Werk

Titel: Trees, valuations, and the Bieri-Neumann-Strebel invariant.

Autor: Brown, K.S.

Jahr: 1987

PURL: https://resolver.sub.uni-goettingen.de/purl?356556735_0090|log28

Kontakt/Contact

Digizeitschriften e.V.
SUB Göttingen
Platz der Göttinger Sieben 1
37073 Göttingen

info@digizeitschriften.de

Invent. math. 90, 479–504 (1987)

Inventiones mathematicae
© Springer-Verlag 1987

Trees, valuations, and the Bieri-Neumann-Strebel invariant

Kenneth S. Brown*
Department of Mathematics, Cornell University, Ithaca, NY 14853, USA

In this paper I will introduce and study *HNN valuations* on groups. These are functions $v: G \to \mathbf{R} \cup \{+\infty\}$ which satisfy axioms resembling those for valuations on rings, and which are also related to decompositions of G as an HNN extension. They arise naturally when one classifies G-$\mathbf{R}$-trees whose hyperbolic length function is abelian.

My main interest in HNN valuations is that they can be used to characterize the geometric invariant $\Sigma = \Sigma(G)$ defined by Bieri, Neumann, and Strebel [7] if G is finitely generated. [Recall that Σ is a certain set of equivalence classes $[\chi]$ of non-zero homomorphisms $\chi: G \to \mathbf{R}$, where two such homomorphisms are equivalent if they are positive scalar multiples of one another. It captures, among other things, complete information as to which normal subgroups of G with abelian quotient are finitely generated.] An HNN valuation v on G gives rise to a homomorphism $\chi_v: G \to \mathbf{R}$, and I will show that Σ is the set of classes $[\chi]$ such that χ does *not* have the form χ_v for any non-trivial HNN valuation v.

This description of Σ, as we will see, unifies two previously known results: On the one hand, the "rational points" of Σ were understood in terms of *HNN* decompositions of G ([7], § 4); on the other hand, Σ was undertstood for metabelian G in terms of *valuations* on commutative rings [6, 3].

The point of view provided by HNN valuations seems to be quite useful for computational purposes, in the same way that ring-theoretic valuations were useful in the metabelian case. I will illustrate this by calculating Σ when G is a one-relator group $\langle \mathfrak{X}; r \rangle$. As a corollary, one can quickly read off from the defining relator r a description of the set of finitely generated normal subgroups of G with abelian quotient.

The paper is organized as follows. I begin by studying in § 1 G-$\mathbf{R}$-trees with abelian length function. This motivates the definition of HNN valuation, which is given in § 2. It is possible to read § 2 independently of § 1, but the definition may then seem strange. In § 3 I look at the connection between HNN valuations

* Partially supported by NSF grant DMS-8502278

and HNN extensions. This discussion leads naturally to a *definition* in terms of HNN valuations of a set that I call Σ.

§ 4 treats one-relator groups. I calculate Σ (as defined in § 3), and I give the algorithm for describing the finitely generated normal subgroups with abelian quotient. For the convenience of the reader who is interested in one-relator groups but is not familiar with the Bieri-Neumann-Strebel invariant, I have written § 4 so as to be independent of [7].

In § 5 I do assume familiarity with [7], and I prove that the set Σ defined in § 3 coincides with the Bieri-Neumann-Strebel Σ; this gives the promised characterization of the latter in terms of HNN valuations. More generally, § 5 contains a characterization of the sets Σ_A of [7], where A is a G-group. I indicate briefly in § 6 how one can recover from this characterization the known connection between Σ_A and ordinary valuations (on rings) when G and A are abelian. In § 7 I complete the discussion of trees begun in § 1 by constructing a G-$\mathbf{R}$-tree associated to any HNN valuation. This, together with §§ 1 and 5, yields a characterization of Σ (Corollary 7.4) in terms of G-$\mathbf{R}$-trees. Finally, § 8 contains two families of examples where Σ is computed via tree actions.

Some of the results of this paper remain valid if $\mathbf{R}$ is replaced by an arbitrary ordered abelian group A. Treating this generalization systematically, however, would have resulted in extra technicalities without adding any substance to the paper. I have therefore confined myself to the case $A=\mathbf{R}$, but, wherever feasible, have written the exposition in such a way that it applies with little change to a general A.

I am very grateful to R. Bieri and R. Strebel, who have had an enormous influence on this paper. In particular, the idea that Σ should be describable in terms of tree actions arose from discussions with Bieri at a conference in Bielefeld in November 1985. And Strebel, upon seeing the functions v that arose from my study of tree actions, urged me to take seriously the analogy with valuations on rings and suggested that this point of view might lead to a calculation of Σ for one-relator groups. In addition, he showed me how to reformulate the result of this calculation in geometric language.

Acknowledgements. I would also like to thank H. Abels for organizing the Bielefeld conference and providing a stimulating atmosphere, and I thank the University of Bielefeld and the National Science Foundation for their financial support.

1. Tree actions with abelian length function

Useful references for this section are [1], [10], and [15]. We will follow the notation and terminology of [1] regarding $\mathbf{R}$-trees and group actions on $\mathbf{R}$-trees, except that all of our actions will be *right* actions.

Let X be a G-$\mathbf{R}$-tree, where G is an arbitrary group, and let $l\colon G\to\mathbf{R}$ be the corresponding hyperbolic length function. Recall that l is said to be *abelian* if $l=|\chi|$ for some homomorphism $\chi\colon G\to\mathbf{R}$; χ is then unique up to a factor of ± 1 ([1], 1.4). If, in addition, G has no fixed points in X and no invariant line, then we will say that X is a *non-trivial* abelian G-$\mathbf{R}$-tree. In this case G fixes a unique end e of X ([1], Theorem 7.5), which can be characterized as

the unique end belonging to all of the characteristic subtrees A_g for $g \in G$. [Recall that A_g is the axis of g if g is hyperbolic and the fixed point set of g if g is elliptic.]

We use the fixed end e, in the non-trivial case, to make a canonical choice of the homomorphism χ above: For any hyperbolic g, we set $\chi(g) = l(g)$ if g translates *away* from e, and $\chi(g) = -l(g)$ otherwise. In other words, if we choose an identification of A_g with $\mathbf{R}$ such that e corresponds to $-\infty$, then $x \cdot g = x + \chi(g)$ for all $x \in A_g$. This canonical χ will be denoted χ_X. Following a method of Tits (cf. [15], last paragraph of proof of Proposition 4), we will now introduce a second function on G, called v, which essentially determines the G-$\mathbf{R}$-tree X.

Still assuming that X is a non-trivial abelian G-$\mathbf{R}$-tree with fixed end e, choose a basepoint $x \in X$, and let Y be the ray $(e, x]$. Let $\iota\colon Y \to \mathbf{R}$ be the unique isometric embedding such that $\iota(y)$ decreases as $y \to e$ and $\iota(x) = 0$; the image $\iota(Y)$ is then an interval $(r, 0]$, where $r \geqq -\infty$. [The definition of "end" given in [1] allows the possibility $r > -\infty$, but this can only happen if $\chi = 0$. The reader may prefer to redefine "non-trivial" so as to exclude this case.] We define a function $v = v_{X,x}\colon G \to (-\infty, 0)$ as follows. For any $g \in G$, the ray Yg^{-1} represents the same end e as Y, hence $Y \cap Yg^{-1}$ is a ray $(e, y]$ for some $y \in Y$. Now set $v(g) = \iota(y)$. In other words, if we view ι as an identification and think of Y as a set of real numbers, then

$$v(g) = \max\{t \in Y\colon\ t \cdot g \in Y\}. \tag{1.1}$$

Letting $\chi = \chi_X$, we can also write

$$t \cdot g = t + \chi(g) \quad \text{for } t \leqq v(g) \text{ in } Y, \tag{1.1'}$$

and $v(g)$ is the largest non-positive real number with this property.

It is useful to give an alternate definition of v in terms of the subtrees A_g. For this purpose note that, since the end e is common to Y and A_g, the intersection $Y \cap A_g$ is a ray $(e, z]$ for some $z \in Y$; we then have $v(g) = \iota(z)$ if $\chi(g) \leqq 0$ and $v(g) = \iota(z) - \chi(g)$ if $\chi(g) \geqq 0$. The following picture illustrates this in case $\chi(g) > 0$. Here y and z are as above, and we have $z = yg$. The arrows indicate the direction of translation of g along its axis.

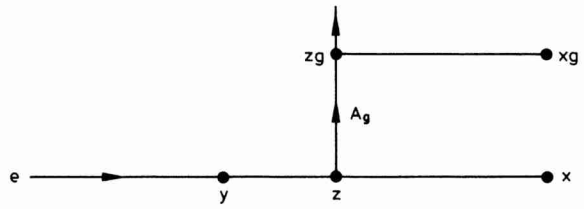

Fig. 1

Letting d be the metric on X, we see from this picture (and similar ones for $\chi(g) \leqq 0$) that

$$v(g) = -d(x, A_g) - \chi(g) \quad \text{if } \chi(g) \geqq 0 \tag{1.2}$$

and

$$v(g) = -d(x, A_g) = v(g^{-1}) + l(g) = v(g^{-1}) - \chi(g) \quad \text{if } \chi(g) \leqq 0. \tag{1.3}$$

Let L be the Lyndon length function associated to the base point x, i.e., $L(g) = d(x, xg)$. Then we have

$$L(g) = l(g) + 2d(x, A_g) = -\chi(g) - 2v(g) \tag{1.4}$$

for all g, where the second equality follows from 1.2 and 1.3. This shows that the (based) G-$\mathbf{R}$-tree X is uniquely determined by χ and v, provided X is spanned by the orbit of x (cf. [1], § 5). Note that χ can be eliminated here. Indeed, it follows from 1.3 that

$$v(g^{-1}) = v(g) + \chi(g) \tag{1.5}$$

for all g, hence $\chi(g) = v(g^{-1}) - v(g)$. Substituting this into 1.4, we obtain

$$L(g) = -(v(g) + v(g^{-1})), \tag{1.6}$$

so the G-$\mathbf{R}$-tree spanned by x is in fact determined by v alone.

1.5 is one of the formal properties of v that we will take as part of the definition of "HNN valuation" in § 2. Before proceeding further, we record two others. The first is essentially a restatement in terms of v of our non-triviality assumption. Let $G_{\chi \leqq 0} = \{g \in G : \chi(g) \leqq 0\}$. In the terminology of [15], this is simply the submonoid of G consisting of the elements g such that the end e is attracting or neutral for g.

(1.7) $v | G_{\chi \leqq 0}$ does not assume a minimum value.

For if $v | G_{\chi \leqq 0}$ assumed a minimum value, then the intersection $Z = \bigcap_{g \in G} A_g$ would be non-empty. But Z would then be a G-invariant subtree of X, which would consist of fixed points if $\chi = 0$ and which would be a line otherwise; either way we contradict the non-triviality assumption on X.

The other property of v is the following inequality, which resembles the familiar ultrametric inequality from valuation theory:

$$v(gh) \geqq \min\{v(g), v(h) - \chi(g)\}. \tag{1.8}$$

To prove this, suppose that $t \leqq \min\{v(g), v(h) - \chi(g)\}$ in Y. By 1.1′, we have $tg = t + \chi(g) \leqq v(h)$, so $tgh = t + \chi(g) + \chi(h) = t + \chi(gh)$. All of this is going on inside Y, so 1.8 follows from the definition 1.1 of v.

We also have (as in valuation theory again):

(1.9) Equality holds in 1.8 if $v(g) \neq v(h) - \chi(g)$.

One can see this directly, as in the proof of 1.8, or one can deduce it formally from 1.5 and 1.8.

We close this section by describing a variant of our definition of v. Instead of taking values in $(-\infty, 0]$, the new function will be allowed to take arbitrary real values, as well as the value $+\infty$. For the rest of this section we will denote by v_0 (resp. Y_0) the function called v above (resp. the ray $(e, x]$ called Y above).

Choose, if possible, a subtree Y containing Y_0 such that ι extends to an isometry (still called ι) of Y onto an interval $(r, +\infty)$ of real numbers, $r \geqq -\infty$. [If $\chi \neq 0$, then $r = -\infty$, and we have simply extended the half-line Y_0 to a line Y.] For example, if x is on the axis A_g of some hyperbolic element g, then we can take $Y = A_g$. As in 1.1 and the discussion preceding it, one sees for any $g \in G$ that either there is a largest $y \in Y$ such that $yg \in Y$, in which case we set $v(g) = \iota(y)$, or else Y is g-invariant [which means $Y = A_g$ if g is hyperbolic], in which case we set $v(g) = \infty$. Thus 1.1 is still valid, provided "max" is replaced by "sup". This function $v = v_{X,Y,x}: G \to R \cup \{+\infty\}$ has the same formal properties 1.5, 1.7, 1.8, and 1.9 as v_0. Note, however, that we now have $v(1) = +\infty$, whereas v_0 satisfies $v_0(1) = 0$.

The following consequence of 1.5 and 1.9 shows that χ can be recovered from v:

(1.10) If $v(g) < \infty$ then $\chi(g) = v(g^{-1}) - v(g)$. If $v(g) = \infty$ then $\chi(g) = v(h) - v(gh)$ for any h with $v(h) < \infty$.

(Note that there must exist such an h by 1.7.)

It now follows easily that this "extended" v, like the original v_0, contains enough information to determine the subtree spanned by the orbit of x. To see this, it suffices to show that we can recover v_0 from v. Now we can certainly recover v_0 from v and χ. For by 1.5 we need only consider those g with $\chi(g) \leqq 0$, and we clearly have:

(1.11) If $\chi(g) \leqq 0$ then $v_0(g) = \min\{v(g), 0\}$.

So our claim follows from 1.10.

Examples. 1. Suppose G is an HNN extension $\langle B, t; (B_1)^t = B_2 \rangle$ and let X be the associated tree, cf. [14]. The stable letter t is hyperbolic in X, and the vertex stabilizers along the axis A_t are the conjugates $N_r = t^{-r} B t^r$, $r \in \mathbf{Z}$. These form an increasing (resp. decreasing) sequence as r increases if and only if the HNN extension is ascending (resp. descending)[1]. It follows easily that the action of G on X is abelian and non-trivial if and only if the HNN extension is either properly ascending (i.e., ascending but not descending) or properly descending (i.e., descending but not ascending).

Suppose, for instance, that it is properly descending. Set $Y = A_t$ and take as basepoint the unique vertex x of Y whose stabilizer is B. The resulting $v = v_{X,Y,x}$ then satisfies $v(t) = \infty$ and $v(n) = \sup\{r \in \mathbf{Z}: n \in N_r\}$ for $n \in N = \bigcup N_r$. These equations determine v, for any $g \in G$ can be written as $g = nt^k$ for some $n \in N$ and $k \in \mathbf{Z}$, and we have $v(nt^k) = v(n)$ by 1.8 and 1.9. The homomorphism χ_X in this case is the *canonical homomorphism* χ associated to the HNN extension, with $X(B) = 0$ and $\chi(t) = 1$. [In the ascending case, on the other hand, $\chi_X = -\chi$.]

2. Let k be a field with a proper non-archimedean valuation $\omega: k \to \mathbf{R} \cup \{+\infty\}$. ("Proper" means that ω takes on values other than 0 and $+\infty$.) Let X be

[1] Recall that G is called an *ascending* (resp. *descending*) HNN extension if $B_2 = B$ (resp. $B_1 = B$)

the corresponding **R**-tree (cf. [1] or [15] or, for the case of a discrete valuation, [14]). It admits an action of $GL_2(k)$. Let G, Q, and A be the following subgroups of $GL_2(k)$:

$$G=\begin{pmatrix}1 & *\\ 0 & *\end{pmatrix} \quad Q=\begin{pmatrix}1 & 0\\ 0 & *\end{pmatrix} \quad A=\begin{pmatrix}1 & *\\ 0 & 1\end{pmatrix}$$

Thus Q can be identified with the multiplicative group k^*, A can be identified with the additive group k, and G is their semi-direct product $A \rtimes Q$. The action of G on X is abelian and non-trivial. One can check that there is a (unique) Q-invariant line Y in X, which we use to define $v: G \to \mathbf{R} \cup \{+\infty\}$ (after choosing a basepoint x). Then $v(Q)=\infty$, and it is possible to choose x so that $v|A=\omega$. Hence $v(aq)=v(a)=\omega(a)$ for $a \in A$ and $q \in Q$. The homomorphism $\chi=\chi_X: G \to \mathbf{R}$ is given by $\chi(aq)=\omega(q)$.

2. HNN valuations

Let G be a group and $\chi: G \to \mathbf{R}$ a homomorphism. Let $\mathbf{R}_\infty$ be the ordered commutative monoid $\mathbf{R} \cup \{+\infty\}$, with $+\infty$ as largest element and $r+(+\infty) = +\infty$ for all r. A function $v: G \to \mathbf{R}_\infty$ will be called an HNN *valuation with associated homomorphism* χ if it satisfies the following two axioms:

(a) $v(g^{-1})=v(g)+\chi(g)$.

(b) $v(gh) \geqq \min\{v(g), v(h)-\chi(g)\}$.

We will say that v is *non-trivial* if, in addition:

(c) $v|G_{\chi \leqq 0}$ does not assume a minimum value.

Axioms (a) and (b) are simply properties 1.5 and 1.8 from the previous section, and (c) is 1.7. As in § 1, (a) and (b) imply the following two properties:

(d) Equality holds in (b) if $v(g) \neq v(h)-\chi(g)$.

(e) The homomorphism χ is uniquely determined by v, except in the case where $v(g)=\infty$ for all g.

In view of (e) we may write $\chi=\chi_v$ if $v(g)<\infty$ for some g, e.g., if v is non-trivial.

It follows from (a) and (b) that $v(1)$ is the largest value taken on by v. In particular, $v(1)=\infty$ if v takes on the value ∞ at all. In case $v(1)<\infty$, we can normalize v by adding a constant to make $v(1)=0$, and then the values of v are in $(-\infty, 0]$, as with the function $v_0=v_{\chi,x}$ of § 1.

Note that we allow the possibility that $\chi=0$ in the definition above. But if G is finitely generated and v is non-trivial, then χ is necessarily non-zero; for if χ were zero, then the minimum value of v on the generators would be the minimum value of v on all of G, contrary to (c).

We now record a few easy consequences of the axioms. Assume, in what follows, that (a) and (b) hold.

(2.1) Proposition. (i) *Given* $g_1, \ldots, g_n$ *in* G, *let* $v_i=v(g_i)-\chi(g_1 \cdots g_{i-1})$, $i=1, \ldots, n$. *Then* $v(g_1 \cdots g_n) \geqq \min\{v_1, \ldots, v_n\}$, *with equality if the sequence* (v_i) *assumes its minimum value only once.*

(ii) *If* $\chi(g)<0$ *then* $v(g^k)=v(g)$ *for any positive integer* k.

(iii) *If* $v(g) \geqq v(h)$ *and* $\chi(h) \leqq 0$, *then* $v(h^g) \geqq v(h)+\chi(g)$, *with equality if* $v(g) > v(h)$.

(iv) *Given* $g, h \in G_{\chi \leqq 0}$ *with* $v(g) \geqq v(h)$, *we have* $v([g,h]) \geqq v(h) + \chi(gh)$, *where* $[g,h] = g^{-1}h^{-1}gh$. *Equality holds if* $v(g) > v(h)$ *and* $\chi(g) < 0$.

(v) *Let* $N = \ker \chi$, *and for any real number* $r \leqq v(1)$ *let* $N_r = \{n \in N\colon v(n) \geqq r\}$. *Then* N_r *is a subgroup of* N, $N_r \supseteq N_s$ *if* $r \leqq s$, *and* $(N_r)^g = N_{r+\chi(g)}$ *for* $r \leqq v(g)$. □

These subgroups N_r have a simple interpretation when v comes from a tree action as in § 1. Assuming, for simplicity, that $\chi \neq 0$, the N_r are simply the stabilizers along a line or half-line leading to the fixed end.

Remark. A comparison of (v) with Example 1 of § 1 suggests that we should think of an HNN valuation v as giving G the structure of "generalized descending HNN extension".

The following consequence of 2.1(i) will be the starting point for our study of one-relator groups in § 4:

(2.2) Corollary. *Suppose* $g_1 \ldots g_n = 1$, *where* $n \geqq 2$. *Then the sequence* (v_i) *of* 2.1(*i*) *assumes its minimum value at least twice.*

Proof. Suppose the minimum μ occurs only once. Then $v(1) = \mu$ by 2.1(i). If $v(1) = \infty$, then we already have a contradiction, since the minimum cannot be ∞ if it only occurs once. If $v(1) < \infty$, then consider the cyclic permutation $g_n \cdot g_1 \cdots g_{n-1}$ of the original relator. Its v-sequence, up to order, is $(v_i - \chi(g_n))$, whose minimum $\mu - \chi(g_n)$ still occurs only once and still must equal $v(1)$ by 2.1(i); hence $\chi(g_n) = 0$. Continuing in this way we see that $\chi(g_i) = 0$ for all i, so that $v_i = v(g_i)$. But now it is clear that the minimum cannot be $v(1)$ if it occurs only once, since $v(1)$ is the largest value taken on by v. □

Next we wish to elaborate on the non-triviality condition (c).

(2.3) Proposition. *Suppose* v *and* χ *satisfy* (*a*) *and* (*b*), *with* $\chi \neq 0$. *Then* (*c*) *is equivalent to each of the following conditions*:

(*c'*) $v|N$ *is not bounded below, where* $N = \ker \chi$.

(*c''*) $v|G'$ *is not bounded below, where* G' *is the commutator subgroup of* G.

Proof. It is trivial that (c'')⇒(c')⇒(c). Suppose (c) holds, and choose $g \in G$ with $\chi(g) < 0$. By hypothesis we can find $h \in G_{\chi \leqq 0}$ with $v(h) < v(g)$. Conjugating h by positive powers of g, we obtain by 2.1(iii) a sequence (h_i) in $G_{\chi \leqq 0}$ such that $v(h_i)$ decreases to $-\infty$. 2.1(iv) now implies that $v([g,h_i]) = v(h_i) + \chi(gh_i) < v(h_i)$, so $v([g,h_i]) \to -\infty$. Hence (c)⇒(c''). □

Examples. 1. Suppose G is a descending HNN extension $\langle B, t;\ B^t = B_2 \rangle$. Let $\chi\colon G \twoheadrightarrow \mathbf{Z} \subset \mathbf{R}$ be the canonical homomorphism, with $\chi(B) = 0$ and $\chi(t) = 1$. Then there is an HNN valuation v on G with $\chi = \chi_v$. We have already seen this in Example 1 of § 1 from the tree point of view, but it can easily be verified directly; just define v by the formula given in that example and check that the axioms hold. This v is non-trivial if and only if the given HNN extension is *properly* descending.

Conversely, suppose v is a non-trivial HNN valuation such that $\chi = \chi_v$ is *discrete*, i.e., has infinite cyclic image. Multiplying v and χ by a positive scalar, we may assume $\chi(G) = \mathbf{Z}$. Choose $t \in G$ with $\chi(t) = 1$, and let $B = N_r$ for any $r \leqq v(t)$. Then the conjugates $t^k B t^{-k} (k \geqq 0)$ are proper subgroups of N which increase

and exhaust N by 2.1(v). Hence G is a properly descending HNN extension and χ is the associated homomorphism.

We can summarize this example as follows: *Given* $\chi: G \twoheadrightarrow \mathbf{Z}$, *there is a non-trivial HNN valuation* v *with* $\chi = \chi_v$ *if and only if* G *admits a decomposition as a properly descending HNN extension with* χ *as associated homomorphism.*

2. Let A be a ring (not necessarily commutative), let Q be a subgroup of the multiplicative group A^* of units of A, and let $\omega: A \to \mathbf{R}_\infty$ be a valuation in the sense of [8], § VI.3.1. Let $G = A \rtimes Q$, where Q acts on the additive group A by *right* multiplication; thus $(a, q)\cdot(a', q') = (a + a'q^{-1}, qq')$. Motivated by Example 2 of § 1, we set $\chi(a, q) = \omega(q)$ and $v(a, q) = \omega(a)$. It is easy to check that χ takes finite values and is a homomorphism and that axioms (a) and (b) hold. Condition (c) will hold if and only if ω does not assume a minimum value; a sufficient condition for this is that $\omega | Q \neq 0$.

[For future reference, we remark that everything we have said here remains valid if ω only satisfies the following weak form of the axioms for a valuation on a ring:

(i) $\omega(0) = \infty$ and $\omega(1) = 0$.

(ii) $\omega(a + b) \geqq \min\{\omega(a), \omega(b)\}$.

(iii) $\omega(-a) = \omega(a)$ and $\omega(aq) = \omega(a) + \omega(q)$ for $a \in A$, $q \in Q$.

To be a valuation, of course, ω would have to satisfy a stronger version of (iii), namely, $\omega(ab) = \omega(a) + \omega(b)$ for all $a, b \in A$.]

3. If G is abelian, then G does not admit any non-trivial HNN valuation, except possibly with $\chi = 0$. This follows from 2.3.

4. At the other extreme, suppose G is a non-abelian free group. Then *any* non-zero $\chi: G \to \mathbf{R}$ can occur as the homomorphism associated to a non-trivial HNN valuation. In fact, let $\mathfrak{X}$ be a basis for G, and let $\varphi: \mathfrak{X} \to \mathbf{R}_\infty$ be an arbitrary function. Extend φ to $\mathfrak{X}^{\pm 1} = \mathfrak{X} \cup \mathfrak{X}^{-1}$ by setting $\varphi(x^{-1}) = \varphi(x) + \chi(x)$ for $x \in \mathfrak{X}$. Then 2.1(i) suggests a way to extend φ to a function v on G: For any reduced $\mathfrak{X}^{\pm 1}$-word $\xi = x_1 \ldots x_n$ with $n \geqq 1$, set $v(\xi) = \min\{\varphi(x_i) - \chi(x_1 \ldots x_{i-1}): 1 \leqq i \leqq n\}$; set $v(1)$ equal to any value $\geqq \varphi(x)$ for all $x \in \mathfrak{X}^{\pm 1}$, e.g., $v(1) = \infty$. It is easy to check that axioms (a) and (b) hold. To make sure that (c) holds, we need only be a little careful in choosing φ. For example, suppose we choose distinct elements $x, y \in \mathfrak{X}$ with $\chi(y) \neq 0$; if we then define φ so that $\varphi(x) = 0$ and $\varphi(y) = \infty$, it is easy to check (c).

5. Let G and χ be arbitrary and let $\mathfrak{X}$ be a set of generators of G. [*Note*: We will take this to mean that G is given as a quotient of $F(\mathfrak{X})$; we do not require that $\mathfrak{X}$ embed in G, although, for simplicity, we will not distinguish notationally between an element of $\mathfrak{X}$ and its image in G.] Let $\varphi: \mathfrak{X} \to \mathbf{R}_\infty$ be arbitrary. We can again try to extend φ to an HNN valuation v with $\chi_v = \chi$, as follows: First extend φ to $\mathfrak{X}^{\pm 1}$ as above; then extend it to $\mathfrak{X}^{\pm 1}$-words $\xi = x_1 \ldots x_n$ by setting $\varphi(\xi) = \min\{\varphi(x_i) - \chi(x_1 \ldots x_{i-1})\}$ if $n \geqq 1$ and $\varphi(\xi) = \sup\{\varphi(x): x \in \mathfrak{X}^{\pm 1}\}$ if $n = 0$. Now set $v(g) = \sup\{\varphi(\xi)\}$, where ξ ranges over all words representing g. It is easy to check that (a) and (b) hold. But there is no guarantee that $v | \mathfrak{X} = \varphi$. All we can say is that $v | \mathfrak{X} \geqq \varphi$ and that v is the smallest HNN valuation with this property (and with χ as associated homomorphism). So if there is *any* extension of φ to an HNN valuation with χ as associated homomorphism, then our v extends φ.

6. Let G be the group of orientation-preserving piecewise linear homeomorphisms of the unit interval. Let $\lambda\colon G \to \mathbf{R}_+^*$ be the derivative at 0, where $\mathbf{R}_+^*$ is the multiplicative group of positive reals, and let $\chi = \log \lambda$. Let $v(g) = \log \varepsilon(g)$, where $\varepsilon(g)$ is the largest $\varepsilon > 0$ such that g is linear on $[0, \varepsilon]$. It is easy to check that χ and v satisfy (a)–(c). A similar definition of v works for $\chi = \log \rho$, where ρ is the derivative at 1. We will return to this example in § 8.

We close this secton with some technical remarks that will be useful later. First, since an HNN valuation is determined by its restriction to $G_{\chi \leqq 0}$, it is natural to try to state the axioms in terms of this restricted function. To this end one easily verifies:

(2.4) Proposition. *Fix a homomorphism $\chi\colon G \to \mathbf{R}$ and a function $v\colon G_{\chi \leqq 0} \to \mathbf{R}_\infty$. Then there is a unique extension of v to G satisfying (a), and this extension satisfies (b) if and only if the original v satisfies (b) and (d).* □

Next we introduce two methods for modifying HNN valuations, both of which are motivated by § 1. Let v be an HNN valuation with associated homomorphism χ, and suppose g is an element of G with $\chi(g) \neq 0$ and $v(g) < \infty$. We wish to change v to a new HNN valuation w with the same χ and with $w(g) = \infty$. Replacing g by g^{-1} if necessary, we may assume $\chi(g) < 0$. Given $h \in G_{\chi \leqq 0}$, let $h_k = g^{-k} h g^k$ for $k \geqq 0$. If there is a k such that $v(h_k) < v(g)$, then choose such a k and set $w(h) = v(h_k) - k\chi(g)$; it follows from 2.1(iii) that this is independent of the choice of k. If there is no such k, then set $w(h) = \infty$. It is easy to verify (b) and (d) for w, as defined on $G_{\chi \leqq 0}$, so w extends to an HNN valuation by 2.4. We will say that w is obtained from v by *change of axis*. [In case v comes from a tree action and is defined with the aid of some line or half-line Y, w is simply the HNN valuation obtained by replacing Y by the axis A_g.]

The second construction, called *truncation*, is motivated by 1.11. Given an HNN valuation v and a real number $M \leqq v(1)$, we define a new HNN valuation w with $w(1) = M$ as follows: Define w on $G_{\chi \leqq 0}$ by $w(g) = \min\{v(g), M\}$, check that (b) and (d) hold, and extend w to G by 2.4.

3. Ascending and descending HNN extensions

Let G be a finitely generated group and $\chi\colon G \twoheadrightarrow \mathbf{Z}$ a surjective homomorphism. In many cases one knows that G necessarily admits an HNN decomposition $G = \langle G, t;\ (B_1)^t = B_2 \rangle$ with χ as associated homomorphism and base group B again finitely generated. This holds, for instance, if G is finitely presented ([4], Theorem A). And if G is a one-relator group, we can even take B to be a one-relator group and B_1 and B_2 to be free subgroups. (The procedure for doing this will be recalled at the beginning of the next section.) It is of interest to know whether this HNN extension is *ascending*, i.e., whether $B_2 = B$. In the one-relator case, for instance, B is then free, and χ therefore yields a particularly simple description of G as a split extension $N \rtimes \langle t \rangle$, where $N = \ker \chi$ is an increasing union of free groups. This description is even simpler if the HNN extension is both ascending and descending, in which case N itself is free.

HNN valuations provide a useful point of view for studying this situation, in view of the following result:

(3.1) Proposition. *Let G be a finitely generated group and $\chi: G \twoheadrightarrow \mathbf{Z}$ a surjective homomorphism. The following three conditions are equivalent:*

(i) *G admits a decomposition as an ascending HNN extension with finitely generated base group B and with χ as associated homomorphism.*

(ii) *G does not admit a decomposition as a properly descending HNN extension with χ as associated homomorphism.*

(iii) *There is no non-trivial HNN valuation v on G with $\chi = \chi_v$.*

Moreover, if these conditions hold then every HNN decomposition of G with χ as associated homomorphism is ascending.

The equivalence of (ii) and (iii) has already been pointed out in Example 1 of § 2. The equivalence of (i) and (ii), as well as the last assertion of the proposition, can be found in § 4 of [7]. The proofs given there, however, are phrased in terms of the Σ-invariant. For the convenience of the reader, we will give here a direct proof:

(i)$\Rightarrow$(ii): Suppose G is an ascending HNN extension $\langle B, t;\ (B_1)^t = B\rangle$ with B finitely generated, $B \subseteq N = \ker \chi$, and $\chi(t) = 1$. If G also admits a descending HNN decomposition $\langle C, s;\ C^s = C_2\rangle$ with $C \subseteq N$ and $\chi(s) = 1$, I will show that $C = N$, so that the decomposition is not *properly* descending. Write $t = ns$, where $n \in N$. Since N is the increasing union of the conjugates $C_k = s^k C s^{-k} (k \geqq 0)$, some C_k must contain n and a finite set of generators for B. But then C_k contains B and is closed under conjugation by t, so C_k contains N and hence $C = N$.

(ii)$\Rightarrow$(i): Choose $t \in G$ with $\chi(t) = 1$. Since G is finitely generated, we can find a finitely generated subgroup $B \subseteq N = \ker \chi$ such that G is generated by B and t. Let C be the t-closure of B, i.e., the subgroup generated by the conjugates $B_k = t^{-k} B t^k (k \geqq 0)$. Then G is a descending HNN extension with base group C, stable letter t, and associated homomorphism χ. In view of the hypothesis (ii), this HNN decomposition cannot be properly descending, so $C = N$. In particular, $tBt^{-1} \subseteq C$, hence tBt^{-1} is contained in the subgroup D generated by $B_0, \ldots, B_k$ for some k. But then D is closed under conjugation by t^{-1}, so G is an ascending HNN extension with base group the finitely generated group D, stable letter t, and associated homomorphism χ. Thus (i) holds.

Proof of the last assertion of 3.1. Suppose we have an arbitrary HNN decomposition $G = \langle B, t;\ (B_1)^t = B_2\rangle$ with $B \subseteq N = \ker \chi$ and $\chi(t) = 1$. As in the previous paragraph, if (ii) holds then the t-closure C of B must equal N. But the normal form theorem for HNN extensions easily implies that C cannot equal N unless $B_2 = B$. □

(3.2) Corollary. *Let $\chi: G \twoheadrightarrow \mathbf{Z}$ be a surjection, where G is finitely generated. Then $\ker \chi$ is finitely generated if and only if χ and $-\chi$ both satisfy the conditions of 3.1.*

Proof. The "only if" part is trivial. Conversely, suppose χ and $-\chi$ satisfy the conditions of 3.1, and write G as an ascending HNN extension $\langle B, t;\ (B_1)^t = B\rangle$ with B finitely generated, $B \subseteq \ker \chi$, and $\chi(t) = 1$. Then G is also a descending HNN extension $\langle B, s;\ B^s = B_1\rangle$, where $s = t^{-1}$. The homomorphism associated

to this HNN decomposition is $-\chi$, so 3.1 (applied to $-\chi$) implies that this descending HNN extension is ascending. Thus $B_1 = B = \ker \chi$, and the latter is therefore finitely generated. □

We wish to understand the set of homomorphisms χ satisfying the conditions of the proposition. It will be convenient here to focus on condition (iii) and to allow arbitrary non-zero χ (not just discrete χ) and arbitrary G (not necessarily finitely generated). Since condition (iii) is not affected if χ is multiplied by a positive scalar, we introduce the set $S = S(G) = (\mathrm{Hom}(G, \mathbf{R}) - \{0\})/\mathbf{R}_+^*$, where $\mathbf{R}_+^*$ is the multiplicative group of positive reals, acting by scalar multiplication. When G is finitely generated, S is the sphere considered in [7].

Now let $\Sigma = \Sigma(G)$ be the set of classes $[\chi] \in S$ such that χ satisfies condition (iii) of 3.1. We will show in § 5 that, for finitely generated G, Σ coincides with the set called Σ in [7][2]. First, however, we wish to compute our Σ for one-relator groups. That will be done in the next section. We close the present section with two observations that will be needed there. Let Σ^c be the complement of Σ in S.

(3.3) Proposition. *Let G be a finitely generated group with finite generating set $\mathfrak{X}$, let $\chi\colon G \to \mathbf{R}$ be a non-zero homomorphism, and let t be an element of $\mathfrak{X}$ such that $\chi(t) \neq 0$. The following three conditions are equivalent:*

(i) $[\chi] \in \Sigma^c$.

(ii) *There is an HNN valuation v on G with $\chi_v = \chi$, $v(t) = \infty$, and $v(x) < \infty$ for some $x \in \mathfrak{X}$.*

(iii) *There is an HNN valuation v on G with $\chi_v = \chi$, $v(t) = \infty$, and $\min\{v(x): x \in \mathfrak{X}\} = 0$.*

Proof. An HNN valuation v as in (ii) is easily seen to be non-trivial (consider the conjugates of $x^{\pm 1}$ by the powers of t); so (ii)⇒(i) by our definition of Σ. Conversely, suppose (i) holds, and let v be a non-trivial HNN valuation with $\chi = \chi_v$. By change of axis we may assume $v(t) = \infty$ (cf. end of § 2; note that change of axis leaves χ unchanged and does not affect the non-triviality assumption). Then $v(x) < \infty$ for some $x \in \mathfrak{X}$, since otherwise we would have $v(G) = \infty$, so (ii)⇒(i). Finally, we may normalize any v as in (ii) (by adding a constant) to make $\min\{v(x)\} = 0$, so (ii)⇔(iii). □

(3.4) Proposition. *If G is finitely generated, then Σ is an open subset of the sphere S.*

Proof. Choose a finite set $\mathfrak{X}$ which generates G. Given $[\chi] \in \Sigma$, choose $t \in \mathfrak{X}$ with $\chi(t) \neq 0$. Let $v = v_\chi$ be the smallest HNN valuation with associated homomorphism χ and with $v(t) = \infty$ and $v(\mathfrak{X}) \geqq 0$; in other words, v is constructed as in Example 5 of § 2 from the function $\varphi = \varphi_\chi\colon \mathfrak{X} \to \mathbf{R}_\infty$ given by $\varphi(t) = \infty$ and $\varphi(x) = 0$ for $x \neq t$. Since $[\chi] \in \Sigma$, we must have $v(\mathfrak{X}) = \infty$ by 3.3. In particular, $v(\mathfrak{X}) > 0$, so the definition of v implies that for each $x \in \mathfrak{X}$ there is a relation $x = \xi_x$ such that $\varphi(\xi_x) > 0$. Now set

$$U = \{[\eta] \in S\colon \eta(t) \neq 0 \quad \text{and} \quad \varphi_\eta(\xi_x) > 0 \quad \text{for all } x \in \mathfrak{X}\}.$$

[2] It is immediate from § 4 of [7] that our Σ has the same discrete points as that of [7]. Since the discrete $[\chi]$'s are dense in S when G is finitely generated, this makes it plausible that the two Σ's are the same.

If $[\eta] \in U$, then 2.1(i) implies that η cannot satisfy condition (iii) of 3.3. Thus U is a neighborhood of $[\chi]$ contained in Σ. □

4. Example: One-relator groups

Let G be a one-relator group $\langle \mathfrak{X}; r \rangle$ with $\mathfrak{X}$ finite and r cyclically reduced and non-trivial in $F(\mathfrak{X})$. Assume, to avoid trivialities, that $\text{card}(\mathfrak{X}) \geqq 2$. We begin by briefly reviewing the standard procedure ("Magnus rewriting") for decomposing G as an HNN extension. See, for instance, [12], § IV.5, for more details.

Suppose that some $t \in \mathfrak{X}$ has exponent sum 0 in the relator r. Then we may rewrite r as a relator r' in the elements $u_k = t^{-k} u t^k$, where $u \in \mathfrak{X} - \{t\}$. For each such u, choose a non-empty interval of integers $I(u) = [\mu(u), \mu'(u)]$ with the following two properties: (a) $I(u)$ contains all the subscripts k such that u_k occurs in r'; and (b) for at least one u which occurs in r, $I(u)$ is the smallest interval satisfying (a). Let B be the one-relator group generated by the u_k ($u \in \mathfrak{X} - \{t\}$, $k \in I(u)$) and having defining relator r'. Let B_1 (resp. B_2) be the free subgroup generated by the u_k with $\mu(u) \leqq k < \mu'(u)$ (resp. $\mu(u) < k \leqq \mu'(u)$). Then G is an HNN extension $\langle B, t; (B_1)^t = B_2 \rangle$, and the associated homomorphism is given by $\chi(t) = 1$ and $\chi(u) = 0$ for $u \in \mathfrak{X} - \{t\}$.

If $\text{card}(\mathfrak{X}) > 2$, then the given basis for B_2 omits at least two of the generators of B, and at least one of these omitted generators occurs in r'. The Freiheitssatz therefore implies that $B_2 < B$, so the HNN extension is not ascending and $[\chi] \notin \Sigma$. If $\text{card}(\mathfrak{X}) = 2$, then I claim that the HNN extension is ascending if and only if u_μ occurs exactly once in r', where u is the element of $\mathfrak{X} - \{t\}$ and $\mu = \mu(u)$. The "if" part is trivial. Conversely, if $B_2 = B$ then there has to be a relation in B expressing u_μ in terms of the u_k with $\mu < k \leqq \mu'$, and such a relation is clearly a defining relation for the free group $B = B_2$ in terms of the generators u_k, $\mu \leqq k \leqq \mu'$. The claim now follows from the fact that the cyclically reduced defining relator for a one-relator group is unique up to cyclic permutation and passage to inverses, cf. [13], Theorem 4.11.

The procedure just outlined can be used to decide for any given discrete χ whether or not $[\chi] \in \Sigma$. For we can do a sequence of "elementary operations" on $\mathfrak{X}$ [e.g., multiplying one generator by a power of another] to get a new one-relator presentation of G such that $\chi(t) = 1$ for some generator t and $\chi(u) = 0$ for $u \neq t$. Then t has exponent sum 0 in the defining relator, and Magnus rewriting is applicable. In particular, since Σ is open in the sphere S and the discrete $[\chi]$ are dense in S, it follows that $\Sigma = \emptyset$ if card $(\mathfrak{X}) > 2$. If $\text{card}(\mathfrak{X}) = 2$, on the other hand, then it is not obvious how to use this method to obtain a global description of Σ, the problem being that each χ requires a different set of generators. We will solve this problem by using HNN valuations, primarily as an aid in *guessing* a characterization of Σ. This guess turns out to be remarkably easy to prove (Theorem 4.2), and it leads to the desired global description of Σ (Theorem 4.4).

Let $\chi\colon G \to \mathbf{R}$ be a non-zero homomorphism. Choose $t \in \mathfrak{X}$ with $\chi(t) \neq 0$ and let u be the other element of $\mathfrak{X}$. By Proposition 3.3 we have $[\chi] \in \Sigma^c$ if and only if there is an HNN valuation v with $\chi_v = \chi$, $v(t) = \infty$, and $v(u) = 0$. Corol-

lary 2.3 now enables us to write down a necessary condition for $[\chi]$ to be in Σ^c. Let $\varphi\colon \mathfrak{X}^{\pm 1} \to \mathbf{R}_\infty$ be given by $\varphi(t^{\pm 1})=\infty$, $\varphi(u)=0$, and $\varphi(u^{-1})=\chi(u)$. Let the relator r be $x_1 \ldots x_n$ $(x_i \in \mathfrak{X}^{\pm 1})$, and set $v_i=\varphi(x_i)-\chi(x_1 \ldots x_{i-1})$. If there exists a v as above, then clearly we must have $n \geqq 2$, and 2.3 implies that χ satisfies:

(*) The sequence (v_i) has a repeated minimum, i.e., it assumes its minimum value more than once.

We will show that (*) is also sufficient for $[\chi]$ to be in Σ^c:

(4.2) Theorem. *Let G be a two-generator one-relator group $\langle \mathfrak{X}; r\rangle$ such that r is cyclically reduced and not the trivial relator, and let $\chi\colon G \to \mathbf{R}$ be a non-zero homomorphism. Choose $t \in \mathfrak{X}$ with $\chi(t) \neq 0$, and let v_i be as above. Then $[\chi] \in \Sigma^c$ if and only if (*) holds.*

Proof. The "only if" part has already been proved, so we will assume (*) and prove that $[\chi] \in \Sigma^c$.

Case 1. $\chi(t)=1$ and $\chi(u)=0$. Then we use Magnus rewriting to exhibit G as an HNN extension. It is easy to check that the subscripts k such that u_k occurs in the rewritten relator r' are precisely the v_i which are $<\infty$; so our assumption (*) says that u_μ occurs more than once in r' and hence, by the discussion above, the HNN extension is not ascending. This proves that $[\chi] \in \Sigma^c$.

Case 2. χ is discrete. Multiplying χ by a positive scalar, we may assume im $\chi = \mathbf{Z}$. Let $\chi(t)=p$ and $\chi(u)=q$. Let G' be obtained by adjoining to G a new generator s with $s^p=t$, and extend χ to $\chi'\colon G' \to \mathbf{Z}$ by setting $\chi'(s)=1$. Then G' is a one-relator group $\langle s, u; r'\rangle$, where r' is obtained by replacing every occurrence of t in r by s^p. Moreover, condition (*) still holds for G' and χ'. If we can construct an HNN valuation v on G' with $v(s)=\infty$, $v(u)=0$, and $\chi_v=\chi'$, then its restriction to G will show $[\chi] \in \Sigma^c$. So we may replace G by G', i.e., we may assume $p=1$.

Now replace the generator u by $u'=ut^{-q}$, so that $\chi(u')=0$. This yields a new one-relator presentation $G=\langle t, u';\ r''\rangle$, where r'' is obtained from r by substituting $u't^q$ for u and then cyclically reducing. Note that the cyclic reduction process only involves cancelling cyclically adjacent occurrences of t and t^{-1}, i.e., there is no cancellation involving u'. It follows easily that (*) still holds for the new presentation, so we have reduced Case 2 to Case 1.

Case 3. χ is not discrete. Then $\chi(t)$ and $\chi(u)$ are linearly independent over $\mathbf{Z}$. Now the numbers v_i which are $<\infty$ have the form $a_i\chi(t)+b_i\chi(u)$, where a_i and b_i are integers depending only on the relator r. So the repeated minimum guaranteed by (*) must occur for formal reasons, i.e., because of coincidences among the a_i and b_i. Hence any η sufficiently close to χ will also satisfy (*) [and will also have $\eta(t) \neq 0$], so $[\eta]$ will be in Σ^c by Case 2 if η is discrete. Since the discrete homomorphisms are dense in $\mathrm{Hom}(G, \mathbf{R})$, the desired result that $[\chi] \in \Sigma^c$ follows from the fact that Σ^c is a closed set. □

The remainder of this section will be devoted to restating Theorem 4.2 in a form that is more convenient to apply in practice. I am grateful to R. Bieri, W. Neumann, and R. Strebel for helping me arrive at this reformulation. For

simplicity, we will only consider the case where r is in the commutator subgroup $F(\mathfrak{X})'$ of $F(\mathfrak{X})$, so that the abelianization G/G' is of rank 2. (Otherwise, the sphere S contains only two points $[\pm\chi]$ and there is no need to reformulate 4.2.) The first step is to get rid of the asymmetric treatment of the two generators of G:

(4.3) Theorem. *Let $G=\langle x, y; r\rangle$ with r a cyclically reduced non-trivial element of $F(x, y)'$, and let $s_1, \ldots, s_n$ be the proper initial segments of the relator r, i.e., $s_i = x_1 \ldots x_{i-1}$. Let $\chi\colon G \to \mathbf{R}$ be a non-zero homomorphism. If $\chi(x)$ and $\chi(y)$ are both non-zero, then $[\chi]\in\Sigma$ if and only if the sequence $(\chi(s_i))$ assumes its maximum value exactly once. Otherwise, $[\chi]\in\Sigma$ if and only if the maximum occurs exactly twice.*

Proof. Replacing r by a cyclic permutation if necessary, we may assume that the maximum value of $(\chi(s_i))$ is $0=\chi(s_1)$. In case $\chi(x)$ or $\chi(y)$ vanishes, then we may also assume that $\chi(s_2)=0$. For the conditions $\chi(s_2)=\chi(x_1)\leqq 0$ and $\chi(s_n) = -\chi(x_n)\leqq 0$ imply (r being cyclically reduced) that either $\chi(x_1)=0$ or $\chi(x_n)=0$; so we may do a cyclic permutation, if necessary, to make $\chi(x_1)=0$. Finally, we may assume in both cases (by renaming the generators) that $x_1 = x$ and hence that $\chi(y)\neq 0$.

Now consider the sequence (v_i) of 4.2, with $t=y$. We have

$$v_i = \begin{cases} -\chi(s_i) & \text{if } x_i = x \\ -\chi(s_{i+1}) & \text{if } x_i = x^{-1} \\ \infty & \text{if } x_i = y^{\pm 1}. \end{cases}$$

If $\chi(x)\neq 0$ (which means that $\chi(x)<0$), it follows easily that the minimum value $0=v_1$ of (v_i) is repeated if and only if $\chi(s_i)=0$ for some $i>1$, i.e., if and only if $(\chi(s_i))$ has a repeated maximum. (The essential point here is that the non-negative sequence $(-\chi(s_i))$ is monotonic between successive occurrences of $x^{\pm 1}$; so if it ever goes down to 0 for $i>1$, then this must happen at one of the values occurring in the v-sequence.) Similarly, if $\chi(x)=0$, then (v_i) has a repeated minimum if and only if $\chi(s_i)=0$ for some $i>2$, i.e., if and only if the maximum of $(\chi(s_i))$ occurs more than twice. The theorem now follows from 4.2. □

We now restate this geometrically. View the abelianization G/G' of G as a lattice in the vector space $V = G/G' \otimes \mathbf{R}$, and view χ as a vector in the dual space V^*. We will identify both V and V^* with the Euclidean plane $\mathbf{R}^2$, and we will identify our sphere S with the unit circle in V^*; thus $[\chi]$ is thought of as the point of S obtained from χ by radial projection.

Let $v_1, \ldots, v_n$ be the images of $s_1, \ldots, s_n$ in V. It is useful to think of these lattice points v_i as the successive vertices of the closed polygonal path obtained by "tracing out" the relator r in the plane V. Suppose, for example, that r is the following relator of length 16: $r = x^{-1}y^{-1}xy^2x^{-1}y^{-1}x^2y^{-1}x^{-1}yx^{-1}yxy^{-1}$. Its trace is pictured in Fig. 2 below; one should start following it at the origin, which is shown as a heavy dot. [For clarity, multiple occurrences of the same segment have been set off from one another.]

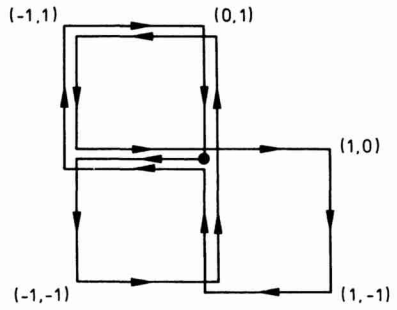

Fig. 2

Let C be the boundary of the convex hull of the v_i. It is a closed, convex polygon, whose vertices form a subset of $\{v_i\}$. Call a vertex *simple* if it equals v_i for exactly one i. The polygon C always contains a horizontal edge at both the top and bottom and a vertical edge on each side. Let e be one of these four edges. If e contains v_i for exactly two i, then e will be called *special* and will play the same role below as the simple vertices. (Note that such an e is necessarily of length 1, and its two vertices are simple.)

It is now easy to restate 4.3 geometrically, since the numbers $\chi(s_i)$ of 4.3 are simply the inner products $\langle v_i, \chi\rangle$. Recall that a line L in V is said to be a *supporting line* of C if C lies on one side of L and intersects L. For any non-zero χ, let L_χ be the supporting line of C such that χ (thought of as a vector) is orthogonal to L_χ and points away from C; thus L_χ intersects C at the points of C where the linear function $\langle -, \chi\rangle$ takes its maximum value. The following restatement of 4.3 is now immediate:

(4.4) Theorem. *If the vector χ is neither horizontal nor vertical, then $[\chi]\in\Sigma$ if and only if $L_\chi\cap C$ is a simple vertex of C; if χ is horizontal or vertical, then $[\chi]\in\Sigma$ if and only if $L_\chi\cap C$ is a special edge of C. Consequently, Σ is a finite union of open arcs of the circle S, one for each simple vertex v or special edge of C. The arc A corresponding to v (or e) is obtained as follows. Let f and f' be the edges preceding and following v (or e) when C is traversed in the counterclockwise direction, and let χ and χ' be the outward pointing unit normals at f and f'. Then A is the (counterclockwise) open arc from χ to χ'.* □

[Note that the arcs A coming from the vertices of a special edge are redunant in this description of Σ. We can exhibit this concretely by extending the sides adjacent to a special edge until they meet (assuming they are not parallel), thereby replacing the special edge and its two vertices by a single vertex, which is then considered simple.]

Returning to the example traced out in Fig. 2, C is the solid polygon in Fig. 3, with the simple vertices shown as heavy dots. The only special edge is the right hand vertical one, and the dotted lines serve as a reminder that the special edge is treated the same as a simple vertex. This yields the picture of Σ in Fig. 4, as a union of the two solid open arcs, from χ_1 to χ_2 and χ_2 to χ_3.

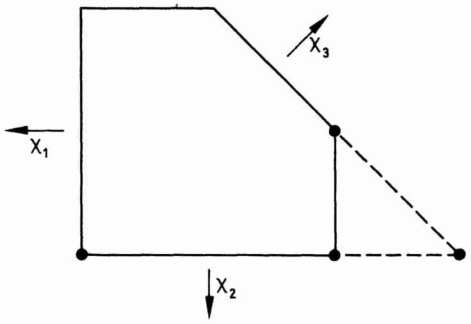

Fig. 3

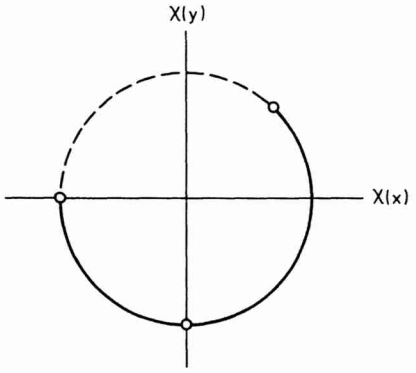

Fig. 4

In particular, we see from this and Corollary 3.2 that a normal subgroup N of G with G/N infinite cyclic is finitely generated if and only if $N = \ker \chi$ for some surjection $\chi\colon G \twoheadrightarrow \mathbf{Z}$ with $0 < \chi(y) < \chi(x)$. Intuitively, then, 1/4 of the normal subgroups with infinite cyclic quotient are finitely generated.

Remark. Suppose G is a one-relator group and we want to determine the finite generation of an arbitrary $N \lhd G$ with G/N abelian. R. Strebel has pointed out to me that the case we have treated, where $G/N \approx \mathbf{Z}$, is essentially the general case; more precisely, if N is finitely generated and G/N is abelian, then G/N necessarily has torsion-free rank $\leqq 1$, except in the trivial case where G is free abelian of rank 2. [If there were a finitely generated N with G/N abelian and of rank > 1, there would be one with $G/N \approx \mathbf{Z}^2$. Then $N \subseteq N'$ for some N' with $G/N \approx Z$, and N' is again finitely generated. Such an N' is necessarily free (cf. proof of 3.2 and first paragraph of § 3). But then the existence of a finitely generated $N \lhd N'$ with $N'/N \approx \mathbf{Z}$ implies that N' is a free of rank 1 and hence that N is trivial.]

5. Connection with Bieri-Neumann-Strebel

Let G be an arbitrary group, $\chi\colon G\to\mathbf{R}$ a non-zero homomorphism, and $\mathfrak{X}$ a set of generators of G. For any $\mathfrak{X}^{\pm 1}$-word $\xi=x_1\dots x_n$, let $t_1,\dots,t_n$ be the non-empty terminal segments of ξ, i.e., $t_i=x_i\dots x_n$. Recall from [7] that the *χ-track* of ξ is the set of real numbers $\{\chi(t_i)\}$, $1\leqq i\leqq n$. Note that if ξ represents an element of $\ker\chi$, then $\chi(t_i)=-\chi(s_i)$, where $s_i=x_1\dots x_{i-1}$ as in § 4.

The following technical lemma will provide the link between HNN valuations and the theory developed in [7]. Let $\mathfrak{X}_-=\mathfrak{X}^{\pm 1}\cap G_{\chi\leqq 0}$.

(5.1) Lemma. *If N is a subgroup of $\ker\chi$, then the following two conditions are equivalent:*

(i) *For any HNN valuation v on G with $\chi_v=\chi$ and $v|\mathfrak{X}_-$ bounded below, $v|N$ is bounded below.*

(ii) *For some real number $r\leqq 0$, every element of N can be represented by an $\mathfrak{X}^{\pm 1}$-word with χ-track bounded below by r.*

If N is normal in G, then (i) and (ii) are also equivalent to:

(iii) *Every element of N can be represented by an $\mathfrak{X}^{\pm 1}$-word with non-negative χ-track.*

Proof. (i)⇒(ii): Let $\varphi\colon\mathfrak{X}^{\pm 1}\to\mathbf{R}$ be given by $\varphi(x)=\min\{0,-\chi(x)\}$; thus $\varphi(\mathfrak{X}_-)=0$ and $\varphi(x^{-1})=\varphi(x)+\chi(x)$ for all $x\in\mathfrak{X}^{\pm 1}$. Let v be the HNN valuation constructed from φ as in Example 5 of § 2. Then $v(1)=0$, and $v(g)$ for $g\neq 1$ equals $\sup\{\varphi(\xi)\}$, where ξ ranges over the words representing g and $\varphi(x_1\dots x_n)=\min\{\varphi(x_i)-\chi(x_1\dots x_{i-1})\}$. Note that $\varphi(\xi)$ is simply the minimum value of the χ-track of ξ if ξ represents an element of $\ker\chi$. Now (i) implies that $v|N$ is bounded below, which says precisely that (ii) holds.

(ii)⇒(i): Suppose v is an HNN valuation with $\chi_v=\chi$ and $v|\mathfrak{X}_-$ bounded below by a real number c. Then 2.1(i) immediately implies that $v(g)\geqq r+c$ for any $g\in\ker\chi$ which is represented by an $\mathfrak{X}^{\pm 1}$-word with χ-track $\geqq r$. Thus (ii) implies that $v(N)\geqq r+c$.

Now suppose $N\lhd G$, and choose $x\in\mathfrak{X}^{\pm 1}$ with $\chi(x)<0$. If (ii) holds and $g\in N$, then for any integer k we can represent $x^{-k}gx^k$ by word ξ with χ-track $\geqq r$. Hence g is represented by $x^k\xi x^{-k}$, which has χ-track $\geqq 0$ if k is large enough. This shows that (ii)⇒(iii), and the converse is trivial. □

We will temporarily denote by Σ^{BNS} the set Σ defined by Bieri, Neumann, and Strebel [7] for finitely generated G. Recall the following characterization of it in terms of χ-tracks with respect to a finite generating set $\mathfrak{X}$ ([7], 3.4): $[\chi]\in\Sigma^{\mathrm{BNS}}$ if and only if every element of $\ker\chi$ can be represented by an $\mathfrak{X}^{\pm 1}$-word with non-negative χ-track. We can therefore conclude from the lemma and Proposition 2.3 that $[\chi]\in\Sigma^{\mathrm{BNS}}$ if and only if every HNN valuation v with $\chi_v=\chi$ and $v|\mathfrak{X}_-$ bounded below is trivial. Now the condition that $v|\mathfrak{X}_-$ be bounded below is vacuous when $\mathfrak{X}$ is finite. We have therefore proved:

(5.2) Theorem. *If G is finitely generated, then $[\chi]\in\Sigma^{BNS}$ if and only if there is no non-trivial HNN valuation v on G with $\chi_v=\chi$. In other words, Σ^{BNS} coincides with the set Σ defined in* § 3. □

Remark. Combining Theorem 5.2 with one of the main results of [7], we obtain a remarkable criterion for finite generation of normal subgroups of G with abelian quotient. Note first that if v is a non-trivial HNN valuation on the finitely generated group G and M is a subgroup such that $G' \subseteq M \subseteq \ker \chi_v$, then M is obviously not finitely generated; for v yields a non-trivial filtration on M by 2.3. The remarkable result, however, is that this is the only way that a normal subgroup with abelian quotient can fail to be finitely generated: If $G' \subseteq M \subseteq G$, then M is finitely generated unless G admits a non-trivial HNN valuation v with $M \subseteq \ker \chi_v$. This is Theorem B1 of [7], restated in terms of HNN valuations.

We close this section by characterizing in terms of HNN valuations the sets $\Sigma_A = \Sigma_A(G)$ considered in [7]. Recall that G is still required to be finitely generated, A is a finitely generated right G-group, and it is assumed that every element of G' acts on A by an inner automorphism of A.

(5.3) Theorem. *Let A be a G-group as above, let $H = A \rtimes G$, and let $\pi: H \twoheadrightarrow G$ be the canonical surjection. Given a non-zero homomorphism $\chi: G \to \mathbf{R}$, $[\chi] \in \Sigma_A(G)$ if and only if every HNN valuation v on H with $\chi_v = \chi \circ \pi$ is bounded below on A.*

Proof. Let $\mathfrak{X}$ be a finite set of generators of G. Recall that, for a suitable finite set of generators $\mathscr{A}$ of A as a G-group, the following condition (*) characterizes the points $[\chi]$ of Σ_A:

(*) Every element of A can be expressed as a product of elements a^g with $a \in \mathscr{A}^{\pm 1}$ and g representable by an $\mathfrak{X}^{\pm 1}$-word with non-negative χ-track.

Here a^g denotes a acted on by g, but it can also be viewed as $g^{-1}ag$ in the group H. As such, it can obviously be represented by an $(\mathfrak{X} \cup \mathscr{A})^{\pm 1}$-word with non-negative $(\chi \circ \pi)$-track if g can be represented by an $\mathfrak{X}^{\pm 1}$-word with non-negative χ-track. Conversely, given an $(\mathfrak{X} \cup \mathscr{A})^{\pm 1}$-word with non-negative $(\chi \circ \pi)$-track which represents an element of A, we can move all the elements of $\mathfrak{X}^{\pm 1}$ to the left (where they must cancel) to get a product of elements a^g as in (*). The theorem now follows at once from Lemma 5.1. □

6. Example: Metabelian groups and valuations on rings

The work of Bieri and Strebel on finite presentation of metabelian groups [5] led them to the study of the invariant $\Sigma_A = \Sigma_A(Q)$, where Q is a finitely generated abelian group and A is a finitely generated abelian Q-group, i.e., a finitely generated $\mathbf{Z}Q$-module. The crucial case to consider is the case of a *cyclic* module A, or, equivalently, the case where A is a quotient ring of $\mathbf{Z}Q$ and Q acts by multiplication. In this case Bieri and Groves ([3], Theorem 8.1), building on earlier work of Bieri and Strebel ([6], Theorem 2.1), showed that Σ_A could be calculated in terms of $\mathbf{R}$-valued valuations on the commutative ring A.

In this section we will indicate briefly how this result falls out of our characterization of Σ_A in terms of HNN valuations. To simplify the notation we will assume that the canonical map $Q \to A^*$ is a monomorphism, and we will identify Q with its image; this involves no loss of generality, since we could simply replace Q by its image.

As in Theorem 5.3, we introduce the metabelian group $G = A \rtimes Q$ and the canonical homomorphism $\pi\colon G \twoheadrightarrow Q$, and we ask which non-zero homomorphisms $\chi\colon Q \to \mathbf{R}$ have the property that $\chi \circ \pi = \chi_v$ for some HNN valuation v on G which is not bounded below on A. Now we have already seen one way to construct such a v. Namely, suppose χ extends to a function $\omega\colon A \to \mathbf{R}_\infty$ satisfying the properties (i)–(iii) mentioned in Example 2 of § 2; then we can set $v(a, q) = \omega(a)$.

Conversely, suppose there exists a v not bounded below on A with $\chi_v = \chi \circ \pi$. By change of axis (cf. end of § 2) we can arrange that $v(0, q) = \infty$ for some $q \in Q$ with $\chi(q) \neq 0$. Q being abelian, it then follows easily from 2.1(iii) that $v(0 \times Q) = \infty$. Now G is generated by $0 \times Q$ and $(1,1)$, where 1 is the multiplicative identity of the ring A (and hence also the identity of the group Q). We must therefore have $v(1,1) < \infty$, so we can normalize v to make $v(1,1) = 0$. Finally, since $(a, q) = (a, 1) \cdot (0, q)$ in G, properties (b) and (d) of HNN valuations (§ 2) imply that $v(a, q) = v(a, 1)$. Thus v has the form $v(a, q) = \omega(a)$, where $\omega(a) = v(a, 1)$, and it is a routine matter to verify that ω satisfies (i)–(iii). This proves:

(6.1) Theorem. *Let A be a commutative ring which is additively generated by a finitely generated subgroup Q of A^*. Then the complement Σ_A^c of Σ_A in the sphere $S(Q)$ is the set of classes $[\chi]$ such that χ extends to a function $\omega\colon A \to \mathbf{R}_\infty$ satisfying properties (i)–(iii) of Example 2 of § 2.* □

(*Note.* The finite generation hypothesis played no role in our proof above; it was included simply because we have not discussed Σ_A otherwise.)

As a corollary, we can easily obtain the result of Bieri-Strebel-Groves cited at the beginning of this section:

(6.2) Corollary. *Let A and Q be as in* 6.1. *Then $[\chi] \in \Sigma_A^c$ if and only if χ extends to a valuation $\omega\colon A \to \mathbf{R}_\infty$.*

Proof. We must show that if χ extends to a function ω_0 satisfying (i)–(iii), then χ extends to a valuation ω. For this purpose we use an analogue of the method of Example 5 of § 2: For $a \in A$ set $p(a) = \sup\{\varphi(\xi)\}$, where ξ ranges over all elements of $\mathbf{Z}Q$ which represent a and $\varphi(\xi) = \min\{\chi(q)\colon q \in \text{support}(\xi)\}$. [Here support$(\xi)$ is the set of $q \in Q$ which occur in ξ with non-zero coefficient.] Clearly $p(a) \leqq \omega_0(a)$, so p, like ω_0, extends χ. It is easy to check that p satisfies (i)–(iii) as well as:

(iv) $p(ab) \geqq p(a) + p(b)$.

Hence p is a *pseudovaluation*. Theorem 2 of [2] now implies that there is a valuation $\omega \geqq p$ such that $\omega | Q = p | Q = \chi$. □

7. The tree associated to an HNN valuation

Recall from § 1 that a non-trivial abelian G-$\mathbf{R}$-tree X yields, after suitable choices, a non-trivial HNN valuation v on G. For our present purposes it will be convenient to describe that construction as follows. Let e be the end of X fixed by G. Let Y be a subtree of X containing a ray $(e, x]$ and admitting an isometry

ι: $Y \xrightarrow{\approx} (r, +\infty)$ for some $r \geqq -\infty$, such that $\iota(y) \to r$ as $y \to e$. Choose such an ι and set $v(g) = \sup\{t \in \iota(Y)\colon \iota^{-1}(t)g \in Y\}$. This function $v = v_{X,Y,\iota}$ is a non-trivial HNN valuation. [Apart from notation, this is the same as what we did in § 1, except that there we insisted on setting $\iota(x) = 0$ when Y was a ray $(e, x]$. By allowing arbitrary ι, we can get v's with $v(1)$ being any finite value, not necessarily 0.]

It is easy to see the effect on v of changing the choices above. First, for fixed Y the isometry ι is unique up to an additive constant; so changing ι simply adds a constant to v. Suppose next that we replace Y by a ray $Y' = (e, x'] \subset Y$ and use the isometry $\iota' = \iota | Y'$. Then the resulting HNN valuation v' will be the truncation of v (as defined at the end of § 2) such that $v'(1) = \iota(x')$. In general, then, we can say that any two HNN valuations coming from X have truncations that differ by a constant.

(7.1) Theorem. *Every non-trivial HNN valuation $v\colon G \to \mathbf{R}_\infty$ has the form $v_{X,Y,\iota}$ for some X, Y, ι as above.*

Proof. Assume first that $v(1) < \infty$. Adding a constant if necessary, we may assume that $v(1) = 0$. Reverting to the notation of § 1, then, we wish to construct X so that $v = v_{X,x}$ for some basepoint x. The construction is essentially forced on us by equation 1.4. Let $L(g) = -(\chi(g) + 2v(g))$, where $\chi = \chi_v$. I claim that L is a Lyndon length function on G. If we accept this for the moment, then the Alperin-Moss generalization of Chiswell's theorem (cf. [1], Theorem 5.4) gives us a G-$\mathbf{R}$-tree X with a basepoint x such that $L(g) = d(x, xg)$ for all $g \in G$.

Note that $L(g^2) = L(g) - \chi(g)$ if $\chi(g) < 0$ (cf. 2.1(ii)) and that $L(g^2) \leqq L(g)$ if $\chi(g) = 0$. In both cases we can conclude from [1], 6.13(c), that the hyperbolic length of g is given by $l(g) = -\chi(g)$; in particular, l is abelian. It also follows that $v(g) = -d(x, A_g)$ if $\chi(g) \leqq 0$ ([1], 6.7).

Next we verify that X is non-trivial, i.e., that $\bigcap_{g \in G} A_g = \emptyset$ (cf. [1], Theorem 7.5). Suppose not, and let $y \in \bigcap A_g$. By the construction of X, y is in $[x, xh]$ for some $h \in G$. Replacing y by yh^{-1} and h by h^{-1} if necessary, we may assume that $\chi(h) \leqq 0$. If h is elliptic, then the midpoint of $[x, xh]$ is the unique point of $[x, xh]$ in A_h, so y must be this midpoint and $d(x, A_h) = d(x, y)$. For any $g \in G$, we then have $d(x, A_g) \leqq d(x, y) = d(x, A_h)$; in view of the previous paragraph, this implies that $v(h)$ is the minimum value of v on $G_{\chi \leqq 0}$, contradicting the non-triviality of v. If h is hyperbolic, on the other hand, then $[x, xh] \cap A_h = [p, ph]$ for some p, so $y \in [p, ph]$ and p is between yh^{-1} and y. Since $\bigcap A_g$ is a G-invariant subtree containing y, it follows that $p \in \bigcap A_g$. Then $v(h) = -d(x, A_h) = -d(x, p)$ is the minimum value of v on $G_{\chi \leqq 0}$, so we again have a contradiction.

It is now clear (cf. 1.2 and 1.3) that we will have $\chi_X = \chi$ and $v_{X,x} = v$ provided $\chi_X(g)$ and $\chi(g)$ have the same sign for all g. Otherwise, we must have $\chi_X = -\chi$ ([1], 1.4), in which case $v_{X,x} = \bar{v}$, where $\bar{v}(g) = v(g^{-1})$. Assuming, as we may, that $\chi \neq 0$, we will show that this leads to a contradiction.

Choose g such that $\chi(g) < 0$. By 2.3 we can find $n \in N = \ker \chi$ such that $v(n) < v(g)$. Replacing n by n^g if necessary, we may assume that we also have $v(n) < \bar{v}(g)$. Then $v(n^g) = v(n) + \chi(g)$ by 2.1(iii). But we may also apply the latter to $\bar{v}$ and its associated homomorphism $-\chi$. Noting that $\bar{v} = v$ on N, we obtain $v(n^g) = v(n) - \chi(g)$. This is the desired contradiction.

We now prove the claim that L is a Lyndon length function. Let δ be the function defined by the equation

$$L(g^{-1}h)=L(g)+L(h)-2\delta(g,h).$$

Explicitly, we have

$$\delta(g,h)=v(g^{-1}h)-v(g)-v(h)-\chi(g). \tag{7.2}$$

To prove the claim, we must verify the following (cf. [1], § 5):

(i) $L(1)=0$.
(ii) $L(g^{-1})=L(g)$.
(iii) $\delta(g,k)\geqq\min\{\delta(g,h),\delta(h,k)\}$.

Now (i) is simply our assumption that $v(1)=0$, and (ii) follows from our axiom (a) for HNN valuations. To prove (iii), we apply properties (a), (b), and (d) of HNN valuations to obtain $v(g^{-1}h)\geqq\min\{v(g)+\chi(g),v(h)+\chi(g)\}$, with equality if $v(g)\neq v(h)$. 7.2 now yields

$$\delta(g,h)\geqq\min\{-v(g),-v(h)\}, \tag{7.3}$$

with equality if $v(g)\neq v(h)$. (iii) follows immediately unless $v(g)=v(h)=v(k)$. But in this case we go back to 7.2 and we find that (iii) reduces to

$$v(g^{-1}k)\geqq\min\{v(g^{-1}h),v(h^{-1}k)-\chi(g^{-1}h)\},$$

which follows from axiom (b). This completes the proof of the theorem in case $v(1)<\infty$.

Suppose now that $v(1)=\infty$. For any real number M, let v_M be the truncation of v such that $v_M(1)=M$. Then v_M is still non-trivial, at least if M is sufficiently large. [This proviso is only necessary if $\chi=0$.] By what we did above, then, v_M comes from (X_M, Y_M, ι_M), where X_M is a non-trivial abelian G-**R**-tree, Y_M is a ray $(e_M, x_M]$ representing the fixed end, and ι_M: $Y_M\hookrightarrow(-\infty, M]$ is the isometry such that $\iota_M(x_M)=M$. We may assume that X_M, as in the proof above, is spanned by the G-orbit of x_M; it is then determined, up to a unique basepoint-preserving G-equivariant isometry, by v_M (cf. [1], Theorem 5.4).

I claim now that there is a canonical embedding of X_M into $X_{M'}$ if $M<M'$. For v_M is a truncation of $v_{M'}$, so the remarks at the beginning of the section imply that v_M is the HNN valuation obtained by using the tree $X_{M'}$ but replacing the ray $Y_{M'}$ by the subray $(e_{M'}, y]$, where y is the point of $Y_{M'}$ at distance $M'-M$ from $x_{M'}$. In view of the last sentence of the previous paragraph, then, we may identify X_M with the subtree of $X_{M'}$ spanned by the orbit of y, whence the claim. The desired X, Y, ι yielding v are now obtained by passage to the limit as $M\to\infty$: Take $X=\bigcup X_M$ and $Y=\bigcup Y_M$, and let ι be the isometry such that $\iota|Y_M=\iota_M$. □

Combining this with Theorem 5.2, we obtain:

(7.4) Corollary. *If G is finitely generated, then the complement Σ^c of the Bieri-Neumann-Strebel invariant Σ is the set of classes $[\chi_X]$ obtained from non-trivial abelian G-**R**-trees X.* □

Remarks. 1. If we take v in the theorem to be the HNN valuation associated to a valuation ω on a ring A as in Example 2 of § 2 (with $Q = A^*$, say), then we obtain a tree $X = X_\omega$, with an action of the group $A \rtimes A^*$. This generalizes to arbitrary rings the construction, well-known for fields, of a tree associated to a valuation. In the case of a field, however, one knows that the action of $A \rtimes A^*$ extends to an action of $GL_2(A)$ on X. I do not know if this is true in general.

2. It is not hard to deduce from the proof of the theorem that abelian G-**R**-trees can be classified by HNN valuations. In order to state this more carefully, we will confine our attention to the case where $\chi_X \neq 0$ and X is *minimal*, i.e., X has no proper G-invariant subtrees. [This simply means that X is the union of the axes of the hyperbolic elements of G.] Then the result is that minimal non-trivial abelian G-**R**-trees X with $\chi_X \neq 0$ are classified by equivalence classes of non-trivial HNN valuations v with $\chi_v \neq 0$, where two such v's are equivalent if they have truncations which differ by a constant.

In practice, however, we do not work directly with this equivalence relation. It is usually more convenient to try to find, for a given G and χ, a canonical way to choose Y and ι in any X with $\chi_X = \chi$. Suppose, for instance, that G is a 2-generator 1-relator group as in § 4. Given χ, let t be a generator with $\chi(t) \neq 0$ and let u be the other generator. Then for any X with $\chi_X = \chi$, we can take $Y = A_t$ and we can choose ι to make the resulting $v = v_{X,Y,\iota}$ satisfy $v(u) = 0$. It follows easily that minimal non-trivial abelian G-**R**-trees X with $\chi_X = \chi$ are classified by HNN valuations v with $v(t) = \infty$, $v(u) = 0$, and $\chi_v = \chi$.

A second example is provided by the metabelian groups $G = A \rtimes Q$ and the homomorphisms of the form $\chi \circ \pi$ considered in § 6. We can then always take Y to be the axis of Q (i.e., the unique Q-invariant line), and we can choose ι so that the resulting v will have $v(1,1) = 0$. It follows easily (cf. proof of 6.1) the minimal non-trivial abelian G-**R**-trees X with $\chi_X = \chi \circ \pi$ for some non-zero homomorphism χ on Q are classified by functions $\omega\colon A \to \mathbf{R}_\infty$ as in Example 2 of § 2, with $\omega | Q = \chi$.

8. Examples

Given a group G, not necessarily finitely generated, we continue to write $\Sigma^c = \Sigma^c(G)$ for the set of classes $[\chi]$ of non-zero homomorphisms $\chi\colon G \to \mathbf{R}$ such that $\chi = \chi_v$ for some non-trivial HNN valuation v, or, equivalently, such that $\chi = \chi_X$ for some non-trivial abelian G-**R**-tree X. In this section we will use tree actions to compute Σ^c in two families of examples. We begin with some general remarks.

Suppose we already know one non-trivial HNN valuation v on G, with $\chi = \chi_v \neq 0$. Adding a constant to v if necessary, we may assume that G contains elements g with $\chi(g) < 0$ and $v(g) \geqq 0$. Let B be the "base group" $N_0 = N \cap v^{-1}([0, +\infty])$, where $N = \ker \chi$. Then for any g as above, the conjugates of B by the positive powers of g increase and exhaust N by 2.1(v). We will use this "generalized HNN structure" on G as an aid in computing Σ^c. More generally, we will use it to analyze actions of G on **R**-trees:

(8.1) Lemma. *Let X be a G-**R**-tree which is abelian as a B-**R**-tree. Then X is either abelian or dihedral. More precisely, X satisfies one of the following mutually exclusive conditions.*

(i) *$X^N \neq \emptyset$, where X^N is the fixed point set of N in X. In this case X is abelian as a G-**R**-tree, and either X is trivial or $\chi_X = 0$.*

(ii) *$X^N = \emptyset$ but $X^B \neq \emptyset$. In this case X is abelian and non-trivial, and $\chi_X = c\chi$ for some real number $c \geqq 0$.*

(iii) *X contains a B-invariant line Y on which B acts as a non-trivial group of translations. In this case Y is also G-invariant.*

(iv) *X is non-trivial as an abelian B-**R**-tree. In this case X is also abelian as a G-**R**-tree (and non-trivial).*

Proof. We will use the method of Tits [15]; see also [1], § 7. By the classification of abelian B-**R**-trees, we have the following three possibilities for the action of B on X: (a) $X^B \neq \emptyset$. (b) X contains a B-invariant line Y on which B acts as a non-trivial group of translations. (c) X is non-trivial as an abelian B-**R**-tree. We examine these cases one at a time.

Case (*a*). If B has fixed points in X then so does any conjugate of B; hence every element of N is elliptic. Thus either $X^N \neq \emptyset$ or else $X^N = \emptyset$ and N has a unique fixed end e. In the first case, the abelian group G/N acts on the tree X^N, and it follows easily that (i) holds. In the second case, e is also fixed by G since $N \lhd G$, hence X is abelian and non-trivial. Assuming $\chi_X \neq 0$, I claim that χ_X is a positive scalar multiple of χ, so that (ii) holds. We already know that $\chi_X(g) = 0$ when $\chi(g) = 0$; the claim will follow if we can show that $\chi_X(g) < 0$ when $\chi(g) < 0$.

Suppose, on the contrary, that $\chi_X(g) \geqq 0$ for some g with $\chi(g) < 0$. Since any conjugate of B has fixed points in X, it follows that the group $B' = N_r$ has fixed points for any $r \leqq 0$. Take $|r|$ large enough that the conjugates of B' by the positive powers of g increase and exhaust N, and choose a point x fixed by B'. Replacing x by a suitable point on the half-line $(e, x]$ if necessary, we may assume that $x \in A_g$. But then the assumption that $\chi_X(g) \geqq 0$ implies that the stabilizer of x is closed under conjugation by g and hence contains N. This contradicts the assumption that $X^N = \emptyset$.

Case (*b*). The line Y in (b) is necessarily the *unique* B-invariant line, since it is the axis of any hyperbolic element of B. Choose $g \in G$ so that the conjugates of B by the positive powers of g increase and exhaust N. Then Yg is the B^g-invariant line and is also B-invariant since $B \subset B^g$; hence $Yg = Y$. Repeating this argument, we see that Y is in fact N-invariant. N being normal in G, the same argument shows that Y is G-invariant, so (iii) holds.

Case (*c*). Let e be the end fixed by B. Arguing as in case (b), we see that N fixes e, hence G fixes e and (iv) holds. □

We now complete the discussion of Example 6 of § 2.

(8.2) Proposition. *Let G be the group of orientation-preserving piecewise linear homeomorphisms of the unit interval. Then every G-**R**-tree X is abelian. If X*

is non-trivial and $\chi_X \neq 0$, *then* $[\chi_X] = [\log \lambda]$ *or* $[\log \rho]$, *where* λ *and* ρ *are the derivative homomorphisms at* 0 *and* 1. *In particular,* $\Sigma^c = \{[\log \lambda], [\log \rho]\}$.

Proof. Let X be a G-$\mathbf{R}$-tree. Suppose first that either $\ker \lambda$ or $\ker \rho$, say $\ker \rho$, contains a hyperbolic element g. If J is a sufficiently small interval $[1-\varepsilon, 1]$ and G_J is the subgroup of G consisting of homeomorphisms supported in J, then G_J commutes with g and hence fixes the two ends of the axis $Y = A_g$. We now apply the lemma, with $\chi = \log \lambda$, v as in Example 6 of § 2 [modified by an additive constant], and $B = G_J$. Since we already know B fixes two ends of X, the lemma is applicable and case (iv) is impossible. Thus X is abelian except possibly in case (iii), and $[\chi_X] = [\log \lambda]$ if X is non-trivial.

To see that X is necessarily abelian if (iii) holds, recall from the proof of the lemma that $N = \ker \lambda$ fixes the two ends of Y. On the other hand, for any $g \in \ker \rho$ we can find a conjugate B' of B such that g commutes with B' and hence with some hyperbolic element of N. So g also fixes the two ends of Y. Since G is generated by $\ker \lambda$ and $\ker \rho$, it follows that G fixes the two ends of Y. Thus X is abelian.

Finally, suppose that $\ker \lambda$ and $\ker \rho$ both consist entirely of elliptic elements. Then we can again apply the lemma, still with $N = \ker \lambda$. This time case (iii) is impossible, hence X is abelian. It follows that the elliptic elements form a subgroup, so our hypothesis implies that every element of G is elliptic. Thus $\chi_X = 0$ if X is non-trivial. □

Remark. The proof applies verbatim to any irreducible subgroup H of G on which λ and ρ are independent, where the terminology is that of [7], § 8. Hence we recover the Bieri-Neumann-Strebel calculation ([7], Theorem 8.1) that $\Sigma^c = \{[\log \lambda | H], [\log \rho | H]\}$ for such an H, provided H is finitely generated (so that their Σ is defined).

For our second and final family of examples we consider a sequence of groups first introducted by Houghton [11]. (See also [9], § 5.) Fix an integer $n \geqq 1$, let $\mathbf{N}$ be the set of positive integers, and let $S = \mathbf{N} \times \{1, \ldots, n\}$. Thus S is the disjoint union of n copies of $\mathbf{N}$. The i-the copy, $\mathbf{N} \times i$, will be called the i-the *prong*. The *n-prong Houghton group* is defined to be the group H of all permutations g of S such that g is eventually a translation on each prong. More precisely, we require that there be an n-tuple $(m_1, \ldots, m_n) \in \mathbf{Z}^n$ such that for each $i \in \{1, \ldots, n\}$ one has $(x, i) \cdot g = (x + m_i, i)$ for all sufficiently large $x \in \mathbf{N}$.

The assignment $g \mapsto m_i$ defines a homomorphism $\alpha_i \colon H \to \mathbf{Z}$, $i = 1, \ldots, n$. If $n \geqq 2$, it is easy to see that α_i is the homomorphism associated to a decomposition of H as a properly ascending HNN extension. As base group we can take the group H_i consisting of those elements of H which fix the i-th prong pointwise. And as stable letter we can take the "translation" $t = t_{ji}$ along the union of the i-th and j-th prongs for any $j \neq i$, this being defined by

$$\begin{aligned} (x, j) \cdot t &= (x-1, j) && \text{for } x > 1 \\ (1, j) \cdot t &= (1, i) \\ (x, i) \cdot t &= (x+1, i) \\ (x, k) \cdot t &= (x, k) && \text{if } k \neq j, i. \end{aligned}$$

Note that H_i is isomorphic to the $(n-1)$-prong Houghton group.

If $n=1$, H is an infinite, locally finite group and $\alpha_1=0$. If $n\geqq 2$, however, it is easy to see that H is finitely generated. And if $n\geqq 3$, then H is generated by the t_{ij}.

(8.3) Proposition. *Every H-**R**-tree X is abelian. If X is non-trivial, then χ_X is a negative scalar multiple of α_i for some i. Hence $\Sigma^c=\{[-\alpha_1], \ldots, [-\alpha_n]\}$.*

This calculation of Σ was also obtained by Bieri and Strebel [unpublished].

Proofs of 8.3. The first assertion follows from the fact that H has a locally finite normal subgroup with abelian quotient. To calculate χ_X if X is non-trivial, we may assume $n\geqq 3$. Suppose χ_X is not a negative multiple of any of the n known α's. Since H is finitely generated, we must have $\chi_X\neq 0$, so $\chi_X(t)\neq 0$ for some $t=t_{ij}$. Apply Lemma 8.1 with $B=H_i$ and $\chi=-\alpha_i$. Since we are *assuming* that X is non-trivial, the only cases that can occur are (ii) and (iv). And since we have also assumed that (ii) does not hold, the only possibility is (iv), i.e., X is non-trivial as an abelian H_i-**R**-tree.

Now let $H_{ij}=H_i\cap H_j$, and apply Lemma 8.1 with $G=H_i$, $B=H_{ij}$, and $\chi=-\alpha_j|H_i$. As above, the only possibilities are (ii) and (iv); but this time it is (iv) that cannot occur, because H_{ij} commutes with t and so fixes the two ends of the axis A_t. Hence $\chi_X|H_i$ is a negative multiple of $\alpha_j|H_i$.

Interchanging i and j in the argument above, we also have that $\chi_X|H_j$ is a negative multiple of $\alpha_i|H_j$. It follows that $\chi_X(t_{ik})\neq 0$ for any k, so we may replace j by k above to conclude that $\chi_X|H_i$ is a negative multiple of $\alpha_k|H_i$. But this cannot be simultaneously true for all $n-1$ values of $k\neq i$, so we have a contradiction. □

References

1. Alperin, R., Bass, H.: Length functions of group actions on Λ-trees. Gersten, S.M., Stallings, J.R. (eds.). Combinatorial group theory and topology. Ann. Math. Stud. **111** (1986)
2. Bergman, G.M.: A weak Nullstellensatz for valuations. Proc. Am. Math. Soc. **28**, 32–38 (1971)
3. Bieri, R., Groves, J.R.J.: The geometry of the set of characters induced by valuations. J. Reine Angew. Math. **347**, 168–195 (1984)
4. Bieri, R., Strebel, R.: Almost finitely presented soluble groups. Comment. Math. Helv. **53**, 258–278 (1978)
5. Bieri, R., Strebel, R.: Valuations and finitely presented metabelian groups. Proc. Lond. Math. Soc. (3) **41**, 439–464 (1980)
6. Bieri, R., Strebel, R.: A geometric invariant for modules over an abelian group, J. Reine Angew. Math. **322**, 170–189 (1981)
7. Bieri, R., Neumann, W.D., Strebel, R.: A geometric invariant of discrete groups. Invent. Math. **90**, 451–477 (1987)
8. Bourbaki, N.: Commutative algebra. Hermann (Paris) and Addison-Wesley (Reading) 1972
9. Brown, K.S.: Finiteness properties of groups. J. Pure Appl. Algebra **44**, 45–75 (1987)
10. Culler, M., Morgan, J.W.: Group actions on **R**-trees (preprint)
11. Houghton, C.H.: The first cohomology of a group with permutation module coefficients. Arch. Math. **31**, 254–258 (1978/1979)
12. Lyndon, R.C., Schupp, P.E.: Combinatorial group theory. Berlin Heidelberg New York: Springer 1977

13. Magnus, W., Karrass, A., Solitar, D.: Combinatorial group theory. 2nd edition, New York: Dover 1976
14. Serre, J.-P.: Trees, Berlin Heidelberg New York: Springer 1980 (Translation of Arbres, amalgames, SL_2, Astérisque 46 (1977))
15. Tits, J.: A "theorem of Lie-Kolchin" for trees. In: Contributions to algebra: A collection of papers dedicated to Ellis Kolchin, Academic Press, New York, 1977, pp. 377–388

Oblatum 4-IX-1986 & 26-VI-1987