structure of a group acting on a tree in terms of the fundamental group of a graph of groups. Section 5 lists some examples of trees arising in nature. Section 6 motivates the main argument of Section 7, which shows the converse of the structure theorem for groups acting on trees, that is, the fundamental group of a graph of groups acts on a tree; some applications in combinatorial group theory are then given. This is continued in Sections 8 and 10, where some important theorems on free groups and free products are proved, while Section 9 gives the structure theorem for groups acting on connected graphs.

1 Groups

The purpose of this section is to recall a list of basic definitions which will be needed throughout.

1.1 Definitions. Let S be a set.

We write $S^{\pm 1}$ for $S \times \{1, -1\}$, and denote an element (s, ε) by s^{ε}.

By a *word* in $S^{\pm 1}$ we mean a finite sequence in $S^{\pm 1}$, possibly empty. The word $(s_1^{\varepsilon_1}, \ldots, s_n^{\varepsilon_n})$ will usually be abbreviated $s_1^{\varepsilon_1} \cdots s_n^{\varepsilon_n}$.

Let $W(S)$ be the set of all words in $S^{\pm 1}$. There is a binary operation $W(S) \times W(S) \to W(S), (w, w') \mapsto ww'$, given by concatenation, and a unary operation $W(S) \to W(S)$, $w \mapsto w^{-1}$, given by $(s_1^{\varepsilon_1} \cdots s_n^{\varepsilon_n})^{-1} = s_n^{-\varepsilon_n} \cdots s_1^{-\varepsilon_1}$.

For any function $\alpha: S \to G, s \mapsto \alpha s$, there is induced a function $\alpha: W(S) \to G$, $s_1^{\varepsilon_1} \cdots s_n^{\varepsilon_n} \mapsto \alpha(s_1)^{\varepsilon_1} \cdots \alpha(s_n)^{\varepsilon_n}$.

Let R be a subset of $W(S)$. We say G has a *presentation* with *generating* *set* S and *relation set* R, and write $G = \langle S | R \rangle$, if the following holds: there is specified a function $\alpha : S \to G$ such that $\alpha(w) = 1$ for all $w \in R$, having the property that for any group H and function $\beta : S \to H$ such that $\beta(w) = 1$ for all $w \in R$, there exists a unique group homomorphism $\phi : G \to H$ such that $\beta = \phi\alpha$. Even though α need not be injective, we usually suppress α and use the same symbol to denote an element of S and its image in G, hoping that the meaning is clear from the context. In essence, S can be thought of as a family of elements of G, possibly having repetitions.

Variations of the prose are: $\langle S | R \rangle$ presents G; G has a presentation with generators $s \in S$ and relators $r \in R$, or relations $r = 1$, $r \in R$. In the latter formulation it is often convenient to write a relation of the form $w_1 w_2 = 1$ as $w_1 = w_2^{-1}$.

Given any subset R of $W(S)$ there exists a group presented by $\langle S | R \rangle$; to prove this, one considers the intersection of all equivalence relations induced on $W(S)$ by the various possible β's, and takes as G the set of equivalence classes, with multiplication induced by concatenation.

Any two groups presented by $\langle S | R \rangle$ are isomorphic, and the isomorphism is unique if the family S is respected.

Conversely, G always has some presentation, for example $\langle G | R \rangle$ where $R = \{((a,1),(b,1),(ab,-1)) \in W(G) | a, b \in G\}$; we refer to the elements of the latter set as *the relations for* G.

In specific cases, it is usual to list the elements of S and R, casually omitting the set brackets. We also use exponents to indicate repetition. For example, for any $n \geqslant 1$, $\langle s | s^n \rangle$ presents the cyclic group C_n of order n, and $\langle r, s | r^2, s^2, (rs)^n \rangle$ presents the dihedral group D_n of order $2n$. This extends by analogy to $n = \infty$, with $C_\infty = \langle s | \varnothing \rangle$, $D_\infty = \langle r, s | r^2, s^2 \rangle$.

The *rank* of G, denoted rank(G), is the minimum number of generators of G; that is, the least cardinal n such that there exists a presentation $\langle S | R \rangle$ of G with $|S| = n$.

For example, the only group of rank zero is the *trivial* group $G = 1$.

For another example, for any set S, if $R = \{w^2 | w \in W(S)\}$, then $\langle S | R \rangle$ has the structure of a vector space of dimension $|S|$ over the field of two elements; as this cannot be generated by fewer than $|S|$ elements, its rank is $|S|$.

We say that G is a *free group* if it has a presentation of the form $\langle S | \varnothing \rangle$. In this event, G is said to be *freely generated* by S, and that S is a *free generating set* of G. The previous example shows that $|S| = \text{rank}(G)$. For any cardinal n, we write F_n for the free group of rank n.

If S is a subset of G, we write $\langle S \rangle$ for the subgroup of G *generated by* S, that is the smallest subgroup of G containing S. ∎

1.2 Definitions. By a *G-set* X we mean a set given with a function $G \times X \to X$, $(g, x) \mapsto gx$, such that $1x = x$ for all $x \in X$, and $g(g'x) = (gg')x$ for all $g, g' \in G$, $x \in X$. This is equivalent to specifying a group homomorphism from G to Sym X, the group of all permutations of X, written on the left. We say also that G *acts* on X, and that there is a *G-action* on X.

For example, G is a G-set under left multiplication; more generally, if H is any subgroup of G then the set of right cosets, $G/H = \{xH \mid x \in G\}$, is a G-set with G-action given by $g(xH) = (gx)H$. We denote the cardinal of this set by $(G:H)$, called the *index* of H in G.

For another example, G is a G-set under *left conjugation*, given by ${}^g x = gxg^{-1}$.

If $X_i, i \in I$, is a family of G-sets then the disjoint union $\bigvee_{i \in I} X_i$ is a G-set, as is the Cartesian product $\prod_{i \in I} X_i$, where G is said to act *diagonally*.

A function $\alpha: X_1 \to X_2$ between G-sets is said to be a *G-map* if $\alpha(gx) = g(\alpha x)$ for all $g \in G, x \in X_1$. We say X_1, X_2 are *G-isomorphic*, denoted $X_1 \approx X_2$, if there exists a bijective G-map from one to the other.

By a *right G-set* X we mean a set given with a function $X \times G \to X$, $(x, g) \mapsto xg$, such that $x1 = x$ for all $x \in X$, and $(xg)g' = x(gg')$ for all $g, g' \in G$, $x \in X$. This is equivalent to X being a G-set with G-action $gx = xg^{-1}$. For example, we have *right conjugation* $x^g = g^{-1}xg$. ∎

1.3 Definitions. Let X be a G-set.

Let $x \in X$. By the *G-stabilizer* of x we mean the subgroup $G_x = \{g \in G \mid gx = x\}$ of G; if P is any subset or element of G_x we say that x is *stabilized by* P, or is *P-stable*. If $g \in G$, then $G_{gx} = {}^g G_x$, where for a subgroup H of G, we write ${}^g H$ and H^g for the *left conjugate* and *right conjugate* $gHg^{-1}, g^{-1}Hg$, respectively.

We say that G acts *trivially* if $gx = x$ for all $g \in G, x \in X$.

We say that X is a *G-free* G-set if $G_x = 1$ for all $x \in X$. For example, if S is a set with trivial G-action then $G \times S$ is G-free.

Since G acts on the set of subsets of X with $gX' = \{gx \mid x \in X'\}$ for $g \in G, X' \subseteq X$, this terminology extends to subsets of X. If X' is G-stable then we say that X' is a *G-subset* of X.

Similarly, G acts on the set of finite sequences $x_1, \ldots, x_n$ in X, so the notation applies here, and $G_{x_1, \ldots, x_n} = G_{x_1} \cap \cdots \cap G_{x_n}$.

For $x \in X$, the *G-orbit* of x is $Gx = \{gx \mid g \in G\}$, a G-subset of X which is G-isomorphic to G/G_x with $gx \in Gx$ corresponding to $gG_x \in G/G_x$.

By the *quotient set* for the G-set X, we mean $G \backslash X = \{Gx \mid x \in X\}$, the set of G-orbits; there is a natural map $X \to G\backslash X, x \mapsto Gx$. If $G\backslash X$ is finite we say that X is *G-finite*.

By a *G-transversal* in X we mean a subset S of X which meets each G-orbit exactly once, so the composite $S \subseteq X \to G\backslash X$ is bijective. Then X is G-isomorphic to $\bigvee_{s\in S} G/G_s$ with $gG_s \in \bigvee_{s\in S} G/G_s$ corresponding to $gs \in X$, for all $g\in G$, $s\in S$. Hence X is the G-set presented on the generating set S with relations saying that s is G_s-stable for each $s\in S$. ∎

1.4 Remarks. (i) Notice we have a structure theorem for G-sets, which says that a G-set is specified up to G-isomorphism by a G-transversal and the G-stabilizers of the elements of the G-transversal.

For example, a G-set is G-free if and only if it is a disjoint union of copies of G, or equivalently, of the form $G \times S$.

(ii) If $\alpha: X \to Y$ is a map of G-sets then $G_x \subseteq G_{\alpha x}$ for all $x\in X$, and if α is injective then $G_x = G_{\alpha x}$ for all $x\in X$. For example, the only G-sets which have G-maps to free G-sets are the free G-sets.

(iii) Conversely, suppose X, Y, are G-sets, and for each $x\in X$, G_x stabilizes an element of Y. Then we can choose any G-transversal S in X and construct a function $\alpha: S \to Y$ such that $G_s \subseteq G_{\alpha s}$ for all $s\in S$. Now α extends to a well-defined G-map $X \to Y, gs \mapsto g\alpha(s)$. ∎

2 Graphs

We now come to another list of basic concepts, this time somewhat less standard.

2.1 Definitions. By a *G-graph* (X, V, E, ι, τ) we mean a nonempty G-set X with a specified nonempty G-subset V, its complement $E = X - V$, and two G-maps $\iota, \tau: E \to V$. In this event we say simply that X is a G-graph.

For any G-subset Y of X we write $VY = V \cap Y, EY = E \cap Y$. If Y is nonempty, and for each $e\in EY$ both ιe and τe belong to VY, then Y is said to be a *G-subgraph* of X.

In particular, $VX = V, EX = E$. We call V and E the *vertex set* and *edge set* of X, and the elements *vertices* and *edges* of X, respectively. The functions $\iota, \tau: E \to V$ are the *incidence functions* of X.

If e is any edge then ιe and τe are the vertices *incident* to e, and are called the *initial* and *terminal* vertices of e, respectively. The definition allows the possibility that ιe and τe may be equal, in which case e is called a *loop*. In almost all our examples the G-map $(\iota, \tau): E \to V \times V$ will be injective, and here $G_e = G_{\iota e, \tau e} = G_{\iota e} \cap G_{\tau e}$ for all $e\in E$.

For $v \in V$, we define $\mathrm{star}(v) = \iota^{-1}(v) \vee \tau^{-1}(v)$, sometimes called the *neighbourhood* of v. The number of elements in star(v) is called the *valency* of v; the elements of star(v) are the edges *incident to* v, either *going into* v or *going out of* v, depending on whether they belong to $\tau^{-1}(e)$ or $\iota^{-1}(e)$, respectively, possibly both. The vertices joined to v by an edge are called the *neighbours* of v.

If every vertex of X has finite valency then X is said to be *locally finite*.

By a *geometric realization of* X we mean an oriented one-dimensional CW-complex with V the set of zero-cells and E the set of one-cells with each edge e starting at ιe and finishing at τe.

For G-graphs X, Y, a *G-graph map* $\alpha: X \to Y$ is a G-map such that $\alpha(VX) \subseteq VY, \alpha(EX) \subseteq EY$, and for each $e \in EX$, $\alpha(\iota e) = \iota(\alpha e)$, $\alpha(\tau e) = \tau(\alpha e)$.

The terms *G-graph isomorphism* and *G-graph automorphism* are then defined in the natural way.

In all the above phrases, if G is omitted we understand that $G = 1$; in this way we recover the concepts of *graph*, *subgraph* and *graph map*. Thus a G-graph may be viewed as a graph given with a homomorphism from G to its automorphism group.

By the quotient graph $G\backslash X$ we mean the graph $(G\backslash X, G\backslash V, G\backslash E, \bar{\iota}, \bar{\tau})$ where $\bar{\iota}(Ge) = G\iota e$, $\bar{\tau}(Ge) = G\tau e$ for all $Ge \in G\backslash E$; it is straightforward to see that $\bar{\iota}, \bar{\tau}$ are well-defined. There is then a graph map $X \to G\backslash X, x \mapsto Gx$.

The *Cayley graph* of G with respect to a subset S of G, denoted $X(G, S)$, is the G-graph with vertex set G, edge set $G \times S$, and incidence functions $\iota(g, s) = g, \tau(g, s) = gs$ for all $(g, s) \in G \times S$. This is a G-free G-graph. ∎

2.2 Examples. (i) If $G = \langle s | s^4 \rangle = C_4$, $S = \{s\}$, then

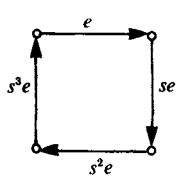

is a geometric realization of $X = X(G, S)$, where $e = (1, s) \in G \times S = EX$. The quotient graph is

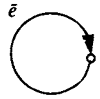

which lifts back to a G-transversal

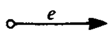

Notice this is not a subgraph, since the terminal vertex is absent.

(ii) If $G = \langle s | \varnothing \rangle = F_1 = C_\infty$, and $S = \{s\}$, then

indicates a geometric realization of $X = X(G, S)$ homeomorphic to $\mathbb{R}$, with $e = (1, s) \in EX$. The quotient graph and G-transversal are as in (i).

(iii) If $G = \langle r, s | r^2, s^2, (rs)^4 \rangle = D_4$, and $S = \{r, s\}$ then

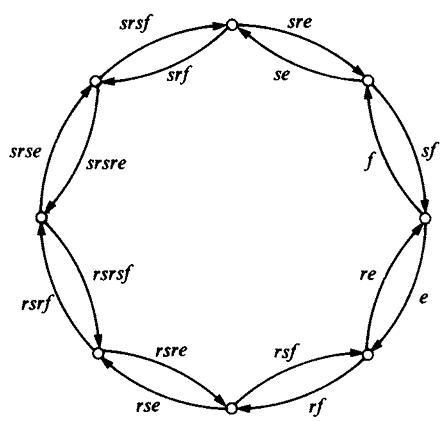

is a geometric realization of $X = X(G, S)$ where $e = (1, r)$, $f = (1, s) \in G \times S = EX$. The quotient graph is

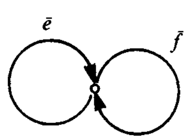

which lifts back to a G-transversal

(iv) If $G = \langle s, r | \varnothing \rangle = F_2$, and $S = \{s, r\}$, then Fig. I.1 indicates a geometric realization of $X = X(G, S)$ omitting the arrows. The quotient graph and G-transversal are essentially as in (iii).

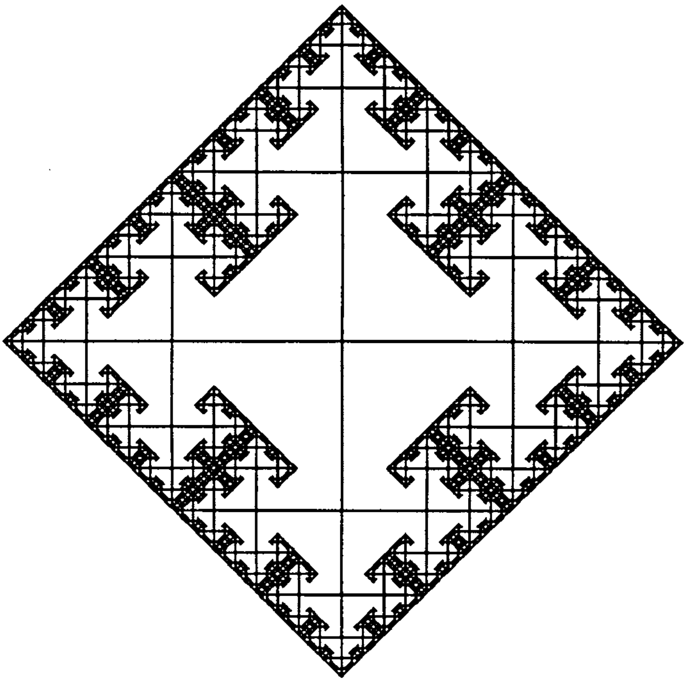

Fig. I.1

(v) If $G = \langle r, s \mid r^2, s^2, (rs)^4 \rangle = D_4$, then

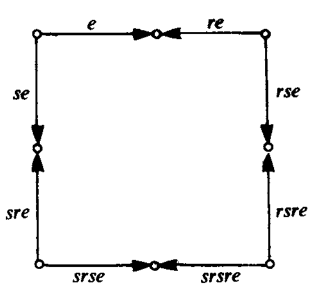

is a geometric realization of a G-graph. The quotient graph is

which lifts back to a G-transversal

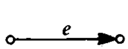

(vi) If $G = \langle r, s \mid r^2, s^2 \rangle = D_\infty$, then

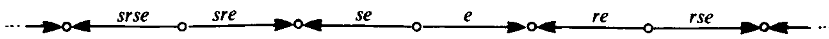

indicates a geometric realization of a G-graph homeomorphic to $\mathbb{R}$; the quotient graph and G-transversal are as in (v). ∎

2.3 Definitions. Let X be a graph.

More incidence functions, again denoted ι, τ, are defined on $EX^{\pm 1}$ by setting $\iota e^1 = \iota e, \tau e^1 = \tau e$, and $\iota e^{-1} = \tau e, \tau e^{-1} = \iota e$ for all $e \in EX$. We think of e^1, e^{-1} as travelling along e the right way and the wrong way, respectively.

A *path* p in X is a finite sequence

(1) $\quad v_0, e_1^{\varepsilon_1}, v_1, \ldots, v_{n-1}, e_n^{\varepsilon_n}, v_n,$

where

$$n \geqslant 0,$$
$$v_i \in VX \quad \text{for each } i \in [0, n],$$
$$e_i^{\varepsilon_i} \in EX^{\pm 1}, \iota e_i^{\varepsilon_i} = v_{i-1}, \tau e_i^{\varepsilon_i} = v_i \quad \text{for each } i \in [1, n].$$

Incidence functions, still denoted ι, τ, are defined on the set of paths in X by setting $\iota p = v_0, \tau p = v_n$; p is said to be a path of *length* n *from* v_0 *to* v_n, and $v_0, \ldots, v_n, e_1, \ldots, e_n, e_1^{\varepsilon_1}, \ldots, e_n^{\varepsilon_n}$ are said to *occur* in p.

It is customary to abbreviate p to $e_1^{\varepsilon_1}, \ldots, e_n^{\varepsilon_n}$. If $n = 0$ then p is said to be *empty*, and here we must specify v_0; if $n \geqslant 1$ the vertices can be recovered from the abbreviated data.

The *inverse* of p, denoted p^{-1}, is the path $v_n, e_n^{-\varepsilon_n}, v_{n-1}, \ldots, v_1, e_1^{-\varepsilon_1}, v_0$. If q is a path with $\iota q = \tau p$ then in an obvious way we can form a path by *concatenation*, denoted p, q.

If for each $i \in [1, n-1]$, $e_{i+1}^{\varepsilon_{i+1}} \neq e_i^{-\varepsilon_i}$ then p is said to be *reduced*. Notice that if $e_{i+1}^{\varepsilon_{i+1}} = e_i^{-\varepsilon_i}$ for some $i \in [1, n-1]$ then $e_1^{\varepsilon_1}, \ldots, e_{i-1}^{\varepsilon_{i-1}}, e_{i+2}^{\varepsilon_{i+2}}, \ldots, e_n^{\varepsilon_n}$ is a path of length $n-2$ from v_0 to v_n.

We say X is a *tree* if for any vertices v, w of X there is a unique reduced path from v to w; this path is then called the X-*geodesic* from v to w. The length of the geodesic is called the *distance* between v and w. For any subset W of V, by the subtree of X *generated by* W we mean the subgraph of X consisting of all edges and vertices which occur in the X-geodesics between the pairs of elements of W.

A subgraph of X which is a tree is called a *subtree* of X.

A path p is said to be a *closed* path *at* a vertex v if $\iota p = \tau p = v$, and is said to be a *simple closed* path if it is nonempty and there are no other repetitions of vertices. Clearly such a path is reduced, and conversely, any reduced closed path is a (possibly empty) sequence of simple closed paths. A graph with no simple closed paths is called a *forest*; equivalently, the only reduced closed paths are the empty ones.

Two elements of X are said to be *connected* in X if there exists a path in X in which they both occur; in this event there is a reduced path in which they both occur. It is straightforward to show that being connected in X is an equivalence relation. The equivalence classes of this relation are called the *components* of X, and they are subgraphs of X. A graph with only one component is said to be *connected*. On VX the relation of being connected in X is the equivalence relation generated by $\{\{\iota e, \tau e\} | e \in EX\}$.

Let E' be a set of edges of X. Write $\bar{E}$ for $EX - E'$ and $\bar{V}$ for the set of components of the graph $X - \bar{E}$ obtained from X by removing $\bar{E}$. There is a natural map $V \to \bar{V}, v \mapsto \bar{v}$, and one can think of $\bar{v}$ as the equivalence class of v relative to the equivalence relation on V generated by $\{(\iota e, \tau e) | e \in E'\}$. Let $\bar{X}$ be the graph with vertex set $\bar{V}$, edge set $\bar{E}$ and incidence functions $\bar{\iota}, \bar{\tau}$ with $\bar{\iota} e = \overline{\iota e}, \bar{\tau} e = \overline{\tau e}$ for all $e \in \bar{E}$. There is a map $X \to \bar{X}, x \mapsto \bar{x}$, which on V is as above, on $\bar{E}$ is the identity, and on E' sends e to $\overline{\iota e} = \overline{\tau e}$; this is not a graph map unless E' is empty. We say $\bar{X}$ is the graph obtained from X by *contracting* all the edges in E', and call $X \to \bar{X}$ the *contracting map*.

For example, if $E' = EX$ then $\bar{E} = \varnothing$ and $\bar{X}$ is a graph with no edges, and vertex set the set of components of X. This provides terminology which is frequently useful for seeing that a graph is connected.

If X is a G-graph and E' is a G-subset of EX then $\bar{X}$ is a G-graph and $X \to \bar{X}$ is a G-map. ∎

2.4 Example. Let S be a subset of G and $X = X(G, S)$.

Let $\bar{X}$ be the graph obtained by contracting all the edges of X, and let $X \to \bar{X}, x \mapsto \bar{x}$, be the contracting map. Then $\bar{X}$ is the G-set with one generator $v = \bar{1}$ and relations $gv = gsv$ for all $(g, s) \in G \times S$, that is, $sv = v$ for all $s \in S$. Hence, G_v is the subgroup of G generated by S, and the components of X correspond to cosets $gG_v \in G/G_v$. Thus X is connected if and only if S generates G.

It will be shown in Theorem 8.2 that X is a tree if and only if S freely generates G; see Examples 2.2(ii), (iv). ∎

2.5 Proposition. *A graph is a tree if and only if it is a connected forest.*

Proof. Let X be a tree. Clearly X is connected. Suppose X has a simple closed path p at some vertex v. Then p and the empty path at v are distinct reduced paths in X from v to itself, which contradicts uniqueness. Hence X is a forest.

Conversely, suppose that X is a connected forest. Let v, w be vertices of X. Since X is connected there is a reduced path from v to w, and it remains to show uniqueness. Suppose that $p = e_1^{\varepsilon_1}, \ldots, e_n^{\varepsilon_n}$ and $q = f_1^{\eta_1}, \ldots, f_m^{\eta_m}$ are reduced paths from v to w. Then $p, q^{-1} = e_1^{\varepsilon_1}, \ldots, e_n^{\varepsilon_n}, f_m^{-\eta_m}, \ldots, f_1^{-\eta_1}$ is a closed path at v. If p, q^{-1} is reduced then it must be empty so clearly $p = q$. If p, q^{-1} is not reduced then $n \geqslant 1, m \geqslant 1$ and $e_n^{\varepsilon_n} = f_m^{\eta_m}$. Here $e_1^{\varepsilon_1}, \ldots, e_{n-1}^{\varepsilon_{n-1}}$ and $f_1^{\eta_1}, \ldots, f_{m-1}^{\eta_{m-1}}$ are reduced paths from v to $\tau e_{n-1}^{\varepsilon_{n-1}}$; by induction on n, these paths are equal. Thus $p = q$ as desired. ∎

We now verify the existence of a very important type of transversal already illustrated in Example 2.2.

2.6 Proposition. *If X is a G-graph and $G\backslash X$ is connected then there exist subsets $Y_0 \subseteq Y \subseteq X$ such that Y is a G-transversal in X, Y_0 is a subtree of X, $VY = VY_0$ and for each $e \in EY, \iota e \in VY = VY_0$.*

We say Y is a *fundamental G-transversal in X, with subtree Y_0*.

Proof. Write $\bar{X} = G\backslash X$ and $\bar{x} = Gx$ for all $x \in X$.

Choose a vertex v_0 of X. By Zorn's Lemma we can choose a maximal subtree Y_0 of X containing v_0 such that the composite $Y_0 \subseteq X \to \bar{X}$ is injective. Let $\bar{Y}_0$ denote the image of Y_0. We claim that $V\bar{Y}_0 = V\bar{X}$. If not, since $\bar{X}$ is connected, any vertex in $\bar{Y}_0$ is connected to any vertex in $\bar{X} - \bar{Y}_0$ by a path in $\bar{X}$, so some edge $\bar{e}$ of $\bar{X}$ has one vertex $\bar{v}$ in $\bar{Y}_0$ and one vertex in $\bar{X} - \bar{Y}_0$. Here $\bar{v}$ comes from an element v of VY_0 and $\bar{e}$ from an edge e of X; since v lies in the same orbit as a vertex of e, it is a vertex of ge for some $g \in G$, and by replacing e with ge we may further assume that v is a vertex of e. Let w be the other vertex of e. Notice e, w do not lie in Y_0, since their images do not lie in $\bar{Y}_0$. But $Y_0 \cup \{e, w\}$ contradicts the maximality of Y_0. This proves the claim that $V\bar{Y}_0 = V\bar{X}$.

For each edge $\bar{e}$ in $E\bar{X} - E\bar{Y}_0$, $\iota\bar{e}$ comes from a unique vertex of Y_0, and as before we can assume $\iota e \in Y_0$. Adjoining the resulting edges to Y_0 gives a subset Y of X such that the composite $Y \subseteq X \to \bar{X}$ is bijective and if $e \in EY$ then $\iota e \in Y$. ∎