Notes on compact nilspaces

Pablo Candela

Received 31 May 2016; Revised 25 May 2017; Published 7 September 2017

Contents

1	Intr	oduction	3
	1.1	Filtered nilmanifolds as compact nilspaces	5
2	Cha	racterization of compact nilspaces	7
	2.1	Topological preliminaries	/
		2.1.1 Continuous abelian bundles	9
		2.1.2 Metrics	J
	2.2	Measure-theoretic preliminaries	2
		2.2.1 Continuous systems of measures	2
		2.2.2 Haar measure on compact abelian bundles and nilspaces	3
		2.2.3 Probability spaces of morphisms	7
	2.3	Topological spaces associated with continuous systems of measures	J
	2.4	Borel measurable cocycles generate compact extensions	8
	2.5	Compact abelian bundles of finite rank, and averaging	3
	2.6	Counting isomorphism classes of compact finite-rank nilspaces	5
	2.7	Compact nilspaces as inverse limits of finite-rank nilspaces	7
		2.7.1 Proof of the inverse limit theorem	9
	2.8	Rigidity of morphisms	3
	2.9	Characterizing compact connected nilspaces of finite rank	
		2.9.1 Toral nilspaces	1

Chapter 1

Introduction

These notes form the second part of a detailed account of the theory of nilspaces developed by Camarena and Szegedy in the paper [9]. The material in these notes expands on the third chapter of their paper, providing more detailed proofs of the main results. To that end we also include several additional results that are implicit in [9]. This material relies strongly on the first part of our exposition, given in [8]. We shall have to assume some familiarity with the basic theory of nilspaces and some of their algebraic properties, but we shall always refer to the relevant results in [8].

Recall the notion of a nilspace from [8, Definition 1.2.1]. The principal objects treated in these notes are *compact nilspaces*. These are nilspaces equipped with a compact topology that is compatible with the cube structure. To give the precise definition, we adopt the following terminology.

Definition 1.0.1. Throughout the sequel, to be concise we shall use the term *compact space* to mean a compact, Hausdorff, second-countable topological space.

In particular, compact spaces in these notes are Polish spaces [28, Theorem 5.3]. More generally, topological spaces will usually be assumed to be Hausdorff and second-countable.

Recall from [8, Definition 1.2.3] the notion of a cubespace.

Definition 1.0.2. A cubespace X is called a *compact cubespace* if X is a compact space and $C^n(X)$ is a closed subset of $X^{\{0,1\}^n}$ for every $n \in \mathbb{N}$. A *compact nilspace* is a compact cubespace satisfying the ergodicity and corner-completion axioms from [8, Definition 1.2.1].

The main goal in the sequel is to characterize compact nilspaces.

A central result in this direction is a description of a general compact nilspace as an inverse limit of simpler spaces, namely compact nilspaces of finite rank (these are defined in Section 2.5). This result is treated in Section 2.7 (see Theorem 2.7.3).

Another central result concerns finite-rank compact nilspaces, and involves *nilmanifolds*, i.e. compact homogeneous spaces of nilpotent Lie groups, the study of which goes back to [31]. To situate the result, let us note that nilmanifolds provide important examples of compact nilspaces. Indeed, a nilmanifold becomes a compact nilspace when, given a filtration on the Lie group (with certain topological properties), the nilmanifold is equipped with the natural cubes associated with the filtration, namely the cubes introduced by Host and Kra [25, 26]. This construction is treated from a purely algebraic viewpoint in [8, Section 2.3], and its basic topological aspects are detailed in Section 1.1 below. The central result in question here goes in the converse direction, and states that if a compact nilspace of finite rank has connected structure groups, ¹ then it is isomorphic to a nilmanifold with a cube structure of the kind mentioned above. This is treated in Section 2.9 (see Theorem 2.9.17).

¹For the notion of the structure groups of a nilspace, recall [8, Definition 3.2.17] and [8, Theorem 3.2.19].

P. CANDELA

The work toward these main theorems yields several other results of inherent interest, including the following: endowing every compact nilspace with a Borel probability measure that generalizes the Haar measure on compact abelian groups (Proposition 2.2.5); an automatic continuity result for Borel morphisms between compact nilspaces (Theorem 2.4.6); a rigidity result for morphisms into compact nilspaces of finite rank (Theorem 2.8.2). Several of the main tools used to obtain these results rely on the theory of *continuous systems of measures*, which provides a natural framework for measure theoretic aspects of compact nilspaces. This is detailed in Section 2.2.

An alternative treatment of compact nilspaces is given by Gutman, Manners, and Varjú in the series of papers [20, 21, 22].

Let us end this introduction by evoking one of the central motivations for the study of compact nilspaces. This motivation concerns the analysis of uniformity norms. These norms were introduced in arithmetic combinatorics by Gowers [15], and independently in ergodic theory (as the analogous uniformity seminorms) by Host and Kra [25]. The uniformity norms, or U^d norms (one such norm for each integer $d \ge 2$), are defined on the space of complex-valued functions on a finite (or more generally a compact) abelian group. These norms provide useful tools to control averages of bounded functions over certain linear configurations in abelian groups, configurations such as arithmetic progressions (such control can be obtained via the so-called generalized Von Neumann theorems); see [17] and [18]. A major topic concerning these norms is the analysis, for each $d \ge 2$, of the harmonics that determine whether the U^d norm of a function is small or large. For the U^2 norm these characteristic harmonics are simply the characters from Fourier analysis. For d > 2, this topic leads to the theory of higher order Fourier analysis. Central to this theory are the results known as *inverse theorems* for the U^d norms, proved by Green, Tao and Ziegler [19], and independently by Szegedy [39]. Essentially, these theorems tell us that, for the U^d norm of functions on $\mathbb{Z}/N\mathbb{Z}$, one can use as characteristic harmonics the functions known as (d-1)-step nilsequences. In the approach to these theorems developed by Szegedy (consisting principally in [39] and his work with Camarena [9]), compact nilspaces play a fundamental role. Indeed, roughly speaking, in this approach a nilsequence on $\mathbb{Z}/N\mathbb{Z}$ is obtained as a composition $F \circ \varphi$, where φ is a nilspace morphism from $\mathbb{Z}/N\mathbb{Z}$ to a compact nilspace X, and F is some continuous function $X \to \mathbb{C}$. In this way, compact nilspaces play a role in higher order Fourier analysis that generalizes the role played by the circle group \mathbb{R}/\mathbb{Z} in the definition of characters in Fourier analysis. In these notes we shall concentrate on the theory of compact nilspaces in itself. For further introduction to higher order Fourier analysis and its applications, we refer the reader to the survey [16].

Acknowledgements. I am very grateful to Balázs Szegedy for discussions that were crucial for my understanding of [9], and to anonymous referees for advice that helped to improve this paper. I also thank Yonatan Gutman, Frederick Manners and Péter Varjú for informing me of their work prior to publication, and for useful comments. The present work was supported by the ERC Consolidator Grant No. 617747.

1.1 Filtered nilmanifolds as compact nilspaces

Recall from [8, Section 2.2] the notion of a filtration G_{\bullet} on a group G and the basic properties of cube sets of the form $C^n(G_{\bullet})$.

Definition 1.1.1. A *nilmanifold* is a quotient space G/Γ where G is a connected nilpotent Lie group and Γ is a discrete, cocompact subgroup. If G_{\bullet} is a filtration on G of degree at most k, with terms G_i being closed subgroups of G, and with each subgroup $\Gamma \cap G_i$ being cocompact in G_i , then we call $(G/\Gamma, G_{\bullet})$ a *filtered nilmanifold* of degree at most k.

The groups in the filtration G_{\bullet} are sometimes required to be connected (in particular in applications of nilmanifolds in arithmetic combinatorics; see for instance [17, Definition 1.3]), but we shall not need this assumption here.

A compact nilspace X is *connected* if the underlying topological space is connected. We now show that every filtered nilmanifold of degree k can be viewed as a k-step connected compact nilspace, in the following sense.

Proposition 1.1.2. Let G be a connected nilpotent Lie group with a degree-k filtration G_{\bullet} of closed subgroups, let $\Gamma \leq G$ be discrete, and for each $i \in [k] = \{1, 2, ..., k\}$ let $\Gamma_i = \Gamma \cap G_i$. Let X denote the coset space G/Γ equipped with the sets $C^n(X) = \{\pi_{\Gamma} \circ c : c \in C^n(G_{\bullet})\}$, where π_{Γ} is the quotient map $G \to G/\Gamma$. Then X is a connected compact nilspace if and only if Γ_i is cocompact in G_i for each $i \in [k]$.

Proof. The nilspace axioms are proved for X in [8, Proposition 2.3.1]. To complete the proof, we show that the compactness of each G_i/Γ_i is necessary and sufficient for each set $C^n(X)$ to be closed in $X^{\{0,1\}^n}$.

To see the sufficiency, it is enough to deduce that $C^n(G_{\bullet}) \cap \Gamma^{\{0,1\}^n}$ is cocompact in $C^n(G_{\bullet})$, as then $C^n(X)$ is homeomorphic to the compact set $C^n(G_{\bullet})/(C^n(G_{\bullet}) \cap \Gamma^{\{0,1\}^n})$ (see [30, Proposition 2.7]). This cocompactness is proved elsewhere in the literature but we include a proof here for completeness, following the argument from [18, Lemma E.10]. By Definition 1.1.1 there is a compact set $K_j \subseteq G_j$ such that the product set $K_j \Gamma_j$ equals G_j . Recall from [8, Lemma 2.2.5] that every cube $c \in C^n(G_{\bullet})$ has a unique factorization $c = g_0^{F_0} g_1^{F_1} \cdots g_{2^n-1}^{F_{2^n-1}}$ with $g_i \in G_{\operatorname{codim}(F_i)}$ for each i, where the F_i are the upper faces of $\{0,1\}^n$, of the form $F_i = F(v_i) = \{v : \operatorname{Supp}(v) \supset \operatorname{Supp}(v_i)\}$ and the v_i are ordered in the colex order. We then have the following inductive formula for the coefficients g_i :

$$g_0 = c(v_0)$$
, and for each $i > 0$, $g_i = \left(\prod_{j < i: \text{Supp}(v_j) \subset \text{Supp}(v_i)} g_j^{F_j}\right)^{-1} \cdot c(v_i)$. (1.1)

Let $H_0 = \mathbb{C}^n(G_{\bullet})$ and for each $i \in [2^n]$ let $H_i = \{c \in \mathbb{C}^n(G_{\bullet}) : g_0 = \dots = g_{i-1} = \mathrm{id}_G\}$. It follows from (1.1) that $c \in H_i$ if and only if $c(v_j) = \mathrm{id}_G$ for all j < i. We deduce that H_i is a normal subgroup of $\mathbb{C}^n(G_{\bullet})$. We shall now use these subgroups to show that every $c \in \mathbb{C}^n(G_{\bullet})$ is a product of an element of $\mathbb{C}^n(G_{\bullet}) \cap \Gamma^{\{0,1\}^n}$ with an element of $K := \{k_0^{F_0} \dots k_{2^n-1}^{F_{2^n-1}} : k_i \in K_i\}$. This will prove the claimed sufficiency, since the set K is compact. Let $i \in [2^n]$ and suppose that $c \in H_{i-1}$. Then $c = g_i^{F_i} c'$, where $c' \in H_i$. We then have $g_i^{F_i} = (k_i \gamma_i)^{F_i}$ where $k_i \in K_{\operatorname{codim}(F_i)}$ and $\gamma_i \in \Gamma_{\operatorname{codim}(F_i)}$. Since H_i is normal, we have $c = k_i^{F_i} c'' \gamma_i^{F_i}$ with $c'' \in H_i$. Starting with any cube in H_0 and repeating this procedure, we eventually obtain the desired factorization showing that $c \in K \cdot \left(\mathbb{C}^n(G_{\bullet}) \cap \Gamma^{\{0,1\}^n}\right)$.

P. CANDELA

To see the necessity of each set G_i/Γ_i being compact, let Q_i denote the set of cubes $c \in C^i(X)$ such that c(v) is the trivial coset Γ for all $v \neq 0^i$. Since $C^i(X)$ is compact, so is Q_i . We have $c = \pi_{\Gamma} \circ c'$ for some $c' \in C^i(G_{\bullet})$ with $c'(v) \in \Gamma$ for $v \neq 0^i$. The fact that $c' \in C^i(G_{\bullet})$ implies that the Gray-code product $\sigma_i(c')$ lies in G_i (recall [8, Definition 2.2.22 and Proposition 2.2.25]). By the formula for the Gray-code product, we deduce that $\sigma_i(c')c'(0^i)^{-1} \in \Gamma$. It follows that $c'(0^i) \in G_i\Gamma$. Thus the projection $\pi_{0^i}: c \to c(0^i)$ maps Q_i into $(G_i\Gamma)/\Gamma$. Conversely, for every coset $g\Gamma$ with $g \in G_i$ there is a cube $c \in Q_i$ with $\pi_{0^i}(c) = g\Gamma$, namely $c = \pi_{\Gamma} \circ c'$ where $c'(0^i) = g$ and $c'(v) = \mathrm{id}_G$ otherwise. It follows that $\pi_{0^i}(Q_i) = (G_i\Gamma)/\Gamma$, so the compactness of Q_i implies that $(G_i\Gamma)/\Gamma$ is compact, whence G_i/Γ_i is compact.

Chapter 2

Characterization of compact nilspaces

2.1 Topological preliminaries

In order to characterize a general compact nilspace, we first have to examine how the basic algebraic structures treated in [8] behave under the additional topological assumptions in Definition 1.0.2.

We begin with a result showing that the closure condition on $C^n(X)$ needs to be checked only for n = k + 1.

Lemma 2.1.1. A k-step nilspace X is a compact nilspace if and only if X is a compact space and $C^{k+1}(X)$ is a closed subset of $X^{\{0,1\}^{k+1}}$.

Proof. Note that for all $n \le m$ we can view $C^n(X)$ as the projection of $C^m(X)$ to $X^{\{0,1\}^n}$ (by the composition axiom, embedding $\{0,1\}^n$ as a face in $\{0,1\}^m$). In particular, if we assume that $C^{k+1}(X)$ is closed then for every n < k+1 the projection $C^n(X)$ is also closed. If n > k+1, then for any (k+1)-dimensional face F let

$$Q_F = \{ f : \{0,1\}^n \to X \mid f \circ \phi_F \in C^{k+1}(X) \},$$

where $f \circ \phi_F$ is the restriction of f to F (recall [8, Definition 1.1.5]). Thus Q_F is the preimage of a closed set under the restriction map $f \mapsto f \circ \phi_F$, and is therefore closed. Then, by [8, Lemma 3.2.13], we have $C^n(X) = \bigcap_F Q_F$ where F runs through (k+1)-dimensional faces of $\{0,1\}^n$.

In the remainder of this section we establish topological properties of the structures in [8, Chapter 3] under the compactness assumption. The main result that we shall obtain along these lines is that if a finite-step nilspace is compact, then the associated abelian bundle obtained in [8, Theorem 3.2.19] is endowed with a topology making it an iterated *continuous* abelian bundle (in a sense that we shall formalize in Definition 2.1.6), with structure groups being compact abelian groups.

We begin by looking at the behaviour of some basic constructions. Recall from [8, Definitions 3.1.19 and 3.1.22] the notions of arrow spaces and the spaces $\partial_x X$.

Lemma 2.1.2. Let X be a compact nilspace, let Z be a compact abelian group, and fix any $x \in X$. Then the arrow spaces $X \bowtie_i X$, $i \ge 1$, the nilspaces $\mathfrak{D}_k(Z)$, $k \ge 1$, and the space $\partial_x X$, are all compact nilspaces.

Proof. The space $X \times X$ underlying $X \bowtie_i X$ is compact and $C^{k+1}(X \bowtie_i X)$ is the image under a continuous map of a closed set of cubes in $C^{k+1+i}(X)$ (see [8, (3.4)]), so it is a compact set as required. The cube set $C^{k+1}(\mathcal{D}_k(Z))$ is also compact, since it is the set of solutions of a linear equation (recall [8, (2.9)]). Finally, the cube set $C^{k+1}(\partial_x X)$ is a closed subset of $C^{k+1}(X \bowtie_1 X)$ (recall [8, Definition 3.1.22]).

We now look at the more involved constructions, especially those using quotients by equivalence relations.

Remark 2.1.3. We shall use the following standard facts.

- (i) Let X, Y be compact spaces (as in Definition 1.0.1). Then a function $f: X \to Y$ is continuous if and only if its graph is closed. (This is a special case of the closed graph theorem [32, Ex. 8, p. 171].)
- (ii) Compact spaces are metrizable. (This follows from Urysohn's metrization theorem [32, Thm. 34.1, p. 215], using that a compact Hausdorff space is regular [32, Ex. 3, p. 205].)
- (iii) Let X be a compact space and let \sim be a closed equivalence relation on X, that is $\{(x,y): x \sim y\}$ is a closed subset of $X \times X$. Then X / \sim is compact in the quotient topology [6, Prop. 8, p. 105].

Recall from [8, Definition 1.2.3] that X is said to be *k-fold ergodic* if $C^k(X) = X^{\{0,1\}^k}$.

Lemma 2.1.4. Let X be a k-step k-fold ergodic compact nilspace. Then X is isomorphic as a compact nilspace to $\mathfrak{D}_k(\mathbf{Z})$ for some compact abelian group Z.

Proof. As seen in [8, Proposition 2.4.1], in the case k = 1, if we fix any $e \in X$ as a distinguished point, then the cube structure on X yields a commutative group operation on X with identity element e, making X a principal homogeneous space of the resulting abelian group Z. Let us check the continuity of addition and inverse on Z. The graph of addition can be written

$$\{(x_{00}, x_{10}, x_{01}, x_{11}) \in \mathbb{C}^2(\mathbb{X}) : x_{00} = e\}.$$

The graph of the inverse is similarly defined by the equation $x_{00} = x_{11} = e$. These are closed sets, so the operations are continuous.

For k > 1, as in [8, Proposition 3.2.14] we consider the 1-step nilspace $\partial_e^{k-1}(X)$, which is here a compact nilspace, and therefore isomorphic to $\mathcal{D}_1(Z)$ for some compact abelian group Z. We then argue just as in the proof of [8, Proposition 3.2.14] to show that $C^n(X) = C^n(\mathcal{D}_k(Z))$ for all n.

Recall from [8, Definition 3.2.3] the relation \sim_k on a nilspace X and the corresponding nilspace factor $\mathcal{F}_k(X)$ from [8, Lemma 3.2.10].

Lemma 2.1.5. Let X be a compact nilspace and let $k \in \mathbb{N}$. Then $\mathcal{F}_k(X)$ with the quotient topology is a compact nilspace.

Proof. As a subset of X × X, the relation \sim_k consists of restrictions to the 1-face $\{0^{k+1}, (1,0,\ldots,0)\}$ of the cubes satisfying the condition in [8, Lemma 3.2.4]. The set of these cubes is

$$\mathbf{C}^{k+1}(\mathbf{X}) \, \cap \, \big\{ f \in \mathbf{X}^{\{0,1\}^{k+1}} \ : \ \forall \nu, \nu' \in \{0,1\}^{k+1} \setminus \{0^{k+1}\}, \ f(\nu) = f(\nu') \big\},$$

and is therefore closed. \Box

We now move toward the main result of this section, namely the addition of suitable topological properties to the bundle decomposition of a nilspace.

2.1.1 Continuous abelian bundles

Recall from [8, Definition 3.2.17] the notion of a *k*-fold abelian bundle.

Definition 2.1.6. Let B be an abelian bundle with base S, structure group Z and projection π . We say that B is *continuous* if the following conditions hold:

- (i) B and S are topological spaces.
- (ii) Z is an abelian topological group.
- (iii) The action $\alpha : Z \times B \to B$ is continuous.
- (iv) A set $U \subset S$ is open if and only if $\pi^{-1}(U)$ is open in B.

We say that B is a *compact* abelian bundle if, in addition to these conditions, B, S, Z are compact spaces.

Remark 2.1.7. Recall that by [8, Definition 3.2.17] the projection π yields a bijection between S and the orbits of α , namely the bijection $s \mapsto \pi^{-1}(s)$. It follows that we obtain a definition equivalent to Definition 2.1.6 if we replace condition (iv) with the following:

(iv') Denoting by B/ α the set of Z-orbits equipped with the quotient topology, we have that the map $S \to B/\alpha$, $s \mapsto \pi^{-1}(s)$ is a homeomorphism.

Thus the base *S* is topologically identified with the orbit space for the Z-action on B. In particular, since the orbit map $B \to B/\alpha$, $x \mapsto x + Z$ is an open map, we have that the projection π is an open map. Note also that if Z is compact, then the orbit map is also a closed map, and hence so is π .

Definition 2.1.8. Let $X = B_k$ be a k-fold abelian bundle, with factors B_i , i = 0, ..., k. We say that X is a k-fold compact abelian bundle if for each $i \in [k]$ the factor B_i is a compact Z_i -bundle with base B_{i-1} . A compact degree-k bundle is a k-fold compact abelian bundle that is also a compact cubespace with respect to the same topology, and such that condition (3.5) from [8, Definition 3.2.18] is satisfied, that is, for every $i \in [0, k-1]$ and every $n \in \mathbb{N}$, we have $C^n(B_i) = \{\pi_i \circ c : c \in C^n(X)\}$, and for every $c \in C^n(B_{i+1})$ we have

$$\{c_2 \in C^n(B_{i+1}) : \pi_i \circ c = \pi_i \circ c_2\} = \{c + c_3 : c_3 \in C^n(\mathcal{D}_{i+1}(Z_{i+1}))\}. \tag{2.1}$$

We can now establish the main result of this subsection.

Proposition 2.1.9. A compact cubespace X is a compact degree-k bundle if and only if X is a k-step compact nilspace.

Proof. The algebraic part of the statement is given by [8, Theorem 3.2.19], so we only need to check the topological properties in each direction.

If X is a compact degree-k bundle then in particular it is a compact cubespace satisfying the nilspace axioms, so it is a k-step compact nilspace. Conversely, if X is a k-step compact nilspace then, to show that the degree-k bundle structure given by [8, Theorem 3.2.19] is compact, it suffices to show that the abelian structure group Z_k can be equipped with a topology making it a compact abelian group such that X is a compact Z_k -bundle with base $\mathcal{F}_{k-1}(X)$ (if we show this then the result follows by repeating the argument for Z_{k-1} and so on by induction). For this it suffices to show that Z_k can be equipped with a

compact topology such that the action of Z_k on X is continuous (since Lemma 2.1.5 then implies the remaining condition (iv') in Definition 2.1.6). Let F be a class of \sim_{k-1} in X. From the proof of [8, Corollary 3.2.16] we know that F together with the F-valued cubes in each $C^n(X)$ is a k-fold ergodic k-step nilspace. Using [8, Lemma 3.2.4] we deduce that F is closed and it follows that F is a compact nilspace. Then by Lemma 2.1.4 there is a compact topology on Z_k such that F is isomorphic to $\mathcal{D}_k(Z_k)$ as a compact nilspace. We now check that the action of Z_k is continuous on F by showing that its graph $\{(a,y_0,y_0+a):a\in Z,y_0\in X\}$ is closed in $Z_k\times X^2$. This graph can be described using the arrow space $Y=X\bowtie_1 X$. Indeed, recall from [8, Lemma 3.2.15] that for $x=(x_0,x_1),y=(y_0,y_1)$ in $F^2\subset Y$ we have $x\sim_{k-1} y$ if and only if $x_0-x_1=y_0-y_1$ in Z_k . It follows that, fixing some $e\in F$, the graph above is homeomorphic to the closed set $\{(x,y)\in F^2\times F^2:x\sim_{k-1} y,x_0=e\}$. We have thus shown that from any given fibre F the group Z_k acquires a compatible compact topology making its action on F continuous. To see that this topology is the same for every fibre, recall that by [8, Lemma 3.2.24] for any two fibres F_0, F_1 and fixed points $x_i\in F_i$, we have an isomorphism $\vartheta:Z_{F_0}\to Z_{F_1}$, defined by $\vartheta(a)=b$ if and only if $(x_0,x_1)\sim_{k-1}(x_0+a,x_1+b)$. We then see that ϑ is continuous by identifying its graph with the closed set $\{y\in F_0\times F_1:y\sim_{k-1}(x_0,x_1)\}$. It follows that ϑ is a homeomorphism.

Recall the notions of a bundle morphism [8, Definition 3.3.1] and of a sub-bundle [8, Definition 3.3.3]. These notions are transferred to the setting of compact abelian bundles in a straightforward manner, by specifying in their definitions that all the maps involved are continuous. The same holds for the kernel of a bundle morphism $\psi : B \to B'$ (recall [8, Definition 3.3.4]). Note that such a kernel K is indeed a compact abelian bundle, since each set K_i in [8, Definition 3.3.4] is a closed subset of $B_i \times B'$ and therefore compact, and every structure group $\ker(\alpha_i)$ is compact.

Recall also that in the purely algebraic setting we had a description of a set of restricted morphisms $hom_f(P,X)$ as a sub-bundle of X^P , namely [8, Lemma 3.3.11]. This has the following version for compact nilspaces.

Lemma 2.1.10. Let P be a subcubespace of $\{0,1\}^n$ with the extension property, let S be a subcubespace of P with the extension property in P, let X be a k-step compact nilspace and let $f: S \to X$ be a morphism. Then $\hom_f(P,X)$ is a compact k-fold abelian bundle that is a sub-bundle of X^P , with factors $\hom_{\pi_i \circ f}(P,X_i)$ and structure groups $\hom_{S \to 0}(P,\mathcal{D}_i(Z_i))$, where Z_i is the i-th structure group of X.

Proof. We just have to check the topological properties. The set $\hom_f(P, X)$ is closed in X^P and therefore compact. For each $i \in [k]$ the abelian group $\hom_{S \to 0}(P, \mathcal{D}_i(Z_i))$ is also compact, as a closed subgroup of Z_i^P , and the action of $\hom_{S \to 0}(P, \mathcal{D}_i(Z_i))$ inherits continuity from that of Z_i . Condition (iv) from Definition 2.1.6 is also inherited from X^P since the topology on each factor $\hom_{\pi_i \circ f}(P, X_i)$ is the subspace topology from X_i^P .

2.1.2 Metrics

By Remark 2.1.3 (ii), we know that a compact nilspace can always be equipped with a compatible metric. We shall often use the fact that this metric can be assumed to be invariant under the action of the last structure group. We record this fact as follows.

Lemma 2.1.11. Let X be a k-step compact nilspace and let d_0 be a metric on X generating its topology. Define $d(x,y) = \int_{z \in Z_k} d_0(x+z,y+z) d\mu_{Z_k}(z)$, for any $x,y \in X$. Then d is a Z_k -invariant metric that generates the same topology on X.

NOTES ON COMPACT NILSPACES

This invariant metric is a basic construction in the theory of G-spaces (see [33, Proposition 1.1.12]).

Proof. It is clear that d is a \mathbb{Z}_k -invariant metric. To see that d generates the same topology as d_0 , note first that the function $\mathbb{X} \times \mathbb{X} \to \mathbb{R}_{\geq 0}$, $(x,y) \mapsto d(x,y)$ is continuous and in particular for a fixed $x \in \mathbb{X}$ we have that $y \mapsto d(x,y)$ is continuous (relative to d_0). It follows that the topology generated by d is a subtopology of the given topology on \mathbb{X} . In particular the former topology is compact. However, that topology is also Hausdorff, so it is maximal among the compact topologies on \mathbb{X} , and so it must be the original topology on \mathbb{X} (see [6, §9.4, Corollary 3]).

From now on we always assume that a metric on X is Z_k -invariant in this sense. Given such a metric d on X, we may define the following quotient metric on $\mathcal{F}_{k-1}(X)$ (see [33, Proposition 1.1.12]):

$$d'(x',y') = \inf \{ d(x,y) : x,y \in X, \ \pi_{k-1}(x) = x', \ \pi_{k-1}(y) = y' \}.$$
 (2.2)

In the sequel we shall use the following fact several times.

Lemma 2.1.12. Let X be a k-step compact nilspace, and let $Cor^{k+1}(X)$ denote the set of (k+1)-corners on X with the subspace topology obtained from the product topology on $X^{\{0,1\}^{k+1}\setminus\{1^{k+1}\}}$. Let $K: Cor^{k+1}(X) \to X$ denote the map sending a (k+1)-corner c' to $c(1^{k+1})$, where c is the unique completion of c'. Then K is continuous.

Proof. We argue by contradiction. If K were not continuous then there would exist some corner c', some open set $U \ni K(c')$ and a sequence of corners c'_n converging to c' such that for all n we have $K(c'_n) \not\in U$. Let c_n denote the completion in $C^{k+1}(X)$ of c'_n . By the compactness of $C^{k+1}(X)$ there exists a subsequence (c_m) of (c_n) such that c_m converges to some cube $c \in C^{k+1}(X)$. This convergence combined with the assumption that $c'_n \to c'$ implies that for each $v \neq 1^{k+1}$ we have c(v) = c'(v), so by uniqueness of completion we have $c(1^{k+1}) = K(c') \in U$. On the other hand, the convergence $c_m \to c$ also implies that $c_m(1^{k+1}) \to c(1^{k+1})$. Since by assumption $c_m(1^{k+1}) = K(c'_m) \not\in U$ for all m, we have $c(1^{k+1}) \not\in U$, a contradiction. □

Let us record also the following consequence of Lemma 2.1.12, to the effect that proximate cubes on $\mathcal{F}_{k-1}(X)$ can always be lifted to proximate cubes on X.

Lemma 2.1.13. Let X be a k-step compact nilspace and let d, d' be the metrics on X, $\mathcal{F}_{k-1}(X)$ given by Lemma 2.1.11 and (2.2) respectively. Then for every $\varepsilon > 0$ there exists $\delta > 0$ such that the following holds. If $c_1, c_2 \in C^{k+1}(\mathcal{F}_{k-1}(X))$ satisfy $d'(c_1(v), c_2(v)) \leq \delta$ for all v, then there exist cubes $\tilde{c}_1, \tilde{c}_2 \in C^{k+1}(X)$ such that $\pi_{k-1} \circ \tilde{c}_i = c_i$ for i = 1, 2 and $d(\tilde{c}_1(v), \tilde{c}_2(v)) \leq \varepsilon$ for all v.

Proof. Let $P = \{0,1\}^{k+1} \setminus \{1^{k+1}\}$ and note that, by definition of d' and [8, Remark 3.2.12], the corners $c_i|_P$ can be lifted to corners c_i' , i=1,2 such that $d\left(c_1'(v),c_2'(v)\right) \leq \delta$ for every $v \in P$. By Lemma 2.1.12, if $\delta \leq \varepsilon$ is sufficiently small then the unique completions \tilde{c}_i of c_i' satisfy $d\left(\tilde{c}_1(1^{k+1}),\tilde{c}_2(1^{k+1})\right) \leq \varepsilon$. \square

2.2 Measure-theoretic preliminaries

In this section we collect the main tools from measure theory that we shall use in the sequel.

A notion that unifies most of these tools is that of a *continuous system of measures*, discussed in Subsection 2.2.1.

Using continuous systems of measures, we shall then define in Subsection 2.2.2 an invariant measure on a continuous abelian bundle, and use this to construct a generalization of the Haar measure for compact nilspaces.

In Subsection 2.2.3 we collect several constructions of probability spaces that will be used repeatedly later in the chapter.

2.2.1 Continuous systems of measures

Recall that a measure μ on a measure space X is said to be *concentrated on* $S \subset X$ if $\mu(X \setminus S) = 0$.

The study of continuous systems of measures goes back at least to Bourbaki (see [7, Ch. V, §3]). A recent treatment is given in [10]. We shall use the following definition.

Definition 2.2.1. Let X,Y be compact spaces and let $\pi: X \to Y$ be continuous. A *continuous system of measures* (CSM) on the map π is a family of Borel measures $\{\mu_y: y \in Y\}$ on X satisfying the following conditions:

- (i) For every $y \in Y$ the measure μ_y is concentrated on $\pi^{-1}(y)$.
- (ii) For every continuous function $f: X \to \mathbb{C}$, the function $Y \to \mathbb{C}$, $y \mapsto \int_{\pi^{-1}(y)} f \, d\mu_y$ is continuous.

As explained in [10], replacing \mathbb{C} by \mathbb{R} in condition (ii) yields an equivalent definition.

Note that all the CSMs that we shall consider in these notes consist of strictly positive measures μ_y , i.e. Borel measures taking positive values on non-empty open sets.

A simple example of a CSM is obtained by taking a product space, as follows.

Lemma 2.2.2. Let V and Y be compact spaces, let μ be a Borel probability measure on V, let $X = V \times Y$ and let $\pi: X \to Y$, $(v,y) \mapsto y$. For each $y \in Y$ let μ_y denote the measure on X concentrated on $\pi^{-1}(y)$, where it equals the pushforward of μ under the map $V \to \pi^{-1}(y)$, $v \mapsto (v,y)$. Then $\{\mu_y : y \in Y\}$ is a CSM on π .

Proof. Let $f: V \times Y \to \mathbb{R}$ be a continuous function. Then f is a uniform limit of functions of the form $F: (v,y) \mapsto \sum_{i=1}^n f_{i,1}(v) f_{i,2}(y)$ where $f_{i,1}: V \to \mathbb{R}$, $f_{i,2}: Y \to \mathbb{R}$ are continuous. Indeed, the algebra of such functions F satisfies the assumptions of the Stone-Weierstrass theorem [37, Appendix A, §A14]. (The fact that this algebra separates points follows from Urysohn's lemma [32, Theorem 33.1].) Now, for a function $f_{i,1}(v) f_{i,2}(y)$ as above, it is clear that integrating over v yields a continuous function of y. It follows that $y \mapsto \int_{\pi^{-1}(y)} f d\mu_y$ is a uniform limit of continuous functions and is therefore continuous. \square

One of our main uses of CSMs below is to equip any compact abelian bundle with a probability measure that is invariant under the action of the structure group. This measure will be obtained as a simple application of the following result concerning general CSMs.

Lemma 2.2.3. Let $\{\mu_y : y \in Y\}$ be a family of Borel probability measures on X forming a CSM on $\pi : X \to Y$, and let V be a Borel probability measure on Y. Then the following function on Borel sets $E \subset X$ is a Borel probability measure:

$$\mu(E) = \int_{Y} \mu_{y} \left(E \cap \pi^{-1}(y) \right) d\nu(y).$$

Proof. This is a special case of [10, Corollary 3.7], using [10, Proposition 2.23].

2.2.2 Haar measure on compact abelian bundles and nilspaces

Recall that every compact abelian group can be equipped with a regular Borel probability measure that is invariant under translation, called the (normalized) Haar measure [37]. In this subsection we shall construct a generalization of this measure for compact nilspaces. As we shall see, there is a natural inductive construction of a probability measure on a *k*-fold compact abelian bundle, starting with the Haar measure on a compact abelian group. This construction will yield the desired measure on a *k*-step compact nilspace, thanks to [8, Theorem 3.2.19].

Our first step, then, is to define a Haar measure on a compact abelian bundle.

Lemma 2.2.4. Let B be a compact abelian bundle with base S, structure group Z and projection π . Let μ_S be a regular Borel probability measure on S. Then there is a unique regular Borel probability measure μ on B that is invariant under the action of Z and satisfies $\mu_S = \mu \circ \pi^{-1}$.

We call μ the *Haar measure* on B (given μ_S).

Proof. For each $s \in S$ the set $\pi^{-1}(s)$ is a principal homogeneous space of Z, so there is a Z-equivariant homeomorphism $f: Z \to \pi^{-1}(s)$. Let μ_s be the pushfoward measure $\mu_Z \circ f^{-1}$ on $\pi^{-1}(s)$, where μ_Z is the Haar measure on Z. If we show that $\{\mu_s: s \in S\}$ is a CSM on π , then by Lemma 2.2.3 we can define a Borel measure μ on B by the following formula:

$$\mu(E) = \int_{S} \mu_s \left(\pi^{-1}(s) \cap E \right) d\mu_S(s), \tag{2.3}$$

for every Borel set $E \subset B$.

To show that $\{\mu_s: s \in S\}$ is a CSM on π we check condition (ii) from Definition 2.2.1. Given any continuous function $f: B \to \mathbb{R}$, we need to show that $g: S \to \mathbb{R}$, $s \mapsto \int_{\pi^{-1}(s)} f \, d\mu_s$ is continuous. Since S has the quotient topology (Definition 2.1.6), it suffices to show that $g \circ \pi: B \to \mathbb{R}$ is continuous. Now $g(\pi(b)) = \int_Z f(a+b) \, d\mu_Z(a) = \int_Z f'(a,b) \, d\mu_Z(a)$, where the function $f': Z \times B \to \mathbb{R}$, $(a,b) \mapsto f(a+b)$ is continuous. Therefore, the continuity of $g \circ \pi$ follows from (the proof of) Lemma 2.2.2 applied with V = Z and Y = B.

We thus obtain the Borel measure μ given by (2.3), and μ is regular by the standard fact that any Borel probability measure on a metric space is regular [5, Theorem 1.1]. The Z-invariance of μ clearly follows from that of each measure μ_s . The uniqueness can be deduced from uniqueness of Haar measure on Z combined with a disintegration result for invariant measures, such as [36, Proposition 2].

As mentioned above, this lemma yields the following result.

Proposition 2.2.5. Let B be a compact k-fold abelian bundle, with factors $B_0, B_1, \ldots, B_k = B$ and structure groups Z_1, \ldots, Z_k . Then there is a unique regular Borel probability measure μ on B with the following property: for each $i \in [k]$, letting π_i be the projection $B \to B_i$, we have that the Borel probability measure $\mu \circ \pi_i^{-1}$ on B_i is invariant under the action of Z_i .

In particular, for every *k*-step compact nilspace X, applying the proposition to the bundle structure we obtain the measure that we shall call the *Haar measure on* X.

Proof. By induction on $i \in [k]$, starting with the Haar probability on \mathbb{Z}_1 and applying Lemma 2.2.4 for $i = 2, \dots, k$.

Recall from [8, Definition 3.3.1] that a bundle morphism $\phi : B \to B'$ between k-fold abelian bundles is said to be *totally surjective* if the structure morphism $\alpha_i : Z_i \to Z_i'$ is surjective for each $i \in [k]$.

The following basic result is the analogue for nilspaces of the fact that continuous surjective homomorphisms between compact abelian groups preserve the Haar measures.

Lemma 2.2.6. Let B,B' be compact k-fold abelian bundles, and let $\phi : B \to B'$ be a totally surjective continuous bundle morphism. Then ϕ preserves the Haar measures, that is, for every Borel set $E \subset B'$ we have $\mu_B(\phi^{-1}(E)) = \mu_{B'}(E)$.

Proof. We argue by induction on k. The case k=1 follows from the fact recalled above. For k>1, recall that by [8, Definition 3.3.1] the map ϕ induces a totally surjective continuous bundle morphism $\phi_{k-1}: B_{k-1} \to B'_{k-1}$. Let $\{\mu_s: s \in B_{k-1}\}$ and $\{\mu'_{s'}: s' \in B'_{k-1}\}$ be the CSMs on $\pi_{k-1}: B \to B_{k-1}$, $\pi'_{k-1}: B' \to B'_{k-1}$ respectively (as in the proof of Lemma 2.2.4). For $s \in B_{k-1}$ and any Borel set $E \subset B'$, let $f_B(s) = \mu_s(\pi^{-1}_{k-1}(s) \cap \phi^{-1}(E))$, and for $s' \in B'_{k-1}$ let $f_{B'}(s') = \mu'_{s'}(\pi'^{-1}_{k-1}(s') \cap E)$. The restriction of ϕ to a fibre $\pi^{-1}_{k-1}(s)$, being a continuous affine homomorphism onto some fibre $\pi'^{-1}_{k-1}(s')$, preserves the measures $\mu_s, \mu'_{s'}$. Hence we have $f_B(s) = f_{B'}(\phi_{k-1}(s))$ for each $s \in B_{k-1}$. Since ϕ_{k-1} is measure-preserving (by the induction hypothesis), we have

$$\mu_{\mathrm{B}}(\phi^{-1}(E)) = \int_{\mathrm{B}_{k-1}} f_{\mathrm{B}} \, \mathrm{d}\mu_{\mathrm{B}_{k-1}} = \int_{\mathrm{B}_{k-1}} f_{\mathrm{B}'} \circ \phi_{k-1} \, \mathrm{d}\mu_{\mathrm{B}_{k-1}} = \int_{\mathrm{B}'_{k-1}} f_{\mathrm{B}'} \, \mathrm{d}\mu_{\mathrm{B}'_{k-1}} = \mu_{\mathrm{B}'}(E). \qquad \Box$$

Recall from [8, Definition 3.3.7] the notion of a fibre-surjective morphism.

Corollary 2.2.7. *Let* X, X' *be k-step compact nilspaces and let* $\phi : X \to X'$ *be a continuous fibre-surjective morphism. Then* ϕ *preserves the Haar measures.*

Proof. We combine Lemma 2.2.6 with [8, Lemma 3.3.8].

In the sequel we shall have to use CSMs on certain maps from cube sets $C^n(X)$ to X, such as the evaluation map $c \mapsto c(0^n)$. These CSMs will be obtained as special cases of the construction in the next lemma.

Recall from [8, Definition 3.2.17] the notion of a *relative k*-fold abelian bundle. We may similarly define a *compact* relative *k*-fold abelian bundle B with ground set B₀. Note that each fibre of the projection $\pi_0: B \to B_0$ is then itself a compact *k*-fold abelian bundle and so it has a Haar measure by Lemma 2.2.4.

Lemma 2.2.8. Let B be a compact relative k-fold abelian bundle with factors B_0, B_1, \dots, B_{k-1} . Then the Haar measures on the fibres of the projection $\pi_0 : B \to B_0$ form a CSM on π_0 .

To prove this we shall have to compose several CSMs, in the following sense.

Definition 2.2.9. Given CSMs $\{\mu_y : y \in Y\}$ on $\pi : X \to Y$ and $\{\mu_z : z \in Z\}$ on $\tau : Y \to Z$, their *composition* is the CSM $\{v_z : z \in Z\}$ on $\tau \circ \pi : X \to Z$ consisting of measures defined for Borel sets $E \subset X$ by

$$v_z(E) = \int_{\tau^{-1}(z)} \mu_y (\pi^{-1}(y) \cap E) d\mu_z(y).$$
 (2.4)

The fact that this composition is indeed a CSM is established in [10, Proposition 3.3].

Proof of Lemma 2.2.8. By the proof of Lemma 2.2.4, for each $i \in [k]$ there is a CSM on the bundle map $\pi_{i-1,i}: B_i \to B_{i-1}$ consisting of the pushforward of the Haar measure on Z_i to each fibre of $\pi_{i-1,i}$. Let us compose these CSMs for $i \in [k]$, and let $\{\mu_x : x \in B_0\}$ be the resulting CSM on π_0 . For each $x \in B_0$, a straightforward inductive argument using (2.4) shows that the measure μ_x on the compact k-fold abelian bundle $\pi^{-1}(x)$ satisfies the properties of the measure in Proposition 2.2.5, so it must be the Haar measure on $\pi^{-1}(x)$.

The following result is an analogue for compact abelian bundles of the quotient integral formula.

Lemma 2.2.10. Let B,B' be compact k-fold abelian bundles, with Haar measures μ , μ' respectively, and let ϕ : B \to B' be a totally surjective continuous bundle morphism. For each $t \in B'$ let μ_t denote the Haar measure on $\phi^{-1}(t)$. Then the measures μ_t , $t \in B'$ disintegrate μ relative to μ' , that is for every Borel set $E \subset B$ we have $\mu(E) = \int_{B'} \mu_t(\phi^{-1}(t) \cap E) d\mu'(t)$.

Recall from [8, Definition 3.3.4] the notion of the kernel of a bundle morphism.

Proof. The kernel of ϕ is a compact relative k-fold abelian bundle, with ground set B', and the fibres $\phi^{-1}(t)$ are compact k-fold abelian bundles (recall [8, Lemma 3.3.6]). By Lemma 2.2.8 the Haar measures μ_t form a CSM on ϕ . By Lemma 2.2.3, with this CSM and μ' we can define a Borel probability v on B by

$$v(E) = \int_{B'} \mu_t(\phi^{-1}(t) \cap E) d\mu'(t).$$
 (2.5)

We have to show that $v = \mu$. For this it suffices to show that v has the invariance properties that characterize the Haar measure μ , namely that for each projection $\pi_i : B \to B_i$ the measure $v \circ \pi_i^{-1}$ on B_i is Z_i -invariant (recall Proposition 2.2.5). We prove this by induction on k.

For k = 1, we have that B,B' are principal homogeneous spaces of compact abelian groups Z,Z' respectively, and we just have to show that ν is Z-invariant. But from the quotient integral formula [12, Theorem 1.5.2] it actually follows that $\nu = \mu$ in this case, so ν is indeed Z-invariant.

Supposing now that the claim holds for $k \ge 1$, we prove it for k+1. First we show that we can suppose by induction that $v \circ \pi_i^{-1}$ is already Z_i -invariant on B_i for each $i \in [k]$. For this it suffices to show that $v \circ \pi_k^{-1}$ is of the same form as v (as in (2.5)), but relative to the totally surjective morphism $\phi_k : B_k \to B'_k$ induced by $\phi : B_{k+1} \to B'_{k+1}$. By construction, the Haar measure μ' has a disintegration relative to the Haar measure μ'_k on B'_k , into measures μ'_s , $s \in B'_k$, each of which is the Haar measure on Z'_{k+1} pushed forward to the fibre ${\pi'_k}^{-1}(s)$. Thus for any Borel set $E \subset B_k$ we have

$$\mathbf{v} \circ \pi_k^{-1}(E) \ = \ \int_{\mathbf{B}'} \mu_t \left(\pi_k^{-1}(E) \cap \phi^{-1}(t) \right) \mathrm{d}\mu'(t) \ = \ \int_{\mathbf{B}'_k} \left(\int_{\pi_k'^{-1}(s)} \mu_t \left(\pi_k^{-1}(E) \cap \phi^{-1}(t) \right) \mathrm{d}\mu_s'(s) \right) \mathrm{d}\mu_k'(s).$$

For each $s \in B'_k$ and each $t \in B'$ with $\pi'_k(t) = s$, note that $\phi^{-1}(t)$ is a (k+1)-fold compact abelian bundle, with k-th factor $\phi_k^{-1}(s)$ having Haar measure denoted μ_s . Then π_k restricted to $\phi^{-1}(t)$ pushes μ_t forward to μ_s . It follows that for any such s,t we have $\mu_t(\pi_k^{-1}(E) \cap \phi^{-1}(t)) = \mu_s(E \cap \phi_k^{-1}(s))$. Hence

$$v \circ \pi_k^{-1}(E) = \int_{B_k'} \left(\int_{\pi_k'^{-1}(s)} \mu_s \left(E \cap \phi_k^{-1}(s) \right) d\mu_s'(t) \right) d\mu_k'(s) = \int_{B_k'} \mu_s \left(E \cap \phi_k^{-1}(s) \right) d\mu_k'(s).$$

The right side here is indeed the measure on B_k constructed in the same way as v in (2.5). Therefore, as mentioned above, by induction it now suffices to show that v is Z_{k+1} -invariant.

Again we use the disintegration of μ' relative to μ'_{k-1} , so that for any Borel set $E \subset B'$ we have

$$\nu(E) = \int_{\mathbf{B}'_k} \left(\int_{\pi'_k^{-1}(s)} \mu_t \big(E \cap \phi^{-1}(t) \big) \, \mathrm{d}\mu'_s(t) \right) \mathrm{d}\mu'_k(s).$$

For each $s \in B'_k$, we claim that the inner integral on the right side here integrates 1_E over a union of orbits of Z_{k+1} . Indeed, we have $\mu_t(E \cap \phi^{-1}(t)) = \int_{\phi^{-1}(t)} 1_E(x) d\mu_t(x)$, and by disintegrating μ_t into the fibres of $\pi_k : \phi^{-1}(t) \to \phi_k^{-1}(s)$, the inner integral in question is written

$$\int_{\pi_{k}^{\prime-1}(s)} \int_{\phi_{k}^{-1}(s)} \int_{\pi_{k}^{-1}(y) \cap \phi^{-1}(t)} 1_{E}(x) d\mu_{\pi_{k}^{-1}(y)}(x) d\mu_{s}(y) d\mu'_{s}(t).$$

By Fubini's theorem we can interchange the two outer integrals, so this equals

$$\int_{\phi_k^{-1}(s)} \int_{\pi_k'^{-1}(s)} \int_{\pi_k^{-1}(y) \cap \phi^{-1}(t)} 1_E(x) \ d\mu_{\pi_k^{-1}(y)}(x) \ d\mu_s'(t) \ d\mu_s(y).$$

For each $y \in \phi_k^{-1}(s)$, note that $\pi_k'^{-1}(s)$ is a Z_{k+1}' -torsor (or principal homogeneous space of Z_{k+1}') and for every t in this torsor we have that $\pi_k^{-1}(y) \cap \phi^{-1}(t)$ is a $\ker(\alpha_{k+1})$ -torsor, where α_{k+1} is the (k+1)-th structure morphism of ϕ (recall [8, Definition 3.3.1 and Lemma 3.3.6]). Fixing any t_0 in the former torsor, and then for each $r \in Z_{k+1}'$ fixing some $x_r \in \pi_k^{-1}(y) \cap \phi^{-1}(t_0 + r)$, the last inner double-integral equals

$$\int_{Z'_{k+1}} \int_{\pi_k^{-1}(y) \cap \phi^{-1}(t_0+r)} 1_E(x) d\mu_{\pi_k^{-1}(y)}(x) d\mu_{Z'_{k+1}}(r) = \int_{Z'_{k+1}} \int_{\ker(\alpha_{k+1})} 1_E(x_r+u) d\mu_{\ker(\alpha_{k+1})}(u) d\mu_{Z'_{k+1}}(r).$$

Note that the sets $\{x_r + u : u \in \ker(\alpha_{k+1})\}$, $r \in Z'_{k+1}$ form a partition of $\pi_k^{-1}(y)$. It follows from the quotient integral formula [12, Theorem 1.5.2] that for any $x_0 \in \pi_k^{-1}(y)$ this integral equals

$$\int_{Z_{k+1}} 1_E(x_0 + z) d\mu_{Z_{k+1}}(z) = \int_{\pi_k^{-1}(y)} 1_E(x) d\mu_{\pi_k^{-1}(y)}(x).$$

This is clearly invariant under shifting E by any fixed $z \in \mathbb{Z}_{k+1}$. The \mathbb{Z}_{k+1} -invariance of ν follows. \square

Recall the fact that the Haar measure on a compact abelian group is strictly positive. This generalizes to compact nilspaces.

Proposition 2.2.11. *Let* X *be a k-step compact nilspace, with Haar measure* μ . *Then for every open set* $U \subset X$ *we have* $\mu(U) > 0$.

Proof. We argue by induction on k. The case k=1 follows from the fact recalled above. Let k>1 and let U be open in X. By Remark 2.1.7 we have that π_{k-1} is an open map, so $B=\pi_{k-1}(U)$ is open in $\mathcal{F}_{k-1}(X)$. Hence, letting μ' denote the Haar measure on $\mathcal{F}_{k-1}(X)$, we have by induction that $\mu'(B)>0$. By regularity of μ' there is a compact subset $C\subset B$ such that $\mu'(C)>0$, and then since π_{k-1} preserves Haar measures we have $\mu(\pi_{k-1}^{-1}(C))>0$. Now by Remark 2.1.7 we also have that π_{k-1} is a proper map, so $\pi_{k-1}^{-1}(C)$ is a compact set. This set is covered by the union of open sets $\bigcup_{z\in Z_k} U+z$, so there is a finite set of translates of U by elements of Z_k that covers $\pi_{k-1}^{-1}(C)$, and it follows that $\mu(U)>0$.

2.2.3 Probability spaces of morphisms

Recall from [8, Definition 3.1.3] that given a cubespace P and a subcubespace S of P, we say that S has the *extension property* in P if for every non-empty nilspace X and every morphism $g': S \to X$ there is a morphism $g: P \to X$ with $g|_S = g'$.

In this subsection we collect several constructions of probability spaces of morphisms between nilspaces, which will be used repeatedly in the sequel. These constructions rely on two basic ideas.

Firstly, every set of restricted morphisms $hom_f(P, X)$ has a compact abelian-bundle structure, as described in Lemma 2.1.10, which we restate here for convenience.

Lemma 2.2.12. Let P be a subcubespace of $\{0,1\}^n$ with the extension property, let S be a subcubespace of P with the extension property in P, let X be a k-step compact nilspace and let $f: S \to X$ be a morphism. Then $\hom_f(P,X)$ is a compact k-fold abelian bundle that is a sub-bundle of X^P , with factors $\hom_{\pi_i \circ f}(P,X_i)$ and structure groups $\hom_{S \to 0}(P,\mathcal{D}_i(Z_i))$, where Z_i is the i-th structure group of X.

As a consequence of this structure we have a Haar measure on $hom_f(P,X)$, by Proposition 2.2.5.

The second idea consists in giving a simple description of the measure-theoretic properties of a restriction map from one set of such morphisms to another. In particular, we want a convenient criterion to decide whether, for a subcubespace S of P, the restriction map $hom(P,X) \to hom(S,X)$ preserves the Haar measures. The following notion provides a general criterion of this type.

Definition 2.2.13. Let P be a cubespace, and let P_1, P_2 be subcubespaces of P. We call the pair of sets P_1, P_2 a *good pair* if P_1 and $P_1 \cap P_2$ both have the extension property in P and, for every abelian group P and every P and P and P and P and P are P and P and P and P and P are P and P and P are P are P and P are P and P are P and P are P are P and P are P are P and P are P and P are P are P and P are P are P and P are P and P are P are P and P are P are P and P are P and P are P are P and P are P are P and P are P and P are P and P are P and P are P are P and P are P are P are P and P are P and P are P and P are P and P are P are P and P are P and P are P are P and P are P are P are P and P are P are P are P are P are P are P and P are P are P and P are P and P are P and P are P are P are P are P are P and P are P are P are P are P are P and P are P and P are P are P are P are P and P are P are

Note that in particular if *S* has the extension property in $P \subset \{0,1\}^n$ and $P_1 = \emptyset$ then P_1, S is a good pair in *P*. The main purpose of this definition is the following result.

Lemma 2.2.14. Let $P \subset \{0,1\}^n$ be a subcubespace with the extension property, and let P_1, P_2 be a good pair in P. Let X be a k-step nilspace and let $f: P_1 \to X$ be a morphism. Then the restriction map

$$\phi: \hom_f(P, X) \to \hom_{f|_{P_1 \cap P_2}}(P_2, X)$$
(2.6)

is a totally-surjective bundle morphism. In particular, if X is a k-step compact nilspace, then ϕ preserves the Haar measures. Moreover, the Haar measures on the fibres $\phi^{-1}(t)$ form a CSM on ϕ that disintegrates the Haar measure on $\hom_f(P,X)$ relative to the Haar measure on $\hom_{f|_{P_1\cap P_2}}(P_2,X)$.

P. CANDELA

Proof. We prove that ϕ is a totally-surjective bundle morphism by induction on k. For k=0 the claim holds trivially, so let $k \geq 1$ and suppose that X is a k-step nilspace and that the lemma holds for (k-1)-step nilspaces. To see that condition (i) from [8, Definition 3.3.1] holds, note that by a straightforward calculation we have that ϕ induces a map ϕ_{k-1} : hom $_{\pi_{k-1}\circ f}(P,X_{k-1})\to \text{hom}_{\pi_{k-1}\circ f|_{P_1\cap P_2}}(P_2,X_{k-1})$ well-defined by $\pi_{k-1}\circ g\mapsto \pi_{k-1}\circ\phi(g)$, so the condition holds for i=k-1. We also have that ϕ_{k-1} is precisely the restriction map on $\text{hom}_{\pi_{k-1}\circ f}(P,X_{k-1})$, so by induction the condition holds for all $i\leq k-1$. Condition (ii) from the same definition is clearly satisfied, with the structure morphism α_i being the restriction $\text{hom}_{P_1\to 0}(P,\mathcal{D}_i(Z_i))\to \text{hom}_{P_1\cap P_2\to 0}(P_2,\mathcal{D}_i(Z_i))$. To check that ϕ is totally surjective, by induction it suffices to check that α_k is surjective, which holds by Definition 2.2.13. Lemma 2.2.6 then implies that ϕ preserves the Haar measures. The last sentence of the lemma follows from Lemma 2.2.10.

Another useful feature of good pairs is the following extension result for morphisms.

Lemma 2.2.15. Let P be a subcubespace of $\{0,1\}^n$ with the extension property, and let P_1, P_2 be a good pair in P. Then $P_1 \cup P_2$ equipped with the union of the cube structures on P_1, P_2 has the following extension property: any morphism $f: P_1 \cup P_2 \to X$ into a non-empty k-step nilspace X extends to a morphism $f': P \to X$.

Proof. We argue again by induction on k. The statement is trivial for k=0, as X is then a 1-point space, so suppose that the lemma holds for (k-1)-step nilspaces and let X be a k-step nilspace. Let $f: P_1 \cup P_2 \to X$ be a morphism and let $f_2: P \to \mathcal{F}_{k-1}(X)$ be an extension of $\pi_{k-1} \circ f$. By [8, Lemma 3.2.11] there is a morphism $f_2': P \to X$ such that $\pi_{k-1} \circ f_2' = f_2$. Let $g = f_2'|_{P_1} - f|_{P_1}$. By Definition 2.2.13, there is an extension $f_3: P \to \mathcal{D}_k(Z_k)$ of g. Let $g_2 = f_2' - f_3$. We have that $g_2|_{P_1} = f|_{P_1}$. Now let g_3 be an extension of $g_2|_{P_2} - f|_{P_2}$ to P with $g_3|_{P_1} = 0$. Then $f' = g_2 - g_3$ is an extension of f to f.

For the remainder of this section we apply the results above to define various probability spaces of morphisms. We begin with the cube sets $C^n(X)$ on a compact nilspace. Note that $C^n(X) = \text{hom}(\{0,1\}^n, X)$ is a compact k-fold abelian bundle and therefore has a Haar measure (by Proposition 2.2.5).

Lemma 2.2.16. Let X, Y be k-step compact nilspaces and let $n \in \mathbb{N}$. Then for every continuous fibre-surjective morphism $\psi : X \to Y$, the induced map $C^n(X) \to C^n(Y)$, $c \mapsto \psi \circ c$ preserves the Haar measures. Moreover, the fibres of this map can all be equipped with the Haar measure.

Proof. By [8, Lemma 3.3.12 (ii)] we have that $c \mapsto \psi \circ c$ is a totally surjective bundle morphism, so it preserves the Haar measures by Lemma 2.2.6. Moreover, by [8, Lemma 3.3.12 (iii)], the fibres of this map are compact k-fold abelian bundles, so the last claim in the lemma follows from Proposition 2.2.5.

We shall often want to use a disintegration of the Haar measure on $C^n(X)$ into a CSM on the projection $c \mapsto c(0^n)$. For this purpose we record the following special case of Lemma 2.2.14.

Lemma 2.2.17. Let X be a k-step compact nilspace with Haar measure μ , let $\pi : C^n(X) \to X$, $c \mapsto c(0^n)$, and let μ_n denote the Haar measure on $C^n(X)$. Then for every $x \in X$, the space $\pi^{-1}(x) = C^n_x(X)$ is a compact k-fold abelian bundle and therefore has a Haar measure μ_x . The family $\{\mu_x : x \in X\}$ is a CSM on π disintegrating μ_n relative to μ .

Proof. By Lemma 2.2.14 with $P = \{0,1\}^n$, with $P_1 = \emptyset$, $P_2 = \{0^n\}$, identifying X with hom (P_2, X) .

NOTES ON COMPACT NILSPACES

We move on to some results concerning the tricube T_n . Recall from [8, Subsection 3.1.3] the definitions of tricubes, of morphisms Ψ_v , and of the outer-point map $\omega_n : \{0,1\}^n \to T_n$.

Lemma 2.2.18. The set $\omega_n(\{0,1\}^n) \subset T_n$ has the extension property in T_n .

Proof. Let X be a nilspace, let $P = \omega_n(\{0,1\}^n)$ and let $f: P \to X$ be a morphism. Let $h: \{-1,0,1\} \to \{0,1\}$ be the map $-1 \mapsto 0$, $0 \mapsto 1$, $1 \mapsto 1$. Then the map $f' = f \circ h^n$ is an extension of f to T_n .

Lemma 2.2.19. For every $v \in \{0,1\}^n$, the subcube $\Psi_v(\{0,1\}^n)$ of T_n has the extension property in T_n .

Proof. The claim is checked for $v = 0^n$ using the map $h : \{-1,0,1\} \to \{0,1\}, -1 \mapsto 1, 0 \mapsto 0, 1 \mapsto 1$. (Indeed then $f \circ h^n$ extends f to T_n .) The lemma then follows by the composition axiom.

Corollary 2.2.20. Let $P \subset T_n$ be $\Psi_v(\{0,1\}^n)$ or $\omega_n(\{0,1\}^n)$, let X be a k-step compact nilspace, and let $f: P \to X$ be a morphism. Then $\hom_f(T_n, X)$ has a Haar measure.

Proof. This follows again from Lemma 2.2.12 and Proposition 2.2.5.

Lemma 2.2.21. The sets $P_1 = \omega_n(\{0,1\}^n) = \{-1,1\}^n$ and $P_2 = \{0,1\}^n$ form a good pair in T_n .

Proof. Note first that we may assume that $\omega_n \in C^n(T_n)$, since by [8, Lemma 3.1.16] this assumption does not affect any set of morphisms from T_n into any nilspace. Now since $P_1 \cap P_2 = \{1^n\}$ is a singleton, it has the extension property. The set P_1 itself has the extension property, by Lemma 2.2.18. Let $f: \{0,1\}^n \to \mathcal{D}_k(Z)$ be a morphism with $f(1^n) = 0$. Let h(-1) = 1, h(0) = 0, h(1) = 1. Then $f' = f \circ h^n$ is an extension of f with $f|_{P_1} = 0$.

Combining this with Lemma 2.2.14, we obtain that a set of morphisms $T_n \to X$ with a constraint at the outer points of T_n is still large enough to cover a whole set of rooted cubes $C_r^n(X)$, in the following sense.

Corollary 2.2.22. Let X be a k-step compact nilspace, let P_1 and P_2 be as in Lemma 2.2.21, and let $f: P_1 \to X$ be a morphism. Then the restriction map $hom_f(T_n, X) \to hom_{0^n \to f(1^n)}(P_2, X) = C^n_{f(1^n)}(X)$, $t \mapsto t \circ \Psi_{0^n}$ preserves the Haar measures.

As a complement to this corollary, we also have that a set of morphisms $T_n \to X$ with a given constraint at just one outer point is always large enough to cover a full cube set $C^n(X)$, as follows.

Lemma 2.2.23. Let X be a k-step compact nilspace, let $x \in X$, and let Q_x denote the compact abelian bundle $\hom_{1^n \mapsto x}(T_n, X)$. Then for every $v \in \{0, 1\}^n \setminus \{0^n\}$, the map $Q_x \to C^n(X)$, $t \mapsto t \circ \Psi_v$ preserves the Haar measures.

Proof. There is a Haar measure on Q_x by Lemma 2.2.12. Let $P_1 = \{1^n\}$ and $P_2 = \Psi_v(\{0,1\}^n)$. The lemma will follow if we show that P_1, P_2 form a good pair in T_n , by Lemma 2.2.19. Since the singleton P_1 has the extension property in T_n and so does $P_1 \cap P_2 = \emptyset$, it suffices to show that any given morphism $f': P_2 \to \mathcal{D}_k(Z)$ can be extended to a morphism $f: T_n \to \mathcal{D}_k(Z)$ satisfying $f|_{P_1} = 0$. We can certainly extend f' to a morphism $f_2: T_n \to \mathcal{D}_k(Z)$, by Lemma 2.2.19. Now $v \neq 0^n$ implies that $\Psi_v(0^n) \neq 1^n$, so this has some coordinate equal to -1, say without loss of generality $\Psi_v(0^n)(1) = -1$. Then the set $F = \{1\} \times \{-1,0,1\}^{n-1}$ is disjoint from P_2 . Let $g: T_n \to \mathcal{D}_k(Z)$ equal $f_2(1^n)$ on F and 0 otherwise. A straightforward calculation shows that g is a morphism, using [8, (2.9)]. Then the function $f = f_2 - g$ is also a morphism $T_n \to \mathcal{D}_k(Z)$, it extends f', and satisfies $f|_{P_1} = 0$.

2.3 Topological spaces associated with continuous systems of measures

Definition 2.3.1. Let X,Y be compact spaces and let μ be a Borel measure on X. We denote by L(X,Y) the quotient of the set of Borel measurable functions $f: X \to Y$ by the equivalence relation \sim defined by $f \sim g \Leftrightarrow \mu(\{x \in X : f(x) \neq g(x)\}) = 0$.

In this section, given a continuous system of measures, we construct an associated topological space that will be crucial in the sequel. In order to motivate the construction, let us outline how it will be used.

The central application will be given in the next section, which is one of the main sections of this chapter. The aim there is to show that given a compact nilspace X and a compact abelian group Z, given a cocycle $\rho: C^k(X) \to Z$ (recall [8, Definition 3.3.14]), if ρ is Borel measurable then the extension $M(\rho)$ constructed in [8, Proposition 3.3.26] can be equipped with a compact nilspace structure compatible with that of X. The purpose of the construction in this section is to provide an ambient continuous abelian bundle from which $M(\rho)$ will inherit the desired topological structure. The construction will be completed with Proposition 2.3.13.

Let us now turn to the definition of the topological space in the construction. To motivate this, consider the following general theme of the application mentioned above: given a Borel measurable function on a space (e.g. the cocycle ρ), we have to build a topological space that is associated with the function in a useful way. A very basic example of such a construction is the following: let Z be a compact abelian group, suppose that we only know the Borel σ -algebra on Z (not the underlying topology), and that we are given a Borel function $f: \mathbb{Z} \to \mathbb{R}$. Then we can build a topology on Z naturally associated with f using the L^1 seminorm on the Borel functions on Z (relative to the Haar measure), by restricting this seminorm to the translates of f, i.e. the functions $f_z: x \mapsto f(x+z), z \in \mathbb{Z}$. More precisely, we can define a pseudometric d on Z by $d(z,z') = ||f_z - f_{z'}||_{L^1(Z)}$, and then take the topology generated by the open balls with respect to d. The central application mentioned above will use a more elaborate version of this idea. Indeed, recall from [8, (3.11)] that the extension generated by the cocycle ρ is $M(\rho) = \bigcup_{x \in X} {\{\rho_x + z : z \in Z\}}$, where for each x we denote by ρ_x the restriction of ρ to the set of rooted cubes $C_x^k(X) = \{c \in C^k(X) : c(0^k) = x\}$. To put a useful topology on $M(\rho)$, we shall first put an analogue of the L^1 -topology on each family of shifts $\rho_x + z$, $z \in \mathbb{Z}$, for each $x \in \mathbb{X}$, and we shall then want to tie together adequately these different topologies over different points x into a global topology. Now the Haar measures on sets $C_r^k(X)$ form a CSM, by Lemma 2.2.17. This leads to the following definition, which concerns CSMs more generally, and which provides a topology that will be adequate for our purposes.

Definition 2.3.2. Let V, W, Z be compact spaces, and let $\{\mu_w : w \in W\}$ be a family of strictly positive Borel probability measures forming a CSM on a continuous map $\pi : V \to W$. For each function f in $\bigcup_{w \in W} L(\pi^{-1}(w), Z)$, let $\tilde{\pi}(f)$ be the element $w \in W$ such that $f \in L(\pi^{-1}(w), Z)$. We define the topological space $\mathcal{L}(V, Z)$ to be the set $\bigcup_{w \in W} L(\pi^{-1}(w), Z)$ equipped with the coarsest topology making the following functions continuous:

$$\varphi_{F_1,F_2}: f \mapsto \int_{\pi^{-1}(\tilde{\pi}(f))} F_1(f(v)) F_2(v) d\mu_{\tilde{\pi}(f)}(v), \tag{2.7}$$

for every pair of continuous functions $F_1: Z \to \mathbb{C}$ and $F_2: V \to \mathbb{C}$.

We shall discuss below how this topology relates to the L^1 topology, but for now let us record some of its basic properties.

Proposition 2.3.3. The topological space $\mathcal{L}(V,Z)$ is regular Hausdorff and second-countable, and the map $\tilde{\pi}$ is continuous.

In particular $\mathcal{L}(V,Z)$ is metrizable, by Urysohn's theorem [32, Theorem 34.1].

Proof. We show first that $\tilde{\pi}$ is continuous. Fix a metric on W, and let U be any open subset of W. Then for any $w \in U$ there is an open ball $B_r(w) \subset U$ of radius r > 0. By Urysohn's lemma [32, Theorem 33.1] there exists a continuous function f_w with value 1 on the closure of $B_{r/2}(w)$ and value 0 outside $B_r(w)$. Let B'_w denote the open set $\{x \in W : f_w(x) > 0\}$, and note that $U = \bigcup_{w \in U} B'_w$. By definition of $\mathcal{L}(V, Z)$, for every $w \in U$ the function $\varphi_{1,f_w \circ \pi}$ is continuous. Since $\varphi_{1,f_w \circ \pi}(f) = f_w(\tilde{\pi}(f))$, this function is positive precisely on $\tilde{\pi}^{-1}(B'_w)$, so this set is open. Therefore $\tilde{\pi}^{-1}(U) = \bigcup_{w \in U} \tilde{\pi}^{-1}(B'_w)$ is an open set.

Next we show that $\mathcal{L}(V,Z)$ is Hausdorff. Let f,f' be distinct points in $\mathcal{L}(V,Z)$. We distinguish two cases. In the first case f,f' lie in different fibres of $\tilde{\pi}$. Then since W is Hausdorff there exist disjoint open sets U,U' containing $\tilde{\pi}(f)$, $\tilde{\pi}(f')$ respectively. The preimages of these sets under $\tilde{\pi}$ are open sets separating f,f'. In the second case we have f,f' in the same fibre $L(\pi^{-1}(w),Z)$. Then there are Borel functions f_0,f'_0 in the classes f,f' respectively such that the set $D=\{v:f_0(v)\neq f'_0(v)\}$ has $\mu_w(D)>0$. For some $\varepsilon<\mu_w(D)$ to be fixed later, by Lusin's theorem [37, Appendix E, §E8] there exist continuous functions g,g' and a set $C\subset\pi^{-1}(w)$ of probability at most ε such that $f_0=g$ and $f'_0=g'$ outside C. Then there is a point $v_0\in D\setminus C$ where $g(v_0)\neq g'(v_0)$. By continuity there exist $r_1,r_2,r_3>0$ such that for all $v\in B_{r_1}(v_0)\subset V$ we have $g(v)\in B_{r_2}(g(v_0))\subset Z, g'(v)\in B_{r_3}(g'(v_0))$, and $B_{r_3}(g'(v_0))\cap B_{2r_2}(g(v_0))=\emptyset$. Urysohn's lemma gives us continuous functions $F_1:Z\to[0,1]$, $F_2:V\to[0,1]$, such that $1_{B_{r_2}(g(v_0))}\leq F_1\leq 1_{B_{2r_2}(g(v_0))}$ and $1_{B_{r_1/2}(v_0)}\leq F_2\leq 1_{B_{r_1}(v_0)}$. We then have $\varphi_{F_1,F_2}(g)\geq \mu_w(B_{r_1/2}(v_0))>0$, whereas $\varphi_{F_1,F_2}(g')=0$. We also have $|\varphi_{F_1,F_2}(g)-\varphi_{F_1,F_2}(f_0)|\leq 2\mu_w(C)$ and similarly for g',f'_0 . Hence, choosing ε small enough, we deduce that $\varphi_{F_1,F_2}(f_0)>\varphi_{F_1,F_2}(f'_0)$.

To see that $\mathcal{L}(V,Z)$ is regular, first note that since by definition the finite intersections of sets $\varphi_{F_1,F_2}^{-1}(U)$ form a base for the topology, we have that for every closed set S and $f \notin S$, there exist $\varphi_{F_1,i,F_2,i}^{-1}(U_i)$, $i \in [N]$ such that $f \in \bigcap_i \varphi_{F_1,i,F_2,i}^{-1}(U_i)$ and $S \subset \bigcup_i \varphi_{F_1,i,F_2,i}^{-1}(\mathbb{C} \setminus U_i)$. Since \mathbb{C} is regular, for each i there exist disjoint open sets U_i', U_i'' such that $\varphi_{F_1,i,F_2,i}(f) \in U_i' \subset U_i$ and $\mathbb{C} \setminus U_i \subset U_i''$. Then the open sets $\bigcap_i \varphi_{F_1,i,F_2,i}^{-1}(U_i')$ and $\bigcup_i \varphi_{F_1,i,F_2,i}^{-1}(U_i'')$ are disjoint, the former contains f, and the latter includes S.

Finally, we show that $\mathcal{L}(V,Z)$ is second-countable. It suffices to find sequences of continuous functions $F_{1,i}:Z\to\mathbb{C},\,F_{2,j}:V\to\mathbb{C}$ and open sets $U_k\subset\mathbb{C},\,i,j,k\in\mathbb{N}$, such that for every φ_{F_1,F_2} , every open set $U\subset\mathbb{C}$, and every $f\in\mathcal{L}(V,Z)$ with $\varphi_{F_1,F_2}(f)\in U$, there exist i,j,k such that $f\in\varphi_{F_1,i,F_2,j}^{-1}(U_k)\subset\varphi_{F_1,F_2}^{-1}(U)$. Indeed, if this holds then the finite intersections of sets $\varphi_{F_1,i,F_2,j}^{-1}(U_k)$ form a countable base for the topology on $\mathcal{L}(V,Z)$. Let $\{U_k:k\in\mathbb{N}\}$ be a base of open discs in \mathbb{C} . The spaces of continuous functions $C(V,\mathbb{C})$ and $C(Z,\mathbb{C})$, equipped with the uniform metric, are separable [28, Theorem (4.19)]. Let $(F_{1,i})_{i\in\mathbb{N}}$ and $(F_{2,j})_{j\in\mathbb{N}}$ be separating sequences for $C(V,\mathbb{C})$ and $C(Z,\mathbb{C})$ respectively. Given any φ_{F_1,F_2} , open set $U\subset\mathbb{C}$, and $f\in\varphi_{F_1,F_2}^{-1}(U)$, for some $\varepsilon>0$ there is a disc $U_k\subset U$ of radius at most ε with center at most $\varepsilon/2$ away from $\varphi_{F_1,F_2}(f)$ and such that the distance from U_k to the complement of U is at least ε . There exist also i,j such that $\|F_1-F_{1,i}\|_{\infty}<\varepsilon/(4(1+\|F_2\|_{\infty}))$ and $\|F_2-F_{2,j}\|_{\infty}<\varepsilon/(4(1+\|F_{1,i}\|_{\infty}))$. Then for every $g\in\mathcal{L}(V,Z)$, letting $w=\tilde{\pi}(g)$, we have

$$|\varphi_{F_1,F_2}(g)-\varphi_{F_{1,i},F_{2,j}}(g)| \leq \int_{\pi^{-1}(w)} |F_1(g(v))F_2(v)-F_{1,i}(g(v))F_{2,j}(v)| \ \mathrm{d}\mu_w(v).$$

By the triangle inequality, this is at most $\|F_2\|_{\infty} \|F_1 - F_{1,i}\|_{\infty} + \|F_{1,i}\|_{\infty} \|F_2 - F_{2,j}\|_{\infty} < \varepsilon/2$. This implies that $f \in \varphi_{F_{1,i},F_{2,j}}^{-1}(U_k) \subset \varphi_{F_1,F_2}^{-1}(U)$.

Remark 2.3.4. There are spaces $\mathcal{L}(V,Z)$ that do not admit any complete metric. For example, let V be the one-point compactification of the half-strip $[0,\infty)\times[0,1]\subset\mathbb{R}^2$, let W be the one-point compactification of $[0,\infty)$, let $\pi:V\to W$ be the projection to the first coordinate, let $Z=\{0,1\}$, and for each $w\in W\setminus\{\infty\}$ let the interval $\pi^{-1}(w)$ be equipped with the Lebesgue probability measure. For each $n\in\mathbb{N}$ let f_n be the function on $\pi^{-1}(\{n\})$ that takes value 0 on $\{n\}\times[0,1/2)$ and 1 on $\{n\}\times[1/2,1]$. The sequence (f_n) is Cauchy for any compatible metric on $\mathcal{L}(V,Z)$ but it does not converge to a function $\{\infty\}\to Z$. To see this, we can view $\mathcal{L}(V,Z)$ as a subspace of a larger space $D(V,Z)=\bigcup_{w\in W}L(\pi^{-1}(w),P(Z))$, where each fibre $L(\pi^{-1}(w),P(Z))$ consists of probability-valued functions \tilde{f} on $\pi^{-1}(w)$, i.e. functions mapping each $v\in\pi^{-1}(w)$ to a probability measure p on Z. (We can identify $L(\pi^{-1}(w),P(Z))$ with $L(\pi^{-1}(w),[0,1])$ by identifying \tilde{f} with the map $v\mapsto p_v(\{0\})$.) The topology on D(V,Z) is the initial topology on $\mathcal{L}(V,Z)$. A straightforward calculation shows that f_n converges in D(V,Z) to the function \tilde{f} mapping ∞ to the uniform distribution on $\{0,1\}$. (I am grateful to Szegedy for providing this example.)

The topology on $\mathcal{L}(V,Z)$ can be viewed as a generalization of the L^1 topology, as illustrated by the following result.

Lemma 2.3.5. Let V, W be compact spaces, let Z be the closed unit disc in \mathbb{C} , and let $w \in W$. Then a sequence (f_n) in $L(\pi^{-1}(w), Z)$ converges to f in L(V, Z) if and only if it converges to f in $L^1(\pi^{-1}(w), Z)$.

This is a special case of a more general result that we shall use in the sequel, namely Lemma 2.3.7 below. This result involves the following generalization of the L^1 distance for functions taking values in a compact metric space.

Definition 2.3.6. Given a compact space Y with a Borel probability measure, and a compact space Z with a compatible metric d, we define the metric d_1 on L(Y,Z) by

$$d_1(f_1, f_2) = \mathbb{E}_{v \in Y} d(f_1(v), f_2(v)). \tag{2.8}$$

Lemma 2.3.7. Let V,W,Z be compact spaces, let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$, let d be a compatible complete metric on Z, and let $w \in W$. Then a sequence (f_n) in $L(\pi^{-1}(w),Z)$ converges to f in $\mathcal{L}(V,Z)$ if and only if $d_1(f,f_n) \to 0$ as $n \to \infty$.

Proof. Let φ_{F_1,F_2} be any of the functions generating the topology on $\mathcal{L}(V,Z)$. We then have

$$|\varphi_{F_1,F_2}(f_n) - \varphi_{F_1,F_2}(f)| \leq ||F_2||_{\infty} \int_{\pi^{-1}(w)} |F_1(f_n(v)) - F_1(f(v))| \,\mathrm{d}\mu_w(v).$$

We claim that if $d_1(f, f_n) \to 0$ then the last integral tends to 0 as $n \to \infty$. To see this fix any $\varepsilon > 0$ and note that by uniform continuity of F_1 there is $\delta_0 > 0$ such that if $d(f_n(v), f(v)) < \delta_0$ then $|F_1(f_n(v)) - F_1(f(v))| < \varepsilon/2$. Fix a positive $\delta < \min(\delta_0, \varepsilon/4 ||F_1||_{\infty})$. For n sufficiently large we have $d_1(f_n, f) < \delta^2$,

whence, by Markov's inequality, the set $D_{\delta} = \{v \in \pi^{-1}(w) : d(f_n(v), f(v)) \geq \delta\}$ has $\mu_w(D_{\delta}) \leq \delta$. It follows that

$$\int_{\pi^{-1}(w)} |F_1(f_n(v)) - F_1(f(v))| \, \mathrm{d}\mu_w \, \leq \, \varepsilon/2 + 2\|F_1\|_\infty \, \mu_w(D_\delta) \, \leq \, \varepsilon.$$

The claim follows, and so $d_1(f_n, f) \to 0$ implies that $f_n \to f$ in $\mathcal{L}(V, Z)$.

For the converse, suppose that $f_n \to f$ in $\mathcal{L}(V,Z)$. By [28, Theorem (4.17)] and [2, Theorem 1], there exists a compact subset B of the real Hilbert space ℓ^2 such that there is a homeomorphism $F:Z\to B$. We have that $F\circ f_n$, $F\circ f$ are Bochner measurable functions in $L^2(V,\ell^2)$. We claim that if $F\circ f_n\to F\circ f$ in $L^2(V,\ell^2)$ (that is $\int_{\pi^{-1}(w)}\|F\circ f_n(v)-F\circ f(v)\|_{\ell^2}\,\mathrm{d}\mu_w\to 0$) then $d_1(f_n,f)\to 0$ as required. Indeed, for any $\varepsilon>0$, the set $E=\{(z,z'):d(z,z')\geq\varepsilon\}$ is closed in $Z\times Z$ so compact, and therefore the continuous function $\frac{d(z,z')}{\|F(z)-F(z')\|_{\ell^2}}$ on E reaches a maximum E, whence for E sufficiently large we have

$$d_1(f_n, f) = \int_{\pi^{-1}(w)} d(f_n(v), f(v)) d\mu_w \leq \varepsilon + C \int_{\pi^{-1}(w)} 1_E(f_n(v), f(v)) \|F \circ f_n(v) - F \circ f(v)\|_{\ell^2} d\mu_w$$

$$\leq 2\varepsilon.$$

This proves our claim. We now show that $F \circ f_n \to F \circ f$ in $L^2(V, \ell^2)$.

Firstly, for every $F_0 \in L^2(V, \ell^2)$, we show that as $n \to \infty$ we have $\int_{\pi^{-1}(w)} \langle F \circ f_n(v), F_0(v) \rangle_{\ell^2} \, \mathrm{d}\mu_w \to \int_{\pi^{-1}(w)} \langle F \circ f(v), F_0(v) \rangle_{\ell^2} \, \mathrm{d}\mu_w$. By definition of the Bochner integral, it suffices to prove this for simple functions F_0 . By linearity, it then suffices to prove it for F_0 of the form $\alpha \, 1_A$ where $A \subset \pi^{-1}(w)$ is measurable and $\alpha \in \ell^2(\mathbb{R})$. But in this case the function $F_1 : Z \to \mathbb{R}, z \mapsto \langle F(z), \alpha \rangle$ is continuous, and the function $1_A : \pi^{-1}(w) \to \mathbb{R}$ can be approximated in $L^1(\mu_w)$ by continuous functions, hence in this case the convergence follows from the definition of the topology on $\mathcal{L}(V,Z)$.

Secondly, we show that $\int_{\pi^{-1}(w)} \|F \circ f_n(v)\|_{\ell^2}^2 d\mu_w$ converges to $\int_{\pi^{-1}(w)} \|F \circ f(v)\|_{\ell^2}^2 d\mu_w$. This follows from the fact that $F_1: Z \to \mathbb{R}$, $z \to \|F(z)\|_{\ell^2}^2$ is continuous, so that letting F_2 be the function with value 1 on all of V we have $\int_{\pi^{-1}(w)} \|F \circ f_n(v)\|_{\ell^2}^2 d\mu_w = \varphi_{F_1,F_2}(f_n)$, which converges to $\varphi_{F_1,F_2}(f)$ as required.

From the last two paragraphs we deduce the desired convergence by a standard argument, namely

$$\int_{\pi^{-1}(w)} \|F \circ f_n(v) - F \circ f(v)\|_{\ell^2}^2 d\mu_w = \int_{\pi^{-1}(w)} \langle F \circ f_n - F \circ f, F \circ f_n - F \circ f \rangle d\mu_w \\
= \int \|F \circ f_n\|_{\ell^2}^2 - 2\langle F \circ f_n, F \circ f \rangle + \|F \circ f\|_{\ell^2}^2 d\mu_w \to \int 2\|F \circ f\|_{\ell^2}^2 - 2\langle F \circ f, F \circ f \rangle d\mu_w = 0. \quad \Box$$

It is a basic fact that if (f_n) is a sequence of real-valued functions on a standard probability space such that $f_n \to f$ in L^1 and each f_n is constant almost surely, then the limit f is also constant almost surely. We shall use the following analogous fact for $\mathcal{L}(V,Z)$.

Lemma 2.3.8. Let V, W, Z be compact spaces, let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$, and let (f_n) be a convergent sequence in $\mathcal{L}(V, Z)$ such that each f_n is constant almost surely on its fibre $\pi^{-1}(w_n)$. Then the limit is also constant almost surely.

Proof. Let f be the limit of the sequence (f_n) . By Urysohn's lemma on Z, it suffices to show that for every continuous function $F: Z \to \mathbb{R}$ the function $F \circ f$ is constant. Now since F and F^2 are both continuous, by definition of the topology on $\mathcal{L}(V,Z)$ we have $\int_{\pi^{-1}(w_n)} F \circ f_n d\mu_{w_n} \to \int_{\pi^{-1}(w)} F \circ f d\mu_w$ and

 $\int_{\pi^{-1}(w_n)} F^2 \circ f_n \, d\mu_{w_n} \to \int_{\pi^{-1}(w)} F^2 \circ f \, d\mu_w$. On the other hand, since each f_n is constant almost everywhere, we have $\left(\int_{\pi^{-1}(w_n)} F \circ f_n \, d\mu_{w_n}\right)^2 - \int_{\pi^{-1}(w_n)} F^2 \circ f_n \, d\mu_{w_n} = 0$. Taking limits as $n \to \infty$, we deduce that $\left(\int_{\pi^{-1}(w)} F \circ f \, d\mu_w\right)^2 - \int_{\pi^{-1}(w)} F^2 \circ f \, d\mu_w = 0$. By the equality case of the Cauchy-Schwarz inequality, this implies that $F \circ f$ is constant almost everywhere.

The following result will be used several times in the next section.

Lemma 2.3.9. Let V, W, Z be compact spaces, and let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$. Suppose that (f_n) and (g_n) are sequences in $\mathcal{L}(V, Z)$ converging to f and g respectively, where f_n, g_n are defined on $\pi^{-1}(w_n)$ and f, g are defined on $\pi^{-1}(w)$. Then the functions $v \mapsto (f_n(v), g_n(v))$ converge to $v \mapsto (f(v), g(v))$ in $\mathcal{L}(V, Z \times Z)$.

Proof. It suffices to show that for any continuous functions $F_1: Z \times Z \to \mathbb{C}$ and $F_2: V \to \mathbb{C}$ we have $\lim_{n\to\infty}\int_{\pi^{-1}(w_n)}F_1(f_n(v),g_n(v))F_2(v)\,\mathrm{d}\mu_{w_n}(v)=\int_{\pi^{-1}(w)}F_1(f(v),g(v))F_2(v)\,\mathrm{d}\mu_{w}(v)$. By the Stone-Weierstrass theorem the function F_1 can be approximated arbitrarily closely in the uniform metric by finite linear combinations of functions of the form $(z_1,z_2)\mapsto F_{1,1}(z_1)F_{1,2}(z_2)$ where $F_{1,1},F_{1,2}:Z\to\mathbb{C}$ are continuous, so we may assume that F_1 is of this form. Fix any $\varepsilon>0$. By Lusin's theorem there exists a continuous function $q':\pi^{-1}(w)\to\mathbb{C}$ such that $\|F_{1,2}\circ g-q'\|_{L^2}<\varepsilon$ and $\|q'\|_\infty\leq \|F_{1,2}\|_\infty$, and then by the Tietze extension theorem [32, Theorem 35.1] there exists a continuous function $q:V\to\mathbb{C}$ such that the restriction $q|_{\pi^{-1}(w)}$ equals q' and $\|q\|_\infty\leq \|F_{1,2}\|_\infty$. For each n let $q_n=q|_{\pi^{-1}(w_n)}$. Now

$$||F_{1,2} \circ g_n - q_n||_{L^2}^2 = \int_{\pi^{-1}(w_n)} |F_{1,2} \circ g_n - q_n|^2 d\mu_{w_n}$$

$$= \int_{\pi^{-1}(w_n)} |F_{1,2} \circ g_n|^2 - 2\operatorname{Re}(F_{1,2} \circ g_n \overline{q_n}) + |q_n|^2 d\mu_{w_n}. \tag{2.9}$$

Using that $g_n \to g$ in $\mathcal{L}(V, Z)$, that $w_n \to w$, and the uniform continuity of q, we deduce that the integral above converges to $||F_{1,2} \circ g - q'||_{L^2}^2$. Setting $h_n = (F_{1,1} \circ f_n)(F_{1,2} \circ g_n - q_n)$, we then have

$$F_1 \circ (f_n, g_n) = (F_{1,1} \circ f_n) \left(q_n + (F_{1,2} \circ g_n - q_n) \right) = (F_{1,1} \circ f_n) q_n + h_n$$

and similarly $F_1 \circ (f,g) = (F_{1,1} \circ f) \, q' + h$, where for n sufficiently large we have $||h_n||_{L^2}$ and $||h||_{L^2}$ both at most $2\varepsilon ||F_{1,1}||_{\infty}$. The convergence $f_n \to f$ also implies that

$$\int_{\pi^{-1}(w_n)} (F_{1,1} \circ f_n) \, q_n F_2 \, \mathrm{d}\mu_{w_n} \, \to \, \int_{\pi^{-1}(w)} (F_{1,1} \circ f) \, q' \, F_2 \, \mathrm{d}\mu_w.$$

Combining the above estimates we deduce that for all *n* sufficiently large we have

$$\left| \int_{\pi^{-1}(w_n)} F_1 \circ (f_n, g_n) F_2 d\mu_{w_n} - \int_{\pi^{-1}(w)} F_1 \circ (f, g) F_2 d\mu_{w} \right| \leq 5\varepsilon \|F_{1,1}\|_{\infty} \|F_2\|_{\infty}.$$

Since $\varepsilon > 0$ was arbitrary, the result follows.

For the main applications of the space $\mathcal{L}(V,Z)$ in the sequel, we will suppose that Z is a compact abelian group. We now show that in this case the group has a continuous action on the space.

Lemma 2.3.10. Let V, W be compact spaces, let Z be a compact abelian group, and let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$. Then we have the following free continuous action of Z on $\mathcal{L}(V, Z)$:

$$\alpha: \mathbb{Z} \times \mathcal{L}(V,\mathbb{Z}) \to \mathcal{L}(V,\mathbb{Z}), \quad \alpha(z,f): v \mapsto f(v) + z.$$

Proof. The topology on $\mathcal{L}(V, \mathbb{Z})$ is generated by the functions φ_{F_1, F_2} , so to check continuity of α it suffices to show that for every φ_{F_1, F_2} the map $\varphi_{F_1, F_2} \circ \alpha$ is continuous. Thus, it suffices to show that for any continuous functions $F_1 : \mathbb{Z} \to \mathbb{C}$ and $F_2 : V \to \mathbb{C}$, the following function is continuous:

$$q: (z,f) \mapsto \int_{\pi^{-1}(\tilde{\pi}(f))} F_1(f(v)+z) F_2(v) d\mu_{\tilde{\pi}(f)}(v).$$
 (2.10)

Applying the Stone-Weierstrass theorem as in previous proofs, the continuous function $Z \times Z \to \mathbb{C}$ defined by $(z_1, z_2) \mapsto F_1(z_1 + z_2)$ can be approximated uniformly by linear combinations of functions of the form $(z_1, z_2) \mapsto h_1(z_1) h_2(z_2)$ where h_1 and h_2 are continuous. Substituting any of these into (2.10), we obtain a function of the form

$$q': (z,f) \mapsto h_2(z) \int_{\pi^{-1}(\tilde{\pi}(f))} h_1(f(v)) F_2(v) d\mu_{\tilde{\pi}(f)}(v) = h_2(z) \varphi_{h_1,F_2}(f).$$

This is continuous by definition of $\mathcal{L}(V, \mathbb{Z})$. Hence q is a uniform limit of continuous functions and is therefore continuous.

In our uses of $\mathcal{L}(V, \mathbb{Z})$, we shall want to view this space as a continuous Z-bundle, for some appropriate base S and projection \mathcal{E} . We would like S, \mathcal{E} to have the following features. Firstly, the projection should send two functions f_1, f_2 in $\mathcal{L}(V, \mathbb{Z})$ to the same point of S if and only if f_1 and f_2 are in the same Z-orbit (i.e. $f_2 = f_1 + z$ for some $z \in \mathbb{Z}$). Secondly, we would like S to be also of the form $\mathcal{L}(Y, \mathbb{Z})$, for some other similar space Y having a projection to W, as this would help to relate the topologies on S and $\mathcal{L}(V, \mathbb{Z})$.

A natural way to obtain these features is to define \mathcal{E} as a difference map, sending f to the function $\mathcal{E}(f): V \times V \to \mathbb{Z}$, $(v_0, v_1) \mapsto f(v_0) - f(v_1)$. We are thus led to the following construction.

Definition 2.3.11. Let V, W be compact spaces and let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$. We define the compact space

$$V \times_W V = \{(v_0, v_1) : v_0, v_1 \in V, \ \pi(v_0) = \pi(v_1)\},\$$

and define the projection $\pi': V \times_W V \to W$ by $\pi'(v_0, v_1) = \pi(v_0)$.

Lemma 2.3.12. The family of measures $\{\mu_w \times \mu_w : w \in W\}$ on $V \times_W V$ is a CSM on π' .

Proof. For every $w \in W$ the measure μ_w is concentrated on $\pi^{-1}(w)$, so in $V \times V$ the measure $\mu_w \times \mu_w$ is concentrated on $\pi^{-1}(w) \times \pi^{-1}(w) = \pi'^{-1}(w)$. For any continuous function $f: V \times_W V \to \mathbb{C}$, the continuity of $w \mapsto \int_{\pi'^{-1}(w)} f \, \mathrm{d}(\mu_w \times \mu_w)$ is proved by an argument similar to previous ones (e.g. in the proof of Lemma 2.3.10), showing that this function is a uniform limit of continuous functions F_n , where F_n is obtained by approximating f to within 1/n in the supremum norm by a finite linear combination of functions of the form $(v_0, v_1) \mapsto h_0(v_0) h_1(v_1)$ with $h_i : V \to \mathbb{C}$ continuous, and then using the continuity of $w \mapsto \prod_{i=0,1} \int_{\pi^{-1}(w)} h_i \, \mathrm{d}\mu_w$ guaranteed by the CSM on π .

We are now able to complete the construction of the continuous abelian bundle announced at the beginning of this section, by showing that the map \mathcal{E} has the required properties to be a suitable projection.

Proposition 2.3.13. Let V, W be compact spaces, let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$, and let Z be a compact abelian group. Let $\mathcal{E} : \mathcal{L}(V, Z) \to \mathcal{L}(V \times_W V, Z)$ be defined by $\mathcal{E}(g) : (v_0, v_1) \mapsto g(v_0) - g(v_1)$. Then $\mathcal{L}(V, Z)$ is a continuous Z-bundle over $\mathcal{E}(\mathcal{L}(V, Z))$ with projection \mathcal{E} .

Proof. Recall that the topology on $\mathcal{L}(V \times_W V, Z)$ is generated by functions φ_{F_1,F_2} sending $f: V \times_W V \to Z$ with $\tilde{\pi}(f) = w$ to $\int_{\pi'^{-1}(w)} F_1(f(v_0,v_1)) F_2(v_0,v_1) \, \mathrm{d}\mu_w(v_0,v_1)$. To see that \mathcal{E} is continuous, fix any such function and suppose that $g_n \to g$ in $\mathcal{L}(V,Z)$. Then, approximating the continuous functions $(z_0,z_1) \mapsto F_1(z_0-z_1)$ on Z^2 and F_2 on $\pi'^{-1}(w) = \pi^{-1}(w) \times \pi^{-1}(w)$ by combinations of continuous functions $h_0 h_1$ as in previous proofs, we approximate $\varphi_{F_1,F_2}(\mathcal{E}(g_n))$ by finite combinations of products $\varphi_0(g_n)\varphi_1(g_n)$, where φ_0, φ_1 are functions among those generating the $\mathcal{L}(V,Z)$ topology. By assumption, each of these products satisfies $\varphi_0(g_n)\varphi_1(g_n) \to \varphi_0(g)\varphi_1(g)$, and it follows that $\varphi_{F_1,F_2}(\mathcal{E}(g_n)) \to \varphi_{F_1,F_2}(\mathcal{E}(g))$. This proves the continuity of \mathcal{E} .

To ensure that we have a continuous Z-bundle, it only remains to check the 'if' part of condition (iv) in Definition 2.1.6. Equivalently, it now suffices to prove that if $\mathcal{E}^{-1}(U)$ is closed then U is closed. For that, we claim it is enough to show that if for a sequence of functions $f_n \in \mathcal{L}(V, \mathbb{Z})$ the images $\mathcal{E}(f_n)$ converge, then there is a convergent subsequence (f_m) of (f_n) . Indeed, suppose that this holds and suppose for a contradiction that U is not closed. Then (using the surjectivity of \mathcal{E}) there is a sequence of functions $\mathcal{E}(f_n) \in U$ converging to a function $\mathcal{E}(f) \not\in U$. But then letting f' be the limit of the guaranteed subsequence (f_m) , the closure of $\mathcal{E}^{-1}(U)$ implies that $f' \in \mathcal{E}^{-1}(U)$, and so by continuity of \mathcal{E} we have $U \ni \mathcal{E}(f') = \lim_{m \to \infty} \mathcal{E}(f_m) = \mathcal{E}(f) \not\in U$, a contradiction. This proves our claim.

So let us suppose that $\mathcal{E}(f_n)$ converges to $\mathcal{E}(f)$ for some $f \in \mathcal{L}(V, \mathbb{Z})$. Let $w_n = \tilde{\pi}(f_n)$ and $w = \tilde{\pi}(f)$, and note that $w_n \to w$ in W by continuity of $\tilde{\pi}$.

Let \mathcal{C} denote the circle group, viewed as the unit circle in \mathbb{C} . We shall prove the following fact:

There exists a subsequence
$$(f_m)$$
 and some $t \in \mathbb{Z}$ such that, for every character $\chi: \mathbb{Z} \to \mathcal{C}$ in $\widehat{\mathbb{Z}}$, we have $\chi \circ f_m \to \chi \circ (f+t)$ in $\mathcal{L}(V,\mathcal{C})$ as $m \to \infty$.

We prove this below, but before that let us show how it implies that f_m converges to f + t in $\mathcal{L}(V, \mathbb{Z})$.

Let φ_{F_1,F_2} be any of the functions generating the $\mathcal{L}(V,\mathbf{Z})$ topology, and note that by the Stone-Weierstrass theorem for every $\varepsilon>0$ there exists a trigonometric polynomial $F_1'=\sum_{j\in[K]}\lambda_j\chi_j$ on \mathbf{Z} such that $\|F_1-F_1'\|_{\infty}\leq \varepsilon$. Therefore it suffices to show that $\varphi_{F_1',F_2}(f_m)\to \varphi_{F_1',F_2}(f+t)$ for any such polynomial F_1' . We have

$$\phi_{F_1',F_2}(f_m) = \sum_{j \in [K]} \lambda_j \int_{\pi^{-1}(w_m)} (\chi_j \circ f_m) \, F_2 \, \mathrm{d} \mu_{w_m} = \sum_{j \in [K]} \lambda_j \, \phi_{1,F_2}(\chi_j \circ f_m).$$

Since φ_{1,F_2} is one of the functions generating the topology on $\mathcal{L}(V,\mathbb{C})$, we have $\varphi_{1,F_2}(\chi_j \circ f_m) \to \varphi_{1,F_2}(\chi_j \circ (f+t))$ by assumption, and the desired convergence follows.

We now prove (2.11). Given $\chi \in \widehat{\mathbb{Z}}$, the function $\chi \circ f$ is in $L^2(\mu_w)$. Continuous functions are dense in this space, and any such function can be extended to a continuous function on V (by the Tietze extension theorem). Hence, for every $\delta > 0$ there exists a continuous function $\chi_{\delta} : V \to \mathcal{C}$ whose restriction to

 $\pi^{-1}(w)$ satisfies $\|\chi_{\delta} - \chi \circ f\|_{L^{2}(\mu_{w})} \leq \delta$. We claim that for every $\chi \in \widehat{\mathbb{Z}}$ and $\varepsilon > 0$, there exist $\delta > 0$, $N \in \mathbb{N}$, and $\lambda \in \mathbb{C}$, such that for all n > N we have $\|(\chi \circ f_{n})\overline{\chi_{\delta}} - \lambda\|_{L^{2}(\mu_{w_{n}})} \leq \varepsilon$.

To prove this claim, let $\mathcal{E}': \mathcal{L}(V,\mathcal{C}) \to \mathcal{L}(V \times_W V,\mathcal{C})$ be defined by $\mathcal{E}'(g): (v_0,v_1) \mapsto g(v_0)\overline{g(v_1)}$. We have that \mathcal{E}' is continuous, by an argument similar to the one used above for \mathcal{E} . It is also clear that for any $g_1,g_2 \in \mathcal{L}(V,\mathcal{C})$ lying in the same fibre of $\tilde{\pi}$ we have $\mathcal{E}'(g_1g_2) = \mathcal{E}'(g_1)\mathcal{E}'(g_2)$, and that $\mathcal{E}'(\chi \circ g) = \chi \circ \mathcal{E}(g)$ for every $\chi \in \widehat{Z}$, $g \in \mathcal{L}(V,Z)$. Now, from the definition of χ_{δ} , the continuity of \mathcal{E}' , and Lemma 2.3.7, it follows that for δ sufficiently small we have $\|\mathcal{E}'((\chi \circ f)\overline{\chi_{\delta}}) - 1\|_{L^2(\mu_w \times \mu_w)} \leq \varepsilon/2$. On the other hand, we have $\mathcal{E}'((\chi \circ f_n)\overline{\chi_{\delta}}) = \mathcal{E}'(\chi \circ f_n)\mathcal{E}'(\overline{\chi_{\delta}}) = \chi \circ \mathcal{E}(f_n)\mathcal{E}'(\overline{\chi_{\delta}})$ and similarly for f. This together with the assumption $\mathcal{E}(f_n) \to \mathcal{E}(f)$ implies, by an argument similar to (2.9), that for g sufficiently large we have $\|\mathcal{E}'((\chi \circ f_n)\overline{\chi_{\delta}}) - 1\|_{L^2(\mu_{w_n} \times \mu_{w_n})} \leq \varepsilon$. Taking a minimum over v_1 in this integral, we deduce that there exists $v_1 \in \pi^{-1}(w_n)$ such that, for $\lambda = (\chi \circ f_n)(v_1)\overline{\chi_{\delta}}(v_1)$, we have indeed $\|(\chi \circ f_n)\overline{\chi_{\delta}} - \lambda\|_{L^2(\mu_{w_n})} \leq \varepsilon$.

Since there are at most countably many elements in \widehat{Z} , it follows by a standard diagonalization argument (using the above claim with smaller and smaller ε for each χ) that for some sequence $(m_n)_{n\in\mathbb{N}}$ of positive integers, for every $\chi\in\widehat{Z}$ there is a constant $\lambda_{\chi}\in\mathbb{C}$ such that $\chi\circ f_{m_n}\to\lambda_{\chi}(\chi\circ f)$ in $\mathcal{L}(V,\mathbb{C})$ as $n\to\infty$. A straightforward calculation shows that the function $\chi\mapsto\lambda_{\chi}$ is a homomorphism, hence a character on \widehat{Z} , which by duality must be of the form $\chi\mapsto\chi(t)$ for some $t\in\mathbb{Z}$. This confirms (2.11) and completes the proof.

We close this section with two technical lemmas that we shall use several times.

Lemma 2.3.14. Let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$. Let K be a compact space with a Borel measure V. Suppose that $f_1 : V \to K$ is continuous and that, for every $w \in W$, the restriction of f_1 to $\pi^{-1}(w)$ satisfies $\mu_w \circ f_1^{-1} = V$. Let Z be a compact space and let $f_2 : K \to Z$ be a Borel function. Let $h : W \to \mathcal{L}(V, Z)$ be the map $w \mapsto f_2 \circ f_1|_{\pi^{-1}(w)}$. Then h is continuous.

Proof. By definition of the topology on $\mathcal{L}(V,Z)$ it suffices to show that for every φ_{F_1,F_2} the function $\varphi_{F_1,F_2} \circ h : W \to \mathbb{C}$ is continuous. So let $F_1 : Z \to \mathbb{C}$ and $F_2 : V \to \mathbb{C}$ be continuous and let q denote $\varphi_{F_1,F_2} \circ h$, that is $q : w \mapsto \int_{\pi^{-1}(w)} F_1(f_2 \circ f_1(v)) F_2(v) d\mu_w(v)$. To prove that q is continuous we show that it is a uniform limit of continuous functions.

Fix an arbitrary $\varepsilon > 0$. By Lusin's theorem there exists a continuous function $F_3 : K \to \mathbb{C}$ such that $\|F_1 \circ f_2 - F_3\|_{L^1(v)} \le \varepsilon$. Letting $q' : w \mapsto \int_{\pi^{-1}(w)} F_3(f_1(v)) F_2(v) \, \mathrm{d}\mu_w(v)$, the triangle inequality and the assumption that f_1 is measure-preserving imply that $\sup_{w \in W} |q(w) - q'(w)| \le \varepsilon \|F_2\|_{\infty}$. The function q' is continuous by condition (ii) from Definition 2.2.1. The result follows.

Lemma 2.3.15. Let $\{\mu_w : w \in W\}$ be a CSM on $\pi : V \to W$, let Z be a compact abelian group, let P be a finite set, and for each $r \in P$ let λ_r be an integer. Let $\Sigma_P : \mathcal{L}(V, Z^P) \to \mathcal{L}(V, Z)$, $f \mapsto \sum_{r \in P} \lambda_r \pi_r \circ f$, where π_r is the projection to the r-component on Z^P . Then Σ_P is continuous.

Proof. For any of the functions φ_{F_1,F_2} generating the topology on $\mathcal{L}(V,\mathbf{Z})$, letting $w=\tilde{\pi}(f)$, we have $\varphi_{F_1,F_2}(\Sigma_P(f))=\int_{\pi^{-1}(w)}F_1\circ\Sigma_P(f)F_2\,\mathrm{d}\mu_w$. But $F_1\circ\Sigma_P(f)=F_1'\circ f$, where $F_1':\mathbf{Z}^P\to\mathbb{C}$, $(z_r)_{r\in P}\mapsto F_1\left(\sum_r\lambda_rz_r\right)$ is continuous. Hence, by definition of $\mathcal{L}(V,\mathbf{Z}^P)$, if $f_n\to f$ in this space then $\varphi_{F_1,F_2}(\Sigma_P(f_n))\to \varphi_{F_1,F_2}(\Sigma_P(f))$, and since this holds for every φ_{F_1,F_2} , we have $\Sigma_P(f_n)\to\Sigma_P(f)$ in $\mathcal{L}(V,\mathbf{Z})$.

2.4 Borel measurable cocycles generate compact extensions

Let X be a compact nilspace and let Z be a compact abelian group. The purpose of this section is to show that any measurable cocycle $\rho : C^k(X) \to Z$ yields a compact nilspace that is a Z-bundle over X.

Recall that for each $x \in X$ we denote by $C_x^k(X)$ the set of k-cubes c satisfying $\pi(c) := c(0^k) = x$. By Lemma 2.2.17 each space $C_x^k(X)$ has a Haar measure, denoted μ_x , and the family $\{\mu_x : x \in X\}$ is a CSM on π of strictly positive measures.

Definition 2.4.1. We denote by $\mathcal{L}_k(X,Z)$ the space $\mathcal{L}(V,Z)$ with $V = C^n(X)$, W = X, and $\pi : c \mapsto c(0^k)$.

Thus $\mathcal{L}_k(X,Z) = \bigcup_{x \in X} L(C_x^k(X),Z)$, its topology is second-countable regular Hausdorff, and we have a projection $\tilde{\pi}: \mathcal{L}_k(X,Z) \to X$ defined by $\tilde{\pi}(f) = x$ for $f \in L(C_x^k(X),Z)$. For such a function f, we denote by f + Z the set of functions $\{f + z : z \in Z\}$.

The desired nilspace extending X is found inside $\mathcal{L}_k(X,Z)$ by taking, for each $x \in X$, the orbit of the restricted cocycle $\rho_x = \rho|_{C_x^k(X)}$.

Proposition 2.4.2. Let $\rho: C^k(X) \to Z$ be a Borel measurable cocycle and let $M = \bigcup_{x \in X} (\rho_x + Z)$. Then, as a subspace of $\mathcal{L}_k(X,Z)$ with the inherited Z-action, the space M is a compact Z-bundle over X.

Proof. Let $Y = C^k(X) \times_X C^k(X)$, that is (recalling the notation from Proposition 2.3.13),

$$Y = \left\{ (c_0, c_1) : c_0, c_1 \in C^k(X), c_0(0^k) = c_1(0^k) \right\} = \bigcup_{x \in X} C^k_x(X) \times C^k_x(X).$$

Recall that $\mathcal{E}: \mathcal{L}(\mathbf{C}^k(\mathbf{X}), \mathbf{Z}) \to \mathcal{L}(Y, \mathbf{Z})$ is defined by $\mathcal{E}(f): (\mathbf{c}_0, \mathbf{c}_1) \mapsto f(\mathbf{c}_0) - f(\mathbf{c}_1)$. We now compose the function $x \mapsto \rho_x$ with \mathcal{E} , obtaining the function

$$g': X \rightarrow L(C_x^k(X) \times C_x^k(X), Z) \subset \mathcal{L}(Y, Z)$$

 $x \mapsto \mathcal{E}(\rho_x).$ (2.12)

Note that g' is injective, and that its inverse is the projection $\tilde{\pi}$ on $\mathcal{L}(Y, \mathbb{Z})$, a continuous map. We have

$$M = \mathcal{E}^{-1}(g'(X)). \tag{2.13}$$

Indeed, if $f \in \mathcal{L}(C^k(X), Z)$ satisfies $\mathcal{E}(f) = \mathcal{E}(\rho_x)$, then for every $(c_0, c_1) \in C_x^k(X)^2 \subset Y$ we have $f(c_0) - \rho_x(c_0) = f(c_1) - \rho_x(c_1)$, so $f - \rho_x$ is a constant Z-valued function on $C_x^k(X)$.

Now, by Proposition 2.3.13 the preimage under \mathcal{E} of every point is a compact set (homeomorphic to Z), and \mathcal{E} is also a closed map (using Remark 2.1.7). Hence \mathcal{E} is a proper map (see [6, §10.2, Theorem 1]). In particular, the preimage under \mathcal{E} of any compact set is compact ([6, §10.2, Proposition 6]). Therefore, if we show that g' is continuous, then g'(X) is compact and then M must be compact, by (2.13). (Note that M is then also a continuous Z-bundle over X, since by (2.13) it is such a bundle over g'(X), and g' is a homeomorphism $X \to g'(X)$.)

To prove the continuity of g', the key idea is to express this function in terms of compositions of ρ with certain tricube morphisms $T_k \to X$ (recall [8, Definition 3.1.13]), and thereby to use these morphisms

in a smoothening operation that turns the measurability of ρ into the continuity of g'.

Let
$$Q = \text{hom}(T_k, X)$$
 and $Q_x = \text{hom}_{1^k \to x}(T_k, X)$. For each $v \in \{0, 1\}^k \setminus \{0^k\}$, let

$$g_{\nu}: X \to \mathcal{L}(Q, Z), \ x \mapsto (g_{\nu}(x): t \mapsto \rho(t \circ \Psi_{\nu}) \in L(Q_{x}, Z)).$$

We claim that g_v is continuous. Indeed, by Lemma 2.2.23 if $v \neq 0^k$ then $t \mapsto t \circ \Psi_v$ is measure-preserving from Q_x to $C^k(X)$, so we can deduce the continuity of g_v by applying Lemma 2.3.14 (setting V = Q, $K = C^k(X)$, W = X, $f_2 = \rho$, $h = g_v$).

Now consider the following function:

$$g: X \to \mathcal{L}(Q, \mathbb{Z}), \quad x \mapsto \sum_{\nu \in \{0,1\}^k \setminus \{0^k\}} (-1)^{|\nu|} g_{\nu}.$$

The continuity of each function g_{ν} implies that g is continuous. Indeed, letting $P = \{0,1\}^k \setminus \{0^k\}$, by Lemma 2.3.9 the function $X \to \mathcal{L}(Q, Z^P)$, $x \mapsto \prod_{v \in P} g_v(x)$ is continuous, and then Lemma 2.3.15 gives the continuity of g. Now, by [8, Lemma 3.3.31], for each $x \in X$ we have that g(x) is the function $t \mapsto \rho_x(t \circ \omega_k) - \rho_x(t \circ \Psi_{0^k})$ (defined on Q_x). In particular, given any $t \in Q_x$, we have that $c_0 := t \circ \omega_k$ and $c_1 := t \circ \Psi_{0^k}$ are both cubes in $C_x^k(X)$ (c_0 is a cube by [8, Lemma 3.1.16]), and $g(x)(t) = \rho_x(c_0) - \rho_x(c_1) = g'(x)(c_0,c_1)$. Given this equality, we now claim that continuity of g' follows from continuity of g. Indeed, we know that g' is continuous if for any of the functions φ_{F_1,F_2} generating the topology on $\mathcal{L}(Y,Z)$ we have that $X \to \mathbb{C}$, $x \mapsto \varphi_{F_1,F_2}(g'(x))$ is continuous. Letting $Y_x = C_x^k(X) \times C_x^k(X)$, we have

$$\varphi_{F_1,F_2}(g'(x)) = \int_{Y_x} F_1(g'(x)(c_0,c_1)) F_2((c_0,c_1)) d\mu((c_0,c_1)),$$

where $F_1: \mathbb{Z} \to \mathbb{C}$ and $F_2: Y_x \to \mathbb{C}$ are both continuous. Now let $\theta: Q_x \to Y_x$, $t \mapsto (t \circ \omega_k, t \circ \Psi_{0^k})$. If we show that this map is surjective and measure-preserving, then since $g'(x) \circ \theta(t) = g(x)(t)$, we would conclude that

$$\varphi_{F_1,F_2}(g'(x)) = \int_{O_x} F_1(g(x)(t)) F_2(\theta(t)) d\mu(t)$$

and continuity of g' would then indeed follow from that of g (via approximating the measurable function $F_2 \circ \theta$ by continuous functions as in previous proofs).

The surjectivity of θ follows from combining Lemma 2.2.21 and 2.2.15. (To apply the latter result recall that we are implicitly embedding T_k as a subcubespace in $\{0,1\}^{2k}$, using [8, Lemma 3.1.17].)

To see that θ is measure-preserving, first note that by Lemma 2.2.14 there is a CSM on the map $Q_x \to C_x^k(X)$, $t \mapsto t \circ \omega_k$ consisting of the Haar probabilities on the fibres, giving a disintegration of the Haar probability on Q_x . On each of these fibres, of the form $\lim_{c \to \omega_k^{-1}} (T_k, X)$, $c \in C_x^k(X)$, it follows from Corollary 2.2.22 that the map $\lim_{c \to \omega_k^{-1}} (T_k, X) \to C_x^k(X)$, $t \mapsto t \circ \Psi_{0^k}$ preserves the Haar measures. Using this, a straightforward calculation shows that for product sets $A_1 \times A_2 \subset C_x^k(X) \times C_x^k(X)$ we have $\mu_{Q_x} \circ \theta^{-1}(A_1 \times A_2) = \mu_{C_x^k(X) \times C_x^k(X)} (A_1 \times A_2)$, and it follows that θ is measure-preserving as claimed. \square

Let us now recall the definition of cubes on M from [8, Definition 3.3.25]. Note that, by definition of M, for every function $f: \{0,1\}^k \to M$ there is some function $a = a_f: \{0,1\}^k \to \mathbb{Z}$ such that $\rho_x(v) = f(v) + a(v)$, where $x = \tilde{\pi}(f(v)) \in \mathbb{X}$.

Definition 2.4.3. We define $C^k(M)$ to be the set of functions $f: \{0,1\}^k \to M$ such that $\tilde{\pi} \circ f \in C^k(X)$ and

$$\rho(\tilde{\pi} \circ f) = \sigma_k(a). \tag{2.14}$$

For $n \neq k$, a function $f : \{0,1\}^n \to M$ is declared to be in $\mathbb{C}^n(M)$ if $\tilde{\pi} \circ f \in \mathbb{C}^n(X)$ and every k-dimensional face-restriction of f is in $\mathbb{C}^k(M)$.

We can now complete the main task of this section.

Proposition 2.4.4. The space M together with the cube sets $C^n(M)$ is a compact nilspace.

Proof. By [8, Proposition 3.3.26] we have that M is a nilspace, and by Proposition 2.4.2 we have that M is a compact space. Arguing as in the proof of Lemma 2.1.1, we see that it suffices to show that $C^k(M)$ is a closed subset of $M^{\{0,1\}^k}$. We use the probability space $Q = \hom(T_k, X)$ again, but this time we let $\pi_o: Q \to C^k(X)$ be the map $t \mapsto t \circ \omega_k$, and apply Lemma 2.2.14 (with $P = T_k$ embedded in $\{0,1\}^{2k}$, $P_1 = \emptyset$, and $P_2 = \omega_k(\{0,1\}^k)$ also embedded in $\{0,1\}^{2k}$) to obtain a CSM on π_o consisting of the Haar measures on the fibres $\pi_o^{-1}(c) = \hom_{c \circ \omega_k^{-1}}(T_k, X)$. Let $\tilde{\pi}$ be the projection $M \to X$, $\rho_x + z \mapsto x$, and let

$$R_k(M) = \{ f : \{0,1\}^k \to M : \tilde{\pi} \circ f \in C^k(X) \}.$$

Since $\tilde{\pi}$ is continuous and $C^k(X)$ is closed, we have that $R_k(M)$ is closed in $M^{\{0,1\}^k}$. We define the following map:

$$\begin{array}{cccc} \phi \ : \ R_k(M) & \to & L\left(\pi_o^{-1}(\tilde{\pi}\circ f), Z\right) \subset \mathcal{L}(Q, Z) \\ f & \mapsto & \phi_f(t) = \sum_{v \in \{0,1\}^k} (-1)^{|v|} f(v)(t \circ \Psi_v). \end{array}$$

Arguing as in the proof of [8, Lemma 3.3.32], we see that ϕ_f is actually a constant function for every $f \in R_k(M)$, namely the constant $\rho(\tilde{\pi} \circ f) - \sigma_k(a) \in \mathbb{Z}$. By [8, (3.16)] we have that $C^k(M) = \phi^{-1}(0)$. Hence it now suffices to show that ϕ is continuous, as then $C^k(M)$ is closed as required.

To see the continuity of ϕ , we first express it as the composition of simpler functions. For each $v \in \{0,1\}^k$, let

$$\phi_v: R_k(M) \longrightarrow L(\pi_o^{-1}(\tilde{\pi} \circ f), \mathbb{Z}) \subset \mathcal{L}(Q, \mathbb{Z})$$
 $f \mapsto \phi_{v,f}(t) = f(v)(t \circ \Psi_v).$

We claim that ϕ_v is continuous. To see this, we suppose that $f_n \to f$ in $R_k(M)$, which implies that the component $f_n(v)$ converges to f(v) in M, and we want to show that then $\phi_{v,f_n} \to \phi_{v,f}$ in $\mathcal{L}(Q,\mathbf{Z})$. For this it suffices to show that $\phi_{F_1',F_2'}(\phi_{v,f_n}) \to \phi_{F_1',F_2'}(\phi_{v,f})$ for any continuous functions $F_1': \mathbf{Z} \to \mathbb{C}$, $F_2': Q \to \mathbb{C}$, where

$$\varphi_{F_1',F_2'}(\phi_{v,f}) = \int_{\pi_o^{-1}(\tilde{\pi}\circ f)} F_1'(\phi_{v,f}(t)) \ F_2'(t) \, \mathrm{d}\mu_{\tilde{\pi}\circ f}(t) = \int_{\pi_o^{-1}(\tilde{\pi}\circ f)} F_1'\big(f(v)(t\circ \Psi_v)\big) \ F_2'(t) \, \mathrm{d}\mu_{\tilde{\pi}\circ f}(t).$$

We shall now show that, for one of the functions φ_{F_1,F_2} generating the topology on M, we have $\varphi_{F_1',F_2'}(\phi_{v,f}) = \varphi_{F_1,F_2}(f(v))$. Recall that for any $g \in M$ we have

$$\varphi_{F_1,F_2}(g) = \int_{\pi^{-1}(\tilde{\pi}(g))} F_1(g(c)) F_2(c) d\mu_{\tilde{\pi}(g)}(c).$$
 (2.15)

NOTES ON COMPACT NILSPACES

Now for each $t \in Q$ let Q_t denote the space of morphisms $t' \in Q$ such that $\Psi_v \circ t' = \Psi_v \circ t$, and let μ_{Q_t} denote the Haar measure on this space given by Corollary 2.2.20. These measures form a CSM by Lemma 2.2.14. We then have

$$\begin{split} \phi_{F_1',F_2'}(\phi_{v,f}) &= \int_{\pi_o^{-1}(\tilde{\pi}\circ f)} \int_{\mathcal{Q}_t} F_1'\big(f(v)(t'\circ \Psi_v)\big) \, F_2'(t') \, \mathrm{d}\mu_{\mathcal{Q}_t}(t') \, \mathrm{d}\mu_{\tilde{\pi}\circ f}(t) \\ &= \int_{\pi_o^{-1}(\tilde{\pi}\circ f)} F_1'\big(f(v)(t\circ \Psi_v)\big) \, \Big(\int_{\mathcal{Q}_t} F_2'(t') \, \mathrm{d}\mu_{\mathcal{Q}_t}(t')\Big) \, \mathrm{d}\mu_{\tilde{\pi}\circ f}(t) \\ &= \int_{\pi_o^{-1}(\tilde{\pi}\circ f)} F_1'\big(f(v)(t\circ \Psi_v)\big) \, F_2''(t\circ \Psi_v) \, \mathrm{d}\mu_{\tilde{\pi}\circ f}(t), \end{split}$$

where $F_2''(t \circ \Psi_v) = \int_{t' \in Q_t} F_2'(t') d\mu_{Q_t}(t')$ is a continuous function of $c = t \circ \Psi_v \in C^k(X)$, by Definition 2.2.1. We can now use the fact that $t \mapsto t \circ \Psi_v$ is measure-preserving from $\pi_o^{-1}(\tilde{\pi} \circ f)$ to $\pi^{-1}(\tilde{\pi} \circ f(v))$, to conclude that for $f(v) \in M$ we have

$$\varphi_{F'_1,F'_2}(\phi_{v,f}) = \int_{\pi^{-1}(\tilde{\pi}\circ f(v))} F'_1(f(v)(c)) \ F''_2(c) \, \mathrm{d}\mu_{\tilde{\pi}\circ f}(c).$$

This is of the form $\varphi_{F'_1,F''_2}(f(v))$ as in (2.15), and the continuity of ϕ_v follows.

The continuity of each function ϕ_v implies that the following function is continuous:

$$\phi_{1} : R_{k}(M) \longrightarrow L(\pi_{o}^{-1}(\tilde{\pi} \circ f), Z^{\{0,1\}^{k}}) \subset \mathcal{L}(Q, Z^{\{0,1\}^{k}})
f \mapsto \phi_{1,f}(t) = (f(v)(t \circ \Psi_{v}))_{v \in \{0,1\}^{k}}.$$

Indeed, if $f_n \to f$ in $R_k(M)$ then for each ν by continuity of ϕ_{ν} we have $\phi_{\nu,f_n} \to \phi_{\nu,f}$ in $\mathcal{L}(Q, \mathbb{Z})$. This then implies that $\phi_{1,f_n} \to \phi_{1,f}$ in $\mathcal{L}(Q, \mathbb{Z}^{\{0,1\}^k})$, by Lemma 2.3.9.

The proof can now be completed using Lemma 2.3.15. Indeed, for $P = \{0,1\}^k$ and appropriate choices of coefficients λ_r in that lemma, the original function $\phi : R_k(M) \to \mathcal{L}(Q, \mathbb{Z})$ is the composition $\Sigma_P \circ \phi_1$. The lemma tells us that Σ_P is continuous, and so ϕ is continuous.

Recall from [8, Lemma 3.3.28] that, up to isomorphisms of nilspaces, every extension of X by Z is of the form $M(\rho)$ above for some cocycle ρ . We now want an analogue of this result for compact nilspaces, which requires that the isomorphism be also a homeomorphism. To ensure this, recall first from [8, Lemma 3.3.21] that given a degree-k extension of a nilspace, every cross section s for this extension generates a cocycle, denoted ρ_s . We now show that for compact nilspaces one can take ρ_s to be Borel measurable.

Lemma 2.4.5. Let X be a compact nilspace, and let Y be a compact nilspace that is a degree-k extension of X by a compact abelian group Z. Then there is a Borel cross section s for this extension and therefore a Borel cocycle ρ_s : $C^{k+1}(X) \to Z$.

Proof. Let $\pi: Y \to X$ be the projection of the extension. Let $P = \{(x,y) \in X \times Y : \pi(y) = x\}$. A cross-section for π is a subset of P which happens to be the graph of a function $X \to Y$. Let $\mathcal{P}(Y)$ denote the set of Borel probability measures on Y. By [28, Corollary 18.7], a sufficient condition for a Borel cross section to exist is that for some Borel function $\mu: X \to \mathcal{P}(Y)$ we have $\mu_x(P \cap (\{x\} \times Y)) > 0$.

The measures μ_x from the CSM structure on π satisfy this, and the function $x \to \mu_x$ is continuous, by Definition 2.2.1. There is therefore a Borel cross section $s: X \to Y$. We can then define a Borel function $\rho: C^{k+1}(Y) \to Z$ by $\rho(c') = \sigma_{k+1}(s \circ \pi \circ c' - c')$. Moreover, as is shown in the proof of [8, Lemma 3.3.21], the function ρ is constant on each fibre of the map $C^{k+1}(Y) \to C^{k+1}(X)$, $c' \mapsto \pi \circ c'$, thus yielding the well-defined cocycle $\rho_s: C^{k+1}(X) \to Z$. It can then be seen by general results that ρ_s is also Borel measurable. For instance, one can use that by [8, Definition 3.3.13] the space $C^{k+1}(Y)$ is a continuous abelian bundle over $C^{k+1}(X)$ with Polish structure group $C^{k+1}(D_k(Z))$, so that by [28, Theorem (12.16) and Corollary (15.2)] there is a Borel cross section $c \mapsto c'$ for this bundle.

We shall also need the following useful generalization of a classical automatic-continuity result. The classical result in question is a theorem of Kleppner stating that a Borel measurable homomorphism between two locally compact groups must be continuous [29, Theorem 1]. We extend the case of compact abelian groups as follows.

Theorem 2.4.6. Let X, Y be compact nilspaces of finite step, and let $\phi : X \to Y$ be a Borel measurable morphism. Then ϕ is continuous.

Proof. Suppose that Y is a (k-1)-step nilspace, and consider the CSM on $\pi: C^k(X) \to X$, $c \mapsto c(0^k)$, given by Lemma 2.2.17. For each $v \in \{0,1\}^k \setminus \{0^k\}$ let $\pi_v: C^k(X) \to X$, $c \mapsto c(v)$. We have that π_v is continuous and each of its restrictions to a space $C^k_x(X)$ is measure-preserving, by Lemma 2.2.14 (applied with $P_1 = \{0^k\}$, $P_2 = \{v\}$ and $f: 0^k \mapsto x$, identifying $\lim_{f \to \infty} |P_{1,1} \cap P_{2,1} \cap P_{$

Let $f: X \to \mathcal{L}\left(C^k(X), Y\right)$ be the function mapping $x \in X$ to the constant function with value $\phi(x)$ on $C^k_x(X)$. From the definition of $\mathcal{L}\left(C^k(X), Y\right)$ it follows that if f is continuous then so is ϕ .

To see that f is continuous, recall that $Cor^k(Y)$ denotes the space of k-corners on Y and consider the following function:

$$\xi: \mathbf{X} \to \mathcal{L}\left(\mathbf{C}^k(\mathbf{X}), \mathbf{Cor}^k(\mathbf{Y})\right), \quad x \mapsto \quad \left(\xi_x: \mathbf{c} \in \mathbf{C}^k_x(\mathbf{X}) \ \mapsto \ \left(f_v(x)(\mathbf{c})\right)_{v \neq 0^k} = \left(\phi \circ \mathbf{c}(v)\right)_{v \neq 0^k}\right).$$

Combining the continuity of each function f_v with Lemma 2.3.9, we deduce that ξ is continuous. Now let \mathcal{K} denote the function $\mathcal{L}\left(C^k(X), \operatorname{Cor}^k(Y)\right) \to \mathcal{L}\left(C^k(X), Y\right)$ that sends a function g to $K \circ g$, where $K : \operatorname{Cor}^k(Y) \to Y$ is the unique completion of k-corners on Y. By Lemma 2.1.12 we know that K is continuous, and we claim that \mathcal{K} is therefore continuous. Indeed, suppose we have a sequence g_n in $\mathcal{L}\left(C^k(X),\operatorname{Cor}^k(Y)\right)$ converging to g, and let $F_1: Y \to \mathbb{C}$ and $F_2: C^k(X) \to \mathbb{C}$ be continuous. Then

$$\varphi_{F_1,F_2}(\mathcal{K}(g_n)) \, = \, \int_{C^k_{r_n}(X)} F_1(K \circ g_n(c)) \, F_2(c) \, d\mu(c) \, = \, \varphi_{F_1 \circ K,F_2}(g_n),$$

where $\varphi_{F_1 \circ K, F_2}$ is one of the functions generating the topology on $\mathcal{L}(C^k(X), \operatorname{Cor}^k(Y))$. Hence by assumption this converges to $\varphi_{F_1 \circ K, F_2}(g) = \varphi_{F_1, F_2}(\mathcal{K}(g))$, so \mathcal{K} is continuous.

Finally, note that
$$f = \mathcal{K} \circ \xi$$
, so f is continuous, which completes the proof.

We are now able to prove the analogue of [8, Lemma 3.3.28] for compact nilspaces.

Lemma 2.4.7. Let X be a compact nilspace of finite step. Let Y be a degree-k extension of X by a compact abelian group Z. Let $s: X \to Y$ be a Borel cross-section and let $\rho = \rho_s$ be the associated Borel cocycle. Then Y is isomorphic as a compact nilspace to the extension $M(\rho)$ from Proposition 2.4.2.

Proof. Let $\pi: Y \to X$ be the projection for the extension. The isomorphism is given by the following map:

$$\theta: Y \to M, \ x \mapsto \rho_{\pi(x)} + (x - s \circ \pi(x)).$$
 (2.16)

It is checked in the proof of [8, Lemma 3.3.28] that θ is an isomorphism of nilspaces. Since Y and M are both compact Hausdorff spaces, if we show that θ is continuous then it is a homeomorphism [32, Theorem 26.6]. Since θ is Borel measurable, it is continuous by Theorem 2.4.6.

Recall from [8, Subsection 3.3.3] that $H_k(X,Z)$ denoted the quotient of the abelian group of degree-k cocycles by the subgroup of coboundaries. In the category of compact nilspaces we restrict these groups to the Borel measurable cocycles. Combining the results from this section, we now deduce the following analogue of [8, Corollary 3.3.29] for compact nilspaces.

Corollary 2.4.8. Let Φ denote the map which sends each class $C \in H_k(X, \mathbb{Z})$ to the isomorphism class of $M(\rho)$, for any choice of $\rho \in C$. Then Φ is a surjection from $H_k(X, \mathbb{Z})$ to the set of isomorphism classes of compact degree-k extensions of X by Z.

The proof is similar to that of [8, Corollary 3.3.29].

2.5 Compact abelian bundles of finite rank, and averaging

In this section we define a notion of rank for compact abelian bundles and nilspaces, and we give a first topological description of nilspaces of finite rank (Lemma 2.5.3). We then define an averaging operation for certain functions on these spaces, which enables us to prove a useful rigidity result concerning cocycles (Lemma 2.5.7).

Definition 2.5.1. Let B be a k-fold compact abelian bundle, with structure groups Z_1, \ldots, Z_k . We define the *rank* of B by

$$\mathrm{rk}(\mathrm{B}) = \sum_{i=1}^{k} \mathrm{rk}(\widehat{\mathrm{Z}}_{i})$$

where \widehat{Z}_i is the Pontryagin dual of Z_i and $\operatorname{rk}(\widehat{Z}_i)$ is the minimal number of generators of \widehat{Z}_i . If X is a k-step compact nilspace then we define its rank $\operatorname{rk}(X)$ to be $\operatorname{rk}(B)$ for the associated bundle B given by Proposition 2.1.9.

When working with a compact Z-bundle B over S, we shall often want to ensure that the bundle is locally trivial, meaning that for every $x \in B$ there is an open set $U \subset S$ containing the projection $\pi(x)$, and a homeomorphism $\pi^{-1}(U) \to U \times Z$ of the form $x \mapsto (\pi(x), \varphi(x))$ where $\varphi : \pi^{-1}(U) \to Z$ is Z-equivariant (that is $\varphi(x+z) = \varphi(x) + z$ for every $x \in \pi^{-1}(U)$, $z \in Z$). A theorem of Gleason [14, Theorem 3.3] tells us that this local triviality holds provided that Z is a Lie group (see also [1, Theorem 1.1]). We record the special case of this result that we shall use.

P. CANDELA

Proposition 2.5.2 (Consequence of Gleason's slice theorem). *Let* Z *be a compact abelian Lie group, and let* B *be a continuous* Z*-bundle. Then* B *is locally trivial.*

Recalling that a compact abelian group is a Lie group if and only if it has finite rank, we deduce the following result.

Lemma 2.5.3. Let X be a k-step compact nilspace of finite rank with structure groups $Z_1, ..., Z_k$ and factors $X_0, ..., X_k$. Then X_i is a locally trivial Z_i -bundle over X_{i-1} for each $i \in [k]$.

Thus, topologically a k-step compact nilspace of finite rank is a finite dimensional manifold.

Definition 2.5.4. Recall that every compact abelian Lie group Z is isomorphic to a direct product $F \times \mathbb{T}^n$ for some finite abelian group F and $n \ge 0$. We can then define a natural metric d_2 on Z as follows. For two elements $x, y \in \mathbb{T}^n = \mathbb{R}^n/\mathbb{Z}^n$ we define $d_2(x,y)$ as the minimum of the Euclidean distances between preimages of x and y under the map $\mathbb{R}^n \to \mathbb{R}^n/\mathbb{Z}^n$. If Z is not connected, then for points x, y in different connected components we declare that $d_2(x,y) = \infty$.

We now turn to the definition of averaging for a function on a nilspace taking values in a compact abelian group Z. Here we face the problem that in general the average of a Z-valued function is not well-defined. This is clear already in the simple case of the circle group $Z = \mathbb{T}$, where there is no suitable notion of multiplication by a real number. However, if we assume that the function takes all its values in a short interval, then we can lift this \mathbb{T} -valued function to a real-valued one, then take the usual average, and then project the obtained value back to \mathbb{T} . We capture this idea more generally as follows.

Definition 2.5.5. Let $Z \cong F \times \mathbb{T}^n$ be a compact abelian Lie group. Let f be a Borel measurable function from a probability space X to Z, and suppose there exists $z = (g, \theta) \in Z$, $(\theta \in \mathbb{T}^n)$ such that every value of f lies in the ball $B_{1/4}(z)$ in the metric d_2 . Then, letting π denote the projection $\mathbb{R}^n \to \mathbb{T}^n$, for any fixed lift $\theta' \in \mathbb{R}^n$ there is a unique function $f': X \to \mathbb{R}^n$ such that $f(x) = (g, \pi \circ f'(x))$ and f' takes all its values in $B_{z'}(1/4)$ in the euclidean distance. We then define the average value $\mathbb{E}f$ to be $(g, \pi(\mathbb{E}f'))$.

Note that $\mathbb{E}(f)$ does not depend on the choice of the lift of z. Let us say that a function $f: X \to Z$ is δ -concentrated if there exists $z \in Z$ such that $f(x) \in B_{\delta}(z)$ for all $x \in X$.

Lemma 2.5.6. Let Z be a compact abelian Lie group, and let $m \in \mathbb{N}$. Then for any (1/4m)-concentrated measurable functions f_1, \ldots, f_m from a probability space X to Z, the average of $f_1 + \cdots + f_m$ is well defined and satisfies $\mathbb{E}(f_1 + \cdots + f_m) = \mathbb{E}f_1 + \cdots + \mathbb{E}f_m$.

Proof. The assumption implies that $f_1 + \cdots + f_m$ is 1/4-concentrated, so its average is well-defined. The additivity then follows in a straightforward way from the additivity of the usual average for real-valued functions.

As a first application of this notion of averaging, we obtain the following rigidity result for cocycles. Stronger versions of this result will be very useful in later sections. Recall from [8, Definition 3.3.18] that a function $C^k(X) \to Z$ is called a coboundary if it is of the form $c \mapsto \sigma_k(f \circ c)$ for some function $f: X \to Z$.

Lemma 2.5.7. Let X be a compact ℓ -step nilspace and let Z be a compact abelian group of finite rank. There exists $\varepsilon > 0$ such that, for every Borel measurable cocycle $\rho : C^k(X) \to Z$, if $d_2(\rho(c), 0) \le \varepsilon$ for every $c \in C^k(X)$ then there exists a Borel measurable function $g : X \to Z$ such that ρ is the coboundary $c \mapsto \sigma_k(g \circ c)$.

Proof. For any $x \in X$, consider the space of cubes $C_x^k(X)$ with the Haar probability given by Lemma 2.2.17. By assumption ρ is ε -concentrated on $C_x^k(X)$, so we may define

$$g: X \to Z, x \mapsto \mathbb{E}_{C_x^k(X)} \rho(c).$$

Fix any $c \in C^k(X)$ and let $\Omega = \hom_{c \circ \omega_k^{-1}}(T_k, X)$. For any $t \in \Omega$, by [8, Lemma 3.3.31] we have $\rho(c) = \sum_{v \in \{0,1\}^k} (-1)^{|v|} \rho(t \circ \Psi_v)$. By Corollary 2.2.22, we have that for each v the map $t \mapsto t \circ \Psi_v$, $\Omega \to C^k_{c(v)}(X)$ preserves the Haar measures (where the Haar measure on Ω is given by Corollary 2.2.20). Hence, by averaging over t in the last equation, we obtain $\rho(c) = \sum_{v \in \{0,1\}^k} (-1)^{|v|} g(c(v)) = \sigma_k(g \circ c)$. \square

2.6 Counting isomorphism classes of compact finite-rank nilspaces

In this section we use the parametrization of compact extensions by cocycles, given in Corollary 2.4.8, to prove the following result.

Theorem 2.6.1. There are countably many isomorphism classes of compact nilspaces of finite rank.

To prove this we shall use the following lemma, that strengthens Lemma 2.5.7 by allowing its premise to fail on a null set.

Lemma 2.6.2. Let X be a compact ℓ -step nilspace and Z be a compact abelian group of finite rank. Then there exists $\varepsilon > 0$ such that every Borel measurable cocycle $\rho : C^k(X) \to Z$ satisfying $d_2(\rho(c), 0) \le \varepsilon$ for almost every $c \in C^k(X)$ is a coboundary.

This relies on the following fact.

Lemma 2.6.3. Let X be a compact ℓ -step nilspace and let Z be a compact abelian group of finite rank. Let $\rho: C^k(X) \to Z$ be a Borel measurable cocycle such that $\rho = 0$ for almost every element in $C^k(X)$. Then ρ is a coboundary.

Proof. Let $S = \{x \in X : \rho_x = 0 \text{ almost surely on } C_x^k(X)\}$. It follows from the properties of the CSM on $C^k(X) \to X$, $c \mapsto c(0^k)$ (given by Lemma 2.2.17) that S has probability 1. For every $x \in S$ the set $\rho_x + Z \subset \mathcal{L}_k(X,Z)$ consists of functions that are constant almost everywhere. Since by Proposition 2.2.11 every open set in X has positive measure, we have that S is dense in X. We claim that for every $x \in X$ the function ρ_x equals a constant almost everywhere. Indeed, by density of S there exists a sequence of points $x_n \in S$ converging to x. For each n choose an element $\rho_{x_n} + a_n \in M(\rho)$ equal to a constant almost surely. By Proposition 2.4.2 the extension M is compact and so there is a subsequence of such functions $\rho_{x_m} + a_m$ converging in M. By Lemma 2.3.8, the limit is constant almost everywhere, and (as a point of M) is $\rho_x + a$ for some $a \in Z$, which proves our claim. From this, we deduce that the function $g: X \to Z$,

 $x\mapsto \mathbb{E}_{c\in C^k_x(X)} \rho(c)$ is well-defined and that for each $x\in X$ we have $\rho(c)=g(x)$ for almost every $c\in C^k_x(X)$. We now claim that for every $x\in X$, we have $\rho(c)=\sigma_k(g\circ c)$ for almost every $c\in C^k(X)$. This follows from an argument similar to the second paragraph in the proof of Lemma 2.5.7. Thus, denoting by f the coboundary $c\mapsto \sigma_k(g\circ c)$, the cocycle $\rho'=\rho-f$ has the property that for every $x\in X$ we have $\rho'_x=0$ almost surely on $C^k_x(X)$. This implies that $\rho'=0$ everywhere. Indeed, for an arbitrary fixed $c\in C^k(X)$, and any $t\in \hom_{c\circ \omega_k^{-1}}(T_k,X)$, by [8, Lemma 3.3.31] we have $\rho'(c)=\beta(t,\rho')$, and then, averaging both sides of this equation over t, the right side vanishes since ρ'_x is 0 almost everywhere (using again that each map $t\mapsto t\circ \Psi_v$ is probability-preserving, by Corollary 2.2.22).

Proof of Lemma 2.6.2. We argue as in the proof of Lemma 2.5.7. Let S be the set of elements $x \in X$ for which ρ_x is almost surely within ε of 0 in d_2 . By the assumption, we have that S has probability 1. We define $g: X \to Z$ for $x \in S$ by $g(x) = \mathbb{E}_{C_x^k(X)} \rho(c)$, and for $x \in X \setminus S$ by setting g(x) = 0. Letting $\rho_2(c) = \sigma_k(g \circ c)$, the same averaging argument shows that $\rho_2 = \rho$ almost surely. By Lemma 2.6.3 the difference $\rho - \rho_2$ is a coboundary, whence ρ is a coboundary.

We can now prove the main result of this section.

Proof of Theorem 2.6.1. We claim that for every $k \in \mathbb{N}$, for every compact (k-1)-step nilspace X of finite rank and every compact abelian Lie group Z, there are at most countably many non-isomorphic degree-k extensions of X by Z. This implies the theorem by induction on k.

To prove the claim, we first associate with each Z-extension of X a measurable cocycle ρ_s generated by a piecewise-continuous cross section s for the extension. We construct s as follows.

Since X is a finite-dimensional compact manifold, there exist disjoint open sets U_1, U_2, \ldots, U_r , each homeomorphic to the open unit Euclidean ball in \mathbb{R}^n for some $n \geq 0$, such that $U = \bigcup_{i \in [r]} U_i$ has measure 1 in X. Let M denote the given degree-k extension of X by Z, with projection $\pi: M \to X$. Every Z-bundle over a contractible space is trivial [27, Ch. 4, Corollary 10.3]. Therefore there exists a continuous cross section s_i locally on each U_i , and combining these we obtain a piecewise continuous (hence Borel-measurable) cross section $s: X \to M$, with $s|_{U_i} = s_i$. Let $T \subseteq C^{k+1}(X)$ denote the set of cubes whose vertices are all in U. We can partition the cubes $c \in T$ according to which set U_i contains each value c(v), that is $T = \bigsqcup_{i \in [r]^{\{0,1\}^{k+1}}} T_i$, where $T_i = T \cap \prod_{v \in \{0,1\}^{k+1}} U_{i_v}$. The cocycle $\rho_s: C^{k+1}(X) \to Z$ is defined by $\rho_s(c) = \sigma_{k+1}(s \circ c - c')$ for any $c' \in C^{k+1}(M)$ with $\pi \circ c' = c$ (recall [8, Lemma 3.3.21]). Note that ρ_s is continuous on each part T_i .

Now let $\varepsilon > 0$ be as given by Lemma 2.6.2. Then for any other degree-k extension M', we can generate a cocycle $\rho_{s'}$ as above, and if $d_2(\rho_s(c), \rho_{s'}(c)) < \varepsilon$ for almost every c then by the lemma $\rho_{s'}$ must equal ρ_s plus a coboundary, which implies, by Corollary 2.4.8, that M' is isomorphic to M.

Let d_{∞} denote the uniform metric on the set of continuous functions $C^{k+1}(X) \to Z$, defined by $d_{\infty}(f_1, f_2) = \sup_{c} d_2(f_1(c), f_2(c))$. For each i, the space of continuous functions $T_i \to Z$ with d_{∞} is separable (essentially by separability of the space C(Y,Z) for any compact subset Y of a Euclidean space [28, Theorem (4.19)]). Therefore there exists a countable $\frac{\varepsilon}{2}$ -net $(f_{i,j})_{j\in\mathbb{N}}$ of such continuous functions on each T_i . Then, by the previous paragraph, for each choice of functions $\{f_{i,j_i}: i \in [r]^{\{0,1\}^{k+1}}\}$ from these nets, there can only be at most one isomorphism class of extensions having an associated cocycle ρ_s satisfying $d_{\infty}(\rho_s|_{T_i}, f_{i,j_i}) < \varepsilon/2$ for each i.

2.7 Compact nilspaces as inverse limits of finite-rank nilspaces

In this section we treat one of the central results from [9], namely that every compact nilspace can be expressed as an inverse limit of compact nilspaces of finite rank. Before we give the formal statement, let us detail the inverse limit construction in this category.

Definition 2.7.1. An *inverse system* (or projective system) of compact nilspaces over \mathbb{N} is a family of continuous nilspace morphisms $\{\varphi_{ij}: X_j \to X_i \mid i, j \in \mathbb{N}, i \leq j\}$, where $X_j, j \in \mathbb{N}$ are compact nilspaces, such that φ_{jj} is the identity morphism for all $j \in \mathbb{N}$, and for all $i, j, k \in \mathbb{N}$ with $i \leq j \leq k$ we have $\varphi_{ij} \circ \varphi_{jk} = \varphi_{ik}$. The inverse system is said to be *strict* if every morphism φ_{ij} is fibre-surjective.

Lemma 2.7.2. Let $S = (\varphi_{ij} : X_i \to X_i)_{i \le j}$ be a strict inverse system of compact nilspaces. Let

$$X = \left\{ (x_i) \in \prod_{i \in \mathbb{N}} X_i \mid \varphi_{ij}(x_j) = x_i \ \forall i, j \in \mathbb{N}, \ i \leq j \right\},\,$$

and for each $n \ge 0$ let $C^n(X) = \{(c_i) \in \prod_{i \in \mathbb{N}} C^n(X_i) \mid \varphi_{ij} \circ c_j = c_i \ \forall i, j \in \mathbb{N}, i \le j\}$. The space X together with the cube sets $C^n(X)$ is a compact nilspace, called the inverse limit of S, and denoted $\varprojlim_i X_i$. We have that X is k-step if and only if every X_i is k-step.

We call the maps φ_{ij} the *transition morphisms*. The *projections* on X, denoted by φ_j , are the coordinate projections on $\prod_{i\in\mathbb{N}}X_i$ restricted to X. If the inverse system is strict then the projections are also fibre-surjective morphisms. The ℓ -th structure group of $\varprojlim_i X_i$ is the inverse limit of the system $(\alpha_{\ell,ij}: \mathbb{Z}_{\ell,j} \to \mathbb{Z}_{\ell,i})_{i\leq j}$, where $\mathbb{Z}_{\ell,j}$ is the ℓ -th structure group of X_j and $\alpha_{\ell,ij}$ is the ℓ -th structure morphism of φ_{ij} .

Proof. A straightforward argument with convergent sequences shows that X is a closed subspace of the compact space $\prod_{i\in\mathbb{N}} X_i$, and similarly for each set $C^n(X)$. The ergodicity and composition axioms are clear. Let us prove that X satisfies the completion axiom. Let c' be an n-corner on X. Then for each $i \in \mathbb{N}$ the projection $c'_i = \varphi_i \circ c'$ is an *n*-corner on X_i . Now c'_1 has a completion c_1 and by fibre-surjectivity there exists $\tilde{c}_1 \in C^n(X_2)$ such that $\varphi_{12} \circ \tilde{c}_1 = c_1$ (recall [8, Lemma 3.3.9]). The restriction of \tilde{c}_1 to $P = \{0,1\}^n \setminus \{1^n\}$ is in the same fibre of the map $c' \mapsto \varphi_{12} \circ c'$ as c'_2 , so by [8, Lemma 3.3.12 (ii)] these two corners differ by an *n*-corner on $\mathcal{D}_k(\ker(\alpha_k))$, where α_k is the *k*-th structure morphism of φ_{12} . Now by modifying \tilde{c}_1 , we can obtain a cube c_2 that is still in this fibre (so that $\varphi_{12} \circ c_2 = c_1$) and such that $c_2|_P = c_2'|_P$. Indeed, we can obtain c_2 by adding first an appropriate element of $\ker(\alpha_k)$ to the values of \tilde{c}_1 at the vertices of the 1-face $\{0^n, (1,0,\ldots,0)\}$ so that it agrees with c_2' at 0^n , then repeat this for each subsequent 1-face along a Hamiltonian path in the graph of 1-faces on $\{0,1\}^n$ (similar arguments were used several times in [8], for instance in the proof of [8, Lemma 3.2.25]). We have thus obtained a cube c_2 completing c_2' and such that $\varphi_{12} \circ c_2 = c_1$. Now we repeat this operation to obtain c_3 completing c_3' such that $\varphi_{23} \circ c_3 = c_2$, and so on. This yields a sequence of cubes $c_i \in C^n(X)$ such that for each i we have $\varphi_i \circ c_i$ is a completion of c_i' for all $j \le i$. The limit of a convergent subsequence of (c_i) is a completion of c'. To see the last claim in the lemma, recall from [8, Definition 3.3.7] that a fibre-surjective morphism preserves the property of being k-step.

We can now state the main result of this section.

Theorem 2.7.3. Every k-step compact nilspace is an inverse limit of compact nilspaces of finite rank.

The basic case of this theorem, concerning 1-step nilspaces, is given by the standard result that every compact abelian group is an inverse limit of compact abelian Lie groups (see [24, Corollary 2.43]). This standard result also yields the following example of an inverse limit of nilspaces, which will be used in the proof of Theorem 2.7.3.

Proposition 2.7.4. Let X be a k-step compact nilspace. Then X is an inverse limit of k-step nilspaces X_i , each with k-th structure group a compact abelian Lie group, and with $\mathcal{F}_{k-1}(X_i)$ isomorphic to $\mathcal{F}_{k-1}(X)$.

Proof. By the standard result for compact abelian groups recalled above, for the k-th structure group Z_k of X we have for each $i \in \mathbb{N}$ a compact abelian Lie group $Z_{k,i}$ and a surjective continuous homomorphism $\alpha_i : Z_k \to Z_{k,i}$ such that $Z_k = \varprojlim Z_{k,i}$. (For $i \le j$ the transition morphism $\alpha_{i,j}$ is $x \mapsto x + \ker \alpha_i$ for $x \in Z_{k,j} \cong Z_k / \ker \alpha_j$.) Define for each i the k-step compact nilspace X_i as the image of X under the map φ_i that sends $x \in X$ to the orbit $x + \ker \alpha_i$. A straightforward calculation shows that X_i , with the quotient cube-structure of X by $\ker \alpha_i$ and the quotient topology, is a k-step compact nilspace with k-th structure group $Z_{k,i}$. The projections $X \to \mathcal{F}_{k-1}(X)$ and $X_i \to \mathcal{F}_{k-1}(X_i)$ are both given by the orbit map for the Z_k -action, and it follows that $\mathcal{F}_{k-1}(X) \cong \mathcal{F}_{k-1}(X_i)$ as compact nilspaces. The transition morphisms $\varphi_{ij} : X_j \to X_i$, defined by $x + \ker \alpha_j \mapsto x + \ker \alpha_i$, are all clearly fibre-surjective.

Recall from [8, Lemma 3.3.21] that given a degree-k extension Y of X by Z, and given a cross section $s: X \to Y$, letting $f: Y \to Z$, $y \mapsto s \circ \pi(y) - y$, the cocycle generated by s is the degree-k cocycle $\rho_s: C^{k+1}(X) \to Z$ defined by $\rho_s(c) = \sigma_{k+1}(f \circ c')$, for any $c' \in C^{k+1}(Y)$ such that $\pi \circ c' = c$.

For the proof of Theorem 2.7.3, it will be useful to be able to tell in a simple way whether a given cocycle on a nilspace X induces a well-defined cocycle on another nilspace X' via some given fibre-surjective morphism $X \to X'$. The following definition provides such a criterion.

Definition 2.7.5. Let X be a k-step compact nilspace, and for some compact nilspace X' let Ψ : $\mathcal{F}_{k-1}(X) \to X'$ be a continuous fibre-surjective morphism. We say that a measurable cross section $s: \mathcal{F}_{k-1}(X) \to X$ is consistent with the factor X' if the cocycle $\rho_s: C^{k+1}(\mathcal{F}_{k-1}(X)) \to Z_k$ satisfies $\rho_s(c_1) = \rho_s(c_2)$ for every c_1, c_2 such that $\psi \circ c_1 = \psi \circ c_2$.

Thus if s is consistent with X' then ρ_s induces a cocycle ρ' : $C^{k+1}(X') \to Z_k$, well-defined by $\rho'(c') = \rho_s(c)$ for any c such that $\psi \circ c = c'$. Note that we shall often refer to the image of a fibre-surjective morphism on X as a *fibre-surjective factor* of X.

Lemma 2.7.6. Let $\psi : \mathcal{F}_{k-1}(X) \to X'$ be a fibre-surjective morphism and let $s : \mathcal{F}_{k-1}(X) \to X$ be a cross section consistent with X'. Let \sim be the equivalence relation on X defined by

$$x \sim y \iff (\psi \circ \pi_{k-1}(x) = \psi \circ \pi_{k-1}(y) \text{ and } x - s \circ \pi_{k-1}(x) = y - s \circ \pi_{k-1}(y)).$$

Then the quotient nilspace $X/_{\sim}$ is a fibre-surjective factor of X which is an extension of X' by Z_k .

Proof. Let $\rho': \mathbb{C}^{k+1}(X') \to \mathbb{Z}_k$ be the cocycle induced by ρ_s . Let $M(\rho')$ be the compact nilspace extending X' by \mathbb{Z}_k given by Proposition 2.4.2. Let $f: X \to M(\rho')$ be defined by $f(x) = \rho'_{\psi(\pi_{k-1}(x))} + x - s \circ \pi_{k-1}(x)$. Using Definition 2.4.3 we check that f is a morphism (a similar calculation appears in the proof of [8, Lemma 3.3.28]). It is also checked easily that f is fibre-surjective. Moreover, clearly $f(x_1) = f(x_2)$ if and only if $x_1 \sim x_2$. It follows that f factors through \sim giving a bijective measurable morphism $f': X/_{\sim} \to Y$. Then f' is a continuous isomorphism of compact nilspaces, by Theorem 2.4.6.

The following lemma will be used in the proof of Theorem 2.7.3 and also in Section 2.9. Recall the notation d_2 for the metric from Definition 2.5.4.

Lemma 2.7.7. Let X be a k-step compact nilspace such that the k-th structure group Z_k has finite rank, and let d be a compatible metric on X. Then for any $\varepsilon > 0$ there exists $\delta > 0$ such that the following holds. Let $x_0, x_1, y_0, y_1 \in X$ be such that $\pi_{k-1}(x_i) = \pi_{k-1}(y_i)$ for i = 0, 1 and $d(x_0, x_1), d(y_0, y_1)$ are both at most δ . Then in Z_k we have $d_2(y_0 - x_0, y_1 - x_1) \leq \varepsilon$.

Proof. Consider the local translation $\phi = \phi_{x_0,x_1}$ (recall [8, Definition 3.2.26]). By definition of ϕ we have $y_0 - x_0 = \phi(y_0) - x_1$ in \mathbb{Z}_k (using implicitly the isomorphism from [8, Lemma 3.2.24]), so it suffices to show that $d_2(\phi(y_0) - x_1, y_1 - x_1) \le \varepsilon$ for δ sufficiently small.

Since $d(x_0, x_1)$ is unchanged by adding $y_0 - x_0 = \phi(y_0) - x_1$ to x_0 and x_1 , we have $d(y_0, \phi(y_0)) \le \delta$. This together with $d(y_0, y_1) \le \delta$ implies that $d(\phi(y_0), y_1) \le 2\delta$. Thus, the lemma will follow if we prove the following statement: for every $\varepsilon > 0$ there exists $\delta > 0$ such that whenever $x_1, x_2, x_3 \in X$ are in the same fibre of π_{k-1} and $d(x_2, x_3) \le \delta$, we have $d_2(x_2 - x_1, x_3 - x_1) \le \varepsilon$. This can be shown by a straightforward argument using the local triviality of the Z_k -bundle X and the compactness of X_{k-1} . \square

2.7.1 Proof of the inverse limit theorem

To prove Theorem 2.7.3, we argue by induction on k, starting with the trivial case k=0. For k>0 we suppose that X is a k-step compact nilspace such that the nilspace $Y=\mathcal{F}_{k-1}(X)$ is the inverse limit of a strict inverse system $(\tau_{ij}:Y_j\to Y_i)_{i,j\in\mathbb{N},i\leq j}$. We denote by π the projection $X\to Y$, and for each i we denote by τ_i the projection $Y\to Y_i$ and let $\Omega_i=\{\tau_i^{-1}(Q):Q\subset Y_i,Q\text{ open}\}$. Since by assumption the topology on Y is the initial topology generated by the maps τ_i , we have that the finite intersections of sets in $\bigcup_{i\in\mathbb{N}}\Omega_i$ form a base for this topology.

The inductive step: construction of the sequence S

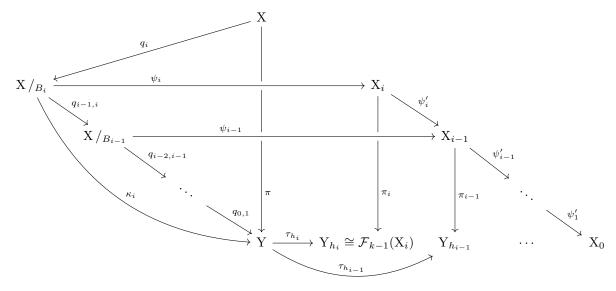
We now apply Proposition 2.7.4 to X. Let $B_0 = \mathbb{Z}_k$ be the k-th structure group of X, and let $B_i = \ker \alpha_i \leq \mathbb{Z}_k$ for each $i \in \mathbb{N}$ (for the homomorphisms α_i from the proof of the proposition). Then $B_0 \geq B_1 \geq \cdots$, with $\bigcap_{i \in \mathbb{N}} B_i = \{0\}$, and $X = \varprojlim_i X/B_i$. Let $q_{i-1,i}$ denote the transition morphism $X/B_i \to X/B_{i-1}$ and q_i the projection $X \to X/B_i$. Note that X/B_i has k-th structure group \mathbb{Z}_k/B_i , of finite rank.

The main goal of the inductive step is to construct, starting with the 1-point nilspace X_0 , a sequence $S = \{(X_i, \psi_i, \psi_i', h_i) : i \in \mathbb{N}\}$, where for each i we have a compact nilspace X_i of finite rank, and fibre-surjective morphisms $\psi_i : X/_{B_i} \to X_i$ and $\psi_i' : X_i \to X_{i-1}$, with the following properties:

- (i) For every *i*, the restriction of ψ_i to each class of \sim_{k-1} in X/B_i is injective. In other words, the *k*-th structure group of X_i is the same as that of X/B_i , namely Z_k/B_i .
- (ii) For every *i*, we have $\psi'_i \circ \psi_i = \psi_{i-1} \circ q_{i-1,i}$.
- (iii) Letting π_i denote the projection $X_i \mapsto \mathcal{F}_{k-1}(X_i)$, and κ_i denote the projection $X/B_i \to Y$, we have $\mathcal{F}_{k-1}(X_i) \cong Y_{h_i}$ and $\pi_i \circ \psi_i = \tau_{h_i} \circ \kappa_i$.

Properties (i) and (iii) imply that any two distinct points in X are separated by a map $\psi_i \circ q_i$ for some i, so the initial topology on X generated by these maps is Hausdorff. Since the original topology on X is compact, it follows that this initial topology equals the original topology (see [6, §9.4, Corollary 3]). The

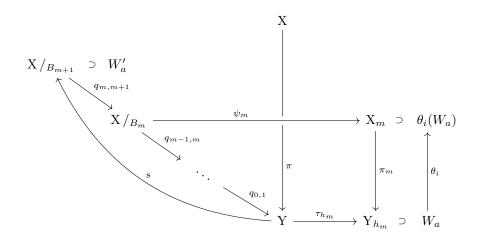
maps $\varphi_{ij} = \psi'_{i+1} \circ \cdots \circ \psi'_j$ and $\varphi_i := \psi_i \circ q_i$ satisfy the relations $\varphi_{ij} \circ \varphi_j = \varphi_i$ thanks to property (ii). Hence X is the inverse limit of the system $(\varphi_{ij} : X_j \to X_i)_{i \le j}$.



We shall now construct the sequence S inductively starting from X_0 . Suppose that $(X_i, \psi_i, \psi_i', h_i)$ has been constructed for $i \le m$. Then, to begin with, we construct a Borel cocycle $C^{k+1}(Y) \to Z_k/B_{m+1}$ generated by a cross section s that is almost consistent with some appropriate factor of Y.

Constructing a measurable cross section s: since X_m is a compact locally-trivial Z_k/B_m -bundle over Y_{h_m} , there is a finite cover of Y_{h_m} by closed subsets W_1, \ldots, W_r such that every preimage $\pi_m^{-1}(W_i) \subset X_m$ is a trivial Z_k/B_m -bundle over W_i . For each $i \in [r]$, let $\theta_i : W_i \to X_m$ be a continuous cross section (thus $\pi_m \circ \theta_i$ is the identity map on W_i).

For each $a \in [r]$ let W'_a denote the preimage of the set $\theta_i(W_a)$ under $\psi_m \circ q_{m,m+1} : X/_{B_{m+1}} \to X_m$. Note that W'_a is a Z_k/B_{m+1} -bundle over $\tau_{h_m}^{-1}(W_a) \subset Y$, so it is locally trivial by Proposition 2.5.2. Let d be a metric on $X/_{B_{m+1}}$ generating its topology.



Now fix an arbitrary $\varepsilon > 0$. We claim that for every $p \in \tau_{h_m}^{-1}(W_a)$ there exists an open set $U_p \subset \tau_{h_m}^{-1}(W_a)$ containing p with the following properties:

- (iv) There exists a continuous cross section $s_p: U_p \to X/_{B_{m+1}}$.
- (v) The diameter of $s_p(U_p)$ in the metric d is at most ε .
- (vi) For some $t(p) \in \mathbb{N}$, the set U_p is in the collection $Q_{t(p)}$ defined above.

Property (iv) is ensured by the local triviality of the bundle, and it can be satisfied together with (v) by letting U_p be a ball of sufficiently small radius, say. Property (vi) can also be satisfied since $\bigcup_i \Omega_i$ generates the topology on Y and for every $i \leq j$ and open $Q \subset Y_i$ we have $\tau_i^{-1}Q = \tau_j^{-1}(\tau_{ij}^{-1}Q) \in \Omega_j$. Since Y is compact there exists a finite cover by such sets U_p , thus for some points p_1, p_2, \ldots, p_n we have $\bigcup_{i=1}^n U_{p_i} \supset \tau_{h_m}^{-1}(W_a)$. Let $t_a = \max\{t(p_i) : i \in [n]\}$, so that for every $i \in [n]$ we have $U_{p_i} \subset \Omega_{t_a}$. Then, by dividing $\tau_{h_m}^{-1}(W_a)$ into the atoms of the Boolean algebra generated by U_{p_1}, \ldots, U_{p_n} , and then using one of the cross sections s_p for each atom, we can construct a Borel measurable cross section $s_a : \tau_{h_m}^{-1}(W_a) \to X/B_{m+1}$ with the following properties:

- (vii) The cross section s_a is continuous on $\tau_{t_a}^{-1}(v)$ for every $v \in Y_{t_a}$.
- (viii) The diameter of $\mathbf{s}_a(\tau_{t_a}^{-1}(v))$ in d is at most ε for every $v \in \mathbf{Y}_{t_a}$.

Note that we obtain (viii) from (v) because each $\tau_{t_a}^{-1}(v)$ is entirely contained in an atom of the Boolean algebra, which contains $\tau_{t_a}^{-1}(Q)$ for some open $Q \supset v$. Let $t = \max\{t_a : a \in [r]\} \cup \{h_m + 1\}$. Using these partial cross sections, we now claim that we can construct a global Borel cross section $s : Y \to X/B_{m+1}$ with the following properties:

- (ix) The cross section s is continuous on every preimage $\tau_t^{-1}(v)$ where $v \in Y_t$.
- (x) The diameter of $s(\tau_t^{-1}(v))$ in d is at most ε for every $v \in Y_t$.
- (xi) If $v_1, v_2 \in Y$ satisfy $\tau_{h_m}(v_1) = \tau_{h_m}(v_2)$, then $\psi_m \circ q_{m,m+1}(s(v_1)) = \psi_m \circ q_{m,m+1}(s(v_2))$.

Indeed, we can construct s by again dividing Y into the atoms of the Boolean algebra generated by the sets $\tau_{h_m}^{-1}(W_a)$ and then using one of the cross sections s_a on each atom. This immediately gives (ix) and (x), and then (xi) follows from (iii) above.

Let $\rho_s: C^{k+1}(Y) \to Z_k/B_{m+1}$ be the cocycle generated by s.

Now for any fixed $\varepsilon_2 > 0$, we note that if $\varepsilon > 0$ is small enough, then property (x) implies the following fact, which tells us that s is almost consistent with the factor Y_t :

$$\forall c_1, c_2 \in C^{k+1}(Y) \text{ with } \tau_t \circ c_1 = \tau_t \circ c_2, \text{ we have } d_2(\rho_s(c_1) - \rho_s(c_2)) \le \varepsilon_2. \tag{2.17}$$

Indeed $\tau_t \circ c_1 = \tau_t \circ c_2$ tells us that for each v the values $c_1(v), c_2(v)$ are in the same fibre of τ_t and so by (x) we have that $d(s \circ c_1(v), s \circ c_1(v)) \leq \varepsilon$ for all v. Then by continuity of κ_{m+1} the values $c_i(v), i=1,2$ are also close in Y for each v. Therefore by Lemma 2.1.13 there exist $\tilde{c}_1, \tilde{c}_2 \in C^{k+1}(X/B_{m+1})$ such that $d(\tilde{c}_1(v), \tilde{c}_2(v)) \leq \varepsilon_2$ for every v and such that $\kappa_{m+1} \circ \tilde{c}_i = c_i$, i=1,2. Finally, from the expression $\rho_s(c_i) = \sigma_{k+1}(s \circ c_i - \tilde{c}_i)$ we obtain the inequality in (2.17) by applying Lemma 2.7.7.

Note that property (xi) also implies that

$$\forall c_1, c_2 \in C^{k+1}(Y) \text{ with } \tau_{h_m} \circ c_1 = \tau_{h_m} \circ c_2, \text{ we have } \rho_s(c_1) - \rho_s(c_2) \in B_m/B_{m+1}.$$
 (2.18)

Indeed, since ψ_m is injective by (i), we have by (xi) that $\tau_{h_m} \circ c_1 = \tau_{h_m} \circ c_2$ implies $q_{m,m+1} \circ s \circ c_1(v) = q_{m,m+1} \circ s \circ c_2(v)$, so $s \circ c_1(v)$, $s \circ c_2(v)$ are in the same orbit of B_m/B_{m+1} for every v. In particular we have $\rho_s(c_1) - \rho_s(c_2) = \sigma_{k+1}(s \circ c_1 - s \circ c_2)$, whence (2.18) follows.

To complete the proof we shall now use ρ_s to find a new cross section $s': Y \to X/B_{m+1}$ that is consistent with the factor Y_t , and that is also congruent to s modulo B_m , that is $q_{m,m+1} \circ s' = q_{m,m+1} \circ s$. If we find s' then, firstly, thanks to the consistency, we can apply Lemma 2.7.6 with map $\psi = \tau_t$, obtaining a new fibre-surjective factor X_{m+1} of X/B_{m+1} that is an extension of Y_t by Z_k/B_{m+1} , and so (i) holds. Secondly $\psi_m \circ q_{m,m+1}$ factors through X_{m+1} , which yields ψ'_{m+1} satisfying (ii). Indeed, by definition of the relation \sim in Lemma 2.7.6 we have that $x, y \in X/B_{m+1}$ are in the same fibre of ψ_{m+1} if and only if $\tau_t \circ \kappa_{m+1}(x) = \tau_t \circ \kappa_{m+1}(y)$ and $x - s' \circ \kappa_{m+1}(x) = y - s' \circ \kappa_{m+1}(y)$. The first equality here implies, thanks to the congruence mod B_m and (xi), that $\psi_m \circ q_{m,m+1}(s' \circ \kappa_{m+1}(x)) = \psi_m \circ q_{m,m+1}(s' \circ \kappa_{m+1}(y))$, and this together with the second equality gives $\psi_m \circ q_{m,m+1}(x) = \psi_m \circ q_{m,m+1}(y)$. Finally, property (iii) holds, since $x \sim y \Rightarrow \tau_t \circ \kappa_{m+1}(x) = \tau_t \circ \kappa_{m+1}(y)$ implies that $\tau_t \circ \kappa_{m+1}$ factors through ψ_{m+1} , and it must then factor as π_{m+1} , by definition of κ_{m+1} . We will have thus found the next term in the sequence S.

Averaging ρ_s to obtain a fibre-surjective factor X_{m+1} of X/B_{m+1} : to find s', the idea is firstly to define a new function ρ' by averaging ρ_s , in such a way that the approximate equality in (2.17) becomes an equality for ρ' , and secondly then to show that ρ' is itself a cocycle generated by some cross section, which will be the desired s'.

Recall that, by [8, Lemma 3.3.12 (ii) & (iii)], for each $i \in \mathbb{N}$ the map $\beta_i : C^i(Y) \to C^i(Y_t)$, $c \mapsto \tau_t \circ c$ is a totally surjective bundle morphism, and for every $c \in C^i(Y_t)$ the preimage $\beta_i^{-1}(c)$ is a (k-1)-fold sub-bundle of $C^i(Y)$. (Recall also Lemmas 2.2.16 and 2.2.17.) We define $\rho' : C^{k+1}(Y) \to Z_k/B_{m+1}$ by

$$\rho'(c) = \mathbb{E}_{c' \in \beta_{k+1}^{-1}(\beta_{k+1}(c))} \rho_s(c'). \tag{2.19}$$

The averaging here is relative to the Haar measure on $\beta_{k+1}^{-1}(\beta_{k+1}(c))$, and is well-defined for ε_2 sufficiently small, thanks to (2.17) (recall Definition 2.5.5). By (2.18) and the fact that $t > h_m$, we have that $\rho'(c) - \rho_s(c) \in B_m/B_{m+1}$ for every $c \in C^{k+1}(Y)$. Note also that it follows clearly from (2.19) that

$$\forall c_1, c_2 \in C^{k+1}(Y) \text{ with } \tau_t \circ c_1 = \tau_t \circ c_2, \text{ we have } \rho'(c_1) = \rho'(c_2). \tag{2.20}$$

We shall now prove that ρ' is a cocycle by showing that it inherits the required properties from ρ_s thanks to the linearity of the averaging in (2.19). The required properties are the two axioms from [8, Definition 3.3.14]. Axiom (i) follows straight from the same property for ρ_s and linearity using the fact that any set $\beta_{k+1}^{-1}(\beta_{k+1}(c))$ is globally invariant under composition with automorphisms in $\operatorname{Aut}(\{0,1\}^{k+1})$. To see the second axiom we take two adjacent cubes $c_1, c_2 \in C^{k+1}(Y)$ with concatenation c_3 and embed them into a (k+2)-dimensional cube $c \in C^{k+2}(Y)$ as restrictions to two adjacent (k+1)-dimensional faces F_1 and F_2 in $\{0,1\}^{k+2}$. The concatenation of F_1 and F_2 is a diagonal subcube F_3 . We have that the concatenation of c_1 and c_2 is the restriction of c to F_3 . Note that the existence of c is guaranteed by simplicial completion [8, Lemma 3.1.5]. By [8, Lemma 3.3.12 (iii)], we have a Haar probability on $\Omega = \beta_{k+2}^{-1}(\beta_{k+2}(c))$. Then, by [8, Lemma 3.3.12 (iv)], the probability spaces $\beta_{k+1}^{-1}(\beta_{k+1}(c_i))$, i = 1, 2, 3 are faithfully embedded as factors and coupled in the big probability space Ω . Using the concatenation property for ρ in Ω (when the random cube is restricted to F_1, F_2, F_3) and linearity of expectation we obtain that $\rho'(c_3) = \rho'(c_1) + \rho'(c_2)$.

It remains to prove that the cocycle ρ' is generated by a cross section. Let ρ'' denote the cocycle $\rho' - \rho_s$, which takes values in B_m/B_{m+1} (as noted just before (2.20)). It follows from (2.17) that $d_2(\rho''(c),0) \le \varepsilon_2$ for every $c \in C^{k+1}(Y)$, so for ε_2 sufficiently small we have that ρ'' is a coboundary, by Lemma 2.5.7. Thus for some Borel function $g: Y \to B_m/B_{m+1}$ we have $\rho''(c) = \sigma_{k+1}(g \circ c)$. Hence, defining the Borel cross section $s': Y \to X/B_{m+1}$ by s'(y) = s(y) + g(y), we have $\rho' = \rho_{s'}$.

By (2.20) we have that s' is consistent with the factor Y_t . Applying Lemma 2.7.6 as described above, we find the new term in the sequence S.

This completes the proof of Theorem 2.7.3.

2.8 Rigidity of morphisms

Given a set X, a metric space (Y,d), and maps $\phi, \phi': X \to Y$, we say that ϕ' is an ε -modification of ϕ if for every $x \in X$ we have $d(\phi(x), \phi'(x)) \le \varepsilon$.

Definition 2.8.1. Let X be a compact nilspace, and let Y be a k-step compact nilspace with a metric d generating its topology. A map $\phi: X \to Y$ is a δ -quasimorphism if for every $c \in C^{k+1}(X)$ there exists $c' \in C^{k+1}(Y)$ such that $d(\phi \circ c(v), c'(v)) \le \delta$ for every $v \in \{0, 1\}^{k+1}$.

In this section we establish the following result.

Theorem 2.8.2. Let Y be a k-step compact nilspace of finite rank, with a metric d generating its topology. Then for every $\varepsilon > 0$ there exists $\delta > 0$ such that the following holds. If X is a compact nilspace and $\phi : X \to Y$ is a Borel δ -quasimorphism, then ϕ has an ε -modification that is a continuous morphism.

In the proof we shall use the fact that the metric d can be assumed to be invariant under the action of the structure group Z_k (Lemma 2.1.11). We shall also use the following 'rectification result' for cubes.

Lemma 2.8.3. Let X be a k-step compact nilspace with metric d. For every $\varepsilon > 0$ there exists $\delta > 0$ such that the following holds. If $c \in C^{k+1}(X)$ satisfies $d'(\pi_{k-1} \circ c(v,0), \pi_{k-1} \circ c(v,1)) \leq \delta$ for every $v \in \{0,1\}^k$, then c has an ε -modification $c' \in C^{k+1}(X)$ such that $\pi_{k-1} \circ c'(v,0) = \pi_{k-1} \circ c'(v,1)$ for every $v \in \{0,1\}^k$.

Proof. Suppose for a contradiction that for some $\varepsilon > 0$ there is a sequence of cubes $c_n \in C^{k+1}(X)$ such that $d'(\pi_{k-1} \circ c_n(v,0), \pi_{k-1} \circ c_n(v,1)) \leq 1/n$ for every $v \in \{0,1\}^k$ and yet, for every cube $c' \in C^{k+1}(X)$ such that $\pi_{k-1} \circ c'(v,0) = \pi_{k-1} \circ c'(v,1)$ for all $v \in \{0,1\}^k$, there is some $v \in \{0,1\}^{k+1}$ such that $d(c_n(v),c'(v)) > \varepsilon$. By compactness of $C^{k+1}(X)$, there is a subsequence (c_m) of (c_n) such that $c_m \to c^* \in C^{k+1}(X)$. The topology on $C^{k+1}(X)$ can be metrized by $d_\infty(c_1,c_2) := \max_{v \in \{0,1\}^{k+1}} d(c_1(v),c_2(v))$. It then follows from $d_\infty(c_m,c^*) \to 0$ and the triangle inequality for d' that $\pi_{k-1} \circ c^*(v,0) = \pi_{k-1} \circ c^*(v,1)$ for every $v \in \{0,1\}^k$. However, it also follows from $c_m \to c^*$ that $d(c_m(v),c^*(v)) \le \varepsilon$ for all $v \in \{0,1\}^{k+1}$, for sufficiently large m, and this contradicts our initial assumption since $\pi_{k-1} \circ c^*(\cdot,0) = \pi_{k-1} \circ c^*(\cdot,1)$. \square

Proof of Theorem 2.8.2. We argue by induction on k, starting from the trivial case k=0. For k>0, suppose that the result holds for k-1, and fix any $\varepsilon>0$. Let $\delta>0$ be a parameter to be determined later, suppose that $\phi: X \to Y$ is a δ -quasimorphism, and let $\phi_1 = \pi_{k-1} \circ \phi$. Then ϕ_1 is a δ -quasimorphism into $\mathcal{F}_{k-1}(Y)$. By induction, for some $\delta_1(\varepsilon)>0$ to be fixed later, we have that if δ is small enough then ϕ_1 has a δ_1 -modification $\phi_2: X \to \mathcal{F}_{k-1}(Y)$ that is a continuous morphism.

We shall now obtain a Borel measurable map $\phi_3: X \to Y$ that lifts ϕ_2 , i.e. such that $\pi_{k-1} \circ \phi_3 = \phi_2$, and which is a $(\delta + \delta_1)$ -modification of ϕ , so that ϕ_3 is a δ_2 -quasimorphism with $\delta_2 = \delta_1 + 2\delta$. To obtain this, the idea is first to obtain such a lift locally for each neighbourhood in some appropriate cover of X, and then combine the local lifts into a global one. To this end we use the following set:

$$G = \{(x, y) : \pi_{k-1}(y) = \phi_2(x)\} \subset X \times Y.$$

We claim that G is compact. Indeed, on one hand the map $f_1 \times f_2 : X \times Y \to \mathcal{F}_{k-1}(Y) \times \mathcal{F}_{k-1}(Y)$, $(x,y) \mapsto (\phi_2(x), \pi_{k-1}(y))$ is continuous, and on the other hand $(x,y) \in G$ if and only if $f_1 \times f_2(x,y)$ is in the diagonal of $\mathcal{F}_{k-1}(Y) \times \mathcal{F}_{k-1}(Y)$, which is a closed set in the product topology since $\mathcal{F}_{k-1}(Y)$ is Hausdorff. Hence G is closed in $X \times Y$ and therefore compact. We also claim that G is a Z_k -bundle over X, where Z_k is the k-th structure group of Y. Indeed, the action of Z_k is given by $a \cdot (x,y) = (x,y+a)$, and the projection map is just $(x,y) \mapsto x \in X$.

By Proposition 2.5.2, the bundle G is locally trivial and so for every $p \in X$ we can find an open set $U_p \subset X$ containing p with a continuous cross section $U_p \to G$, which must then be of the form $x \mapsto (x, \tau(x))$ for some continuous map $\tau : U_p \to Y$. Thus $\phi_2(x) = \pi_{k-1}(\tau(x))$ for every $x \in U_p$. Then, by the definition of d', since for all $x \in U_p$ we have

$$d'(\pi_{k-1}(\tau(x)), \pi_{k-1}(\phi(x))) = d'(\phi_2(x), \phi_1(x)) \le \delta_1,$$

in particular for p there exists $z_p \in \mathbb{Z}_k$ such that $d(\tau(p) + z_p, \phi(p)) \leq \delta_1$. But then, since the function $f_p : x \mapsto \tau(x) + z_p$ is continuous on U_p , there is an open set $U_p' \subset U_p$ containing p such that for all $x \in U_p'$ we have $d(\phi(x), \tau(x) + z_p) < \delta_1 + \delta$. The function f_p on U_p' is the desired local lift of ϕ_2 .

Now by compactness there is a finite covering of X by such open sets U'_p . Let U_1, \ldots, U_s be the atoms of the Boolean algebra generated by the sets U'_p . For each $i \in [s]$ we fix one of the functions f_p defined on U_s and rename it f_i . We then define the Borel function ϕ_3 by $\phi_3(x) = \sum_{i \in [s]} 1_{U_i}(x) f_i(x)$.

We shall now apply an averaging argument to ϕ_3 that will produce a continuous morphism ϕ_4 which is close to ϕ as desired.

Let $P_2 = \{0,1\}^{k+1} \setminus \{0^{k+1}\}$. Note that since ϕ_2 is a morphism $X \to \mathcal{F}_{k-1}(Y)$, it preserves k-cubes, and since ϕ_3 is a lift of ϕ_2 , by [8, Remark 3.2.12] we have that ϕ_3 also preserves k-cubes on Y. This implies that for every $c \in C^{k+1}(X)$ the restriction of $\phi_3 \circ c$ to P_2 is a (k+1)-corner on Y. For any given cube $c \in C^{k+1}(X)$ let us denote by $K_0(c)$ the value at 0^{k+1} of the unique completion in $C^{k+1}(Y)$ of $(\phi_3 \circ c)|_{P_2}$. Let

$$\phi_4(x) = \mathbb{E}_{c \in C_x^{k+1}(X)} K_0(c). \tag{2.21}$$

Let us show that this averaging is well-defined. Note that for every $c \in C_x^{k+1}(X)$ we have $K_0(c) \sim_{k-1} \phi_3(x)$, indeed since the morphism ϕ_2 preserves (k+1)-cubes, and ϕ_3 is a lift of ϕ_2 , we must have $K_0(c)$ and $\phi_3(x)$ both in the fibre $\pi_{k-1}^{-1}(\phi_2(x))$. We have to show that if $\delta_2 > 0$ is small enough then the set of values $K_0(c)$, $c \in C_x^{k+1}(X)$ has small diameter in the fibre. We will show that $K_0(c)$ is actually close to $\phi_3(x)$ for every $c \in C_x^{k+1}(X)$, which will also be useful later. For any $c \in C_x^{k+1}(X)$, since ϕ_3 is a δ_2 -quasimorphism, there is $c_0 \in C^{k+1}(Y)$ such that for all v we have $d(\phi_3 \circ c(v), c_0(v)) \le \delta_2$. In particular, this inequality holds for all $v \in P_2$. By continuity of the corner-completion (Lemma 2.1.12) it follows that for $\lambda > 0$, if δ_2 is small enough then we have $d(K_0(c), c_0(0^{k+1})) \le \lambda$. Hence $d(\phi_3(x), K_0(c)) = d(\phi_3 \circ c(0^{k+1}), K_0(c)) \le \delta_2 + \lambda$. We can now fix λ so that the averaging in (2.21) is indeed well-defined.

Let us now prove that ϕ_4 is a morphism. By [8, Lemma 3.2.13], it suffices to show that for every $c \in C^{k+1}(X)$ we have $\phi_4 \circ c \in C^{k+1}(Y)$. We shall do this using another averaging argument, but working this time with the tricube T_{k+1} . Let $B = T_{k+1} \setminus \{-1,1\}^{k+1}$. For every $t \in \text{hom}_{c \circ \omega_{k+1}^{-1}}(T_{k+1},X)$, there is a unique completion of $\phi_3 \circ t|_B$ to a morphism $t': T_{k+1} \to Y$. Let $c_t = t' \circ \omega_{k+1}$, which is in $C^{k+1}(Y)$ by [8, Lemma 3.1.16]. If δ_2 is small enough then the following function is well defined:

$$c_2: \{0,1\}^{k+1} \to Y, \ \nu \mapsto \mathbb{E}_{t \in \hom_{c \circ \omega_{k+1}^{-1}}(T_{k+1}, X)} \ c_t(\nu).$$

Indeed, by Corollary 2.2.22, for each v this average is equal to $\mathbb{E}_{c' \in C_{c(v)}^{k+1}(X)} K_0(c') = \phi_4(c(v))$. Thus, we have to show that $c_2 \in C^{k+1}(X)$. To see this, recall that for every $c' \in C_{c(v)}^{k+1}(X)$ we have $K_0(c') \sim_{k-1} \phi_3(c(v))$, and so $c_2(v) \sim_{k-1} \phi_3 \circ c(v)$, which implies that $\pi_{k-1} \circ c_2 = \phi_2 \circ c \in C^{k+1}(\mathcal{F}_{k-1}(Y))$. Therefore, by [8, Theorem 3.2.19] and the definition of degree-k extensions, we just have to check that for some (any) t_0 we have $\sigma_{k+1}(c_2-c_{t_0})=0$. This follows from linearity of averaging (Lemma 2.5.6) and the fact that for c_t , c_{t_0} are cubes with the same projection to $\mathcal{F}_{k-1}(Y)$ for every t, so

$$\sigma_{k+1}(c_2-c_{t_0}) \ = \ \sum_{\nu} (-1)^{|\nu|} \ \mathbb{E}_{t \in hom_{c \circ \omega_{k+1}^{-1}}(T_{k+1},X)} \ c_t(\nu) - c_{t_0}(\nu) \ = \ \mathbb{E}_t \ \sigma_{k+1} \big(c_t - c_{t_0} \big) = 0.$$

Finally, let us prove that ϕ_4 is continuous. The argument is similar to the proof of Theorem 2.4.6. Consider the CSM structure on $\pi_0: C^{k+1}(X) \to X$, $c \mapsto c(0^{k+1})$, given by Lemma 2.2.17. For each $v \in \{0,1\}^{k+1} \setminus \{0^{k+1}\}$ let $\pi_v: C^{k+1}(X) \to X$, $c \mapsto c(v)$. We have that π_v is continuous and each of its restrictions to a space $C_x^{k+1}(X)$ is measure-preserving, by Lemma 2.2.14. For each v we let $f_v: X \to \mathcal{L}(C^{k+1}(X), Y)$ be the function sending x to the restriction of $\phi_3 \circ \pi_v$ to $C_x^{k+1}(X)$. Then, just as in the proof of Theorem 2.4.6, we have that f_v is continuous, by Lemma 2.3.14. We now want to express ϕ_4 in terms of the functions f_v . To that end, consider again the map $\xi: X \to \mathcal{L}(C^{k+1}(X), Cor^{k+1}(Y))$, sending x to the function $\xi_x: c \in C_x^{k+1}(X) \mapsto (c(v) = f_v(x)(c))_{v \neq 0^{k+1}}$. Just as in the proof of Theorem 2.4.6, the continuity of each function f_v implies that ξ is continuous. We then use again the continuous map $\mathcal{K}: \mathcal{L}(C^{k+1}(X), Cor^{k+1}(Y)) \to \mathcal{L}(C^{k+1}(X), Y)$ that sends a function g to $g \circ K$, where K is the unique completion of (k+1)-corners on Y. Now note that $\mathcal{K} \circ \xi: X \to \mathcal{L}(C^{k+1}(X), Y)$ is precisely the map that sends each $x \in X$ to the function $c \mapsto K_0(c)$ on $C_x^{k+1}(X)$. This map is continuous, and by the discussion justifying the averaging (2.21), we have that in fact $\mathcal{K} \circ \xi$ takes values in the following subset of $\mathcal{L}(C^{k+1}(X), Y)$:

$$\mathcal{U} = \bigcup_{x \in X} \Big\{ g \in L\big(C_x^{k+1}(X), \pi_{k-1}^{-1}(\phi_2(x))\big) \; : \; \forall c \in C_x^{k+1}(X), \; d\big(g(c), \phi_3(x)\big) \leq \delta_2 + \lambda \Big\}.$$

Now ϕ_4 is the composition of $\mathcal{K} \circ \xi$ with the averaging operator $\mathcal{A}: \mathcal{U} \to \pi_{k-1}^{-1}(\phi_2(x))$ that sends a function g to $\mathbb{E}_{c \in C_x^{k+1}(X)}g(c)$. Thus it now suffices to show that \mathcal{A} is continuous. Note that since $\pi_{k-1}^{-1}(\phi_2(x))$ is homeomorphic to Z_k , it follows from the definition of the metric d_2 on Z_k that if $\delta_2 + \lambda$ is sufficiently small then every function $g \in \mathcal{U}$ can be lifted continuously to a function $g': C_x^{k+1}(X) \to \mathbb{R}^n$, where $Z_k = F \times \mathbb{T}^n$, in such a way that two functions $g_1, g_2 \in \mathcal{U}$ are close in the metric d_1 from (2.8) if and only if their lifts g'_1, g'_2 are close in the same metric but with d now being the Euclidean distance on \mathbb{R}^n . By Lemma 2.3.7, if $g_n \to g$ in the topology on $\mathcal{L}(C^{k+1}(X), Y)$ restricted to $\mathcal{L}(C_x^{k+1}(X), Y)$, then $d_1(g_n, g) \to 0$, whence by the triangle inequality $\mathcal{A}(g_n) \to \mathcal{A}(g)$, so \mathcal{A} is continuous as required.

2.9 Characterizing compact connected nilspaces of finite rank

Recall from [8, Subsection 3.2.4] that given a k-step nilspace X we denote by $\Theta_i(X)$ the group of translations of height i on X (or i-translations). By a slight abuse of notation, when X is a compact nilspace we shall write $\Theta_i(X)$ for the group of i-translations that are also *continuous* functions. A central goal in this section is to show that if X has finite rank then $\Theta_i(X)$ is a Lie group for each i and $\Theta_1(X)$ acts transitively on the connected components of X. This will then enable us to show that if X has connected structure groups then it can be identified with a filtered nilmanifold (Theorem 2.9.17).

Throughout this section we shall abbreviate 'compact and finite-rank' by writing 'CFR'.

Lemma 2.9.1. Let X be a CFR k-step nilspace. Then every element of $\Theta(X)$ is a homeomorphism $X \to X$ preserving the Haar measure.

Proof. Each translation is an invertible continuous map from the compact Hausdorff space X to itself and so its inverse is also continuous [32, Theorem 26.6]. Moreover, the translation is a fibre-surjective automorphism of X, so it is measure preserving by Corollary 2.2.7. \Box

For a compact space X and a space Y with metric d, recall that C(X,Y) denotes the space of continuous functions $f: X \to Y$, with the topology induced by the uniform metric $d_{\infty}(f_1, f_2) = \sup_{x \in X} d(f_1(x), f_2(x))$. We record the following basic fact.

Lemma 2.9.2. Let X be a CFR k-step nilspace and let d be a metric generating the topology on X. Then for every $i \in [k]$, the group $\Theta_i(X)$ equipped with the restriction of the uniform metric on C(X,X) is a Polish group.

Proof. The group H(X) of homeomorphisms on X, with the relative topology from C(X,X), is a Polish group [3, §1.3, Example (v)]. It follows from the definition of translations [8, Definition 3.2.27] and the closure of cube sets that $\Theta_i(X)$ is a closed subgroup of H(X), hence it is also a Polish group. By definition the metric d_{∞} generates the topology on $\Theta_i(X)$.

Recall that by [8, Lemma 3.2.31] a translation maps every class of the relation \sim_{k-1} onto another such class. This enables us to define a translation on the factor $X_{k-1} = \mathcal{F}_{k-1}(X)$, as in the following lemma. Recall also that the topology on X_{k-1} is the quotient topology from X, which can be metrized by the quotient metric d' defined in (2.2). We then denote by d'_{∞} the uniform metric on $\Theta(X_{k-1})$ relative to d'.

Lemma 2.9.3. For each $i \in \mathbb{N}$ let h be the map sending each $\alpha \in \Theta_i(X)$ to the map $h(\alpha)$ on X_{k-1} defined by $h(\alpha)(y) = \pi_{k-1}(\alpha(x))$, for any $x \in X$ such that $\pi_{k-1}(x) = y$. Then h is a continuous homomorphism $\Theta_i(X) \to \Theta_i(X_{k-1})$.

Proof. That h is a homomorphism $\Theta_i(X) \to \Theta_i(X_{k-1})$ follows from the definitions. The continuity follows similarly, thus

$$d'_{\infty}(h(\alpha_{1}),h(\alpha_{2})) = \sup_{y \in X_{k-1}} d'(h(\alpha_{1})(y),h(\alpha_{2})(y)) = \sup_{x \in X} d'(\pi_{k-1}(\alpha_{1}(x)),\pi_{k-1}(\alpha_{2}(x)))$$

$$= \sup_{x \in X} \inf_{\substack{x_{1} \sim_{k-1}\alpha_{1}(x) \\ x_{2} \sim_{k-1}\alpha_{2}(x)}} d(x_{1},x_{2}) \leq \sup_{x \in X} d(\alpha_{1}(x),\alpha_{2}(x)) = d_{\infty}(\alpha_{1},\alpha_{2}).$$

An important result toward our goal in this section is that if an i-translation on X_{k-1} is sufficiently close to the identity then it can be lifted to an i-translation on X, in the following sense.

Lemma 2.9.4. Let X be a CFR k-step nilspace and let $i \in \mathbb{N}$. There exists $\varepsilon > 0$ such that if $\alpha \in \Theta_i(X_{k-1})$ satisfies $d_{\infty}'(\alpha, \mathrm{id}) < \varepsilon$ then there exists $\beta \in \Theta_i(X)$ such that $h(\beta) = \alpha$.

Proof. We shall use notation from [8, Lemma 3.3.38]. Recall that by that lemma the translation bundle $\mathfrak{T}^*(\alpha,X,i)=\mathfrak{F}_{k-1}(\mathfrak{T}(\alpha,X,i))$ is, from the purely algebraic viewpoint, a degree-(k-i) extension of X_{k-1} by Z_k , and note that by the results from Section 2.1 we have that this extension is a continuous Z_k -bundle. We show first that if ε is sufficiently small then this extension has a measurable cross section s such that the corresponding cocycle ρ_s is a coboundary.

Let γ be the projection $\mathfrak{T}^*(\alpha,X,i)\to X_{k-1}$. We claim that if $\varepsilon>0$ is sufficiently small, then we can choose a Borel cross section $s:X_{k-1}\to\mathfrak{T}^*$ with the property that for each s(x), for every (x_0,x_1) in the equivalence class $\pi_{k-1,\mathfrak{T}}^{-1}(s(x))\subset\mathfrak{T}$, we have $d(x_0,x_1)\leq\varepsilon$. (Recall that we have $\alpha(\pi_{k-1}(x_0))=\pi_{k-1}(x_1)$ for $(x_0,x_1)\in\mathfrak{T}$, by [8, Definition 3.3.34].) Indeed, by definition $d'_\infty(\alpha,\mathrm{id})<\varepsilon$ implies that, for each $x\in X_{k-1}$, there exist $x_0,x_1\in X$ with $\pi_{k-1}(x_0)=x$, $\pi_{k-1}(x_1)=\alpha(x)$, and such that $d(x_0,x_1)<\varepsilon$. Then, by Z_k -invariance of d, we have that every pair (y_0,y_1) in the same equivalence class as (x_0,x_1) in \mathfrak{T}^* satisfies $d(y_0,y_1)\leq\varepsilon$. We choose s(x) to be one of these pairs, say (x_0,x_1) . Note that this choice can be made in a Borel measurable way. Indeed, using the fact that X is a locally trivial Z_k -bundle over X_{k-1} (Proposition 2.5.2), we can construct s as a piecewise continuous function.

Now let ρ_s be the cocycle generated by s, defined for $c \in C^{k-i+1}(X_{k-1})$ by $\rho_s(c) = \sigma_{k-i+1}(s \circ c - c')$ for any lift $c' \in C^{k-i+1}(\mathfrak{I}^*)$ of c. We have that ρ_s is measurable (recall the end of the proof of Lemma 2.4.5). We claim that, for some $\varepsilon_2 > 0$ to be fixed later, if ε is sufficiently small then ρ_s is also small in the sense that $d_2(\rho_s(c), 0_{Z_k}) \le \varepsilon_2$ for every $c \in C^{k-i+1}(X_{k-1})$. To show this we first give an alternative expression of the function $s \circ c - c'$. The lift c', being a cube on \mathfrak{T}^* , has itself a lift $\tilde{c} \in C^{k-i+1}(\mathfrak{T})$. By definition we have $\tilde{c} = \tilde{c}_0 \times \tilde{c}_1$ where $\pi_{k-1} \circ \tilde{c}_1(v) = \alpha(\pi_{k-1} \circ \tilde{c}_0(v))$ for all v. By the claim in the previous paragraph, for each v we have that $s \circ c(v)$ is a class of pairs $(x_0, x_1) \in \pi_{k-1}^{-1}(c(v)) \times \pi_{k-1}^{-1}(\alpha \circ c(v))$ satisfying $d(x_0, x_1) \le \varepsilon$ and corresponding to a local translation ϕ , thus $x_1 = \phi(x_0)$ for every such pair. The difference $s \circ c(v) - c'(v)$ is then the element $a_v \in \mathbb{Z}_k$ which has to be added to $\tilde{c}_0(v)$ in order to have $\phi(\tilde{c}_0(v) + a_v) = \tilde{c}_1(v)$, namely $a_v = \phi^{-1}(\tilde{c}_1(v)) - \tilde{c}_0(v)$. Thus, letting a denote the function $\{0,1\}^{k-i+1} \to Z_k, v \mapsto a_v$, we have $\rho_{\rm s}({\rm s}\circ{\rm c}-{\rm c}')=\sigma_{k-i+1}(a)$. To show that this must be a small element of ${\rm Z}_k$, we argue as follows. By assumption on α , the cube \tilde{c} satisfies the premise of Lemma 2.8.3 with $\delta = \varepsilon$. Thus, given $\varepsilon_1 > 0$ to be fixed later, if $\varepsilon < \varepsilon_1$ is sufficiently small then by that lemma there must exist $c^* = c_0^* \times c_1^* \in C^{k-i+1}(\mathfrak{T})$ such that for all v we have $d(c_i^*(v), \tilde{c}_i(v)) \le \varepsilon_1$ for i = 0, 1, and with $\pi_{k-1} \circ c_1^* = \pi_{k-1} \circ c_0^*$. For each v let $a^*(v) = c_1^*(v) - c_0^*(v) \in \mathbb{Z}_k$. It now suffices to show that $d_2(a^*(v), a(v)) \le \varepsilon_2/2^{k-i+1}$ for every v, for then, by virtue of $\langle c_0^*, c_1^* \rangle_1$ being a cube, we must have $\sigma_{k-i+1}(a^*) = 0$ and so $\sigma_{k-i+1}(a) \leq \varepsilon_2$ as claimed. To show that $d_2(a^*(v), a(v))$ is small we can apply Lemma 2.7.7 with $\delta = 2\varepsilon_1$, since we know on one hand that $d(c_0^*(v), \tilde{c}_0(v)) \leq \varepsilon_1$, and on the other hand that $d(\phi^{-1}(\tilde{c}_1(v)), \tilde{c}_1(v)) \leq \varepsilon$ and $d(c_1^*(v), \tilde{c}_1(v)) \leq \varepsilon_1$, so that $d(\phi^{-1}(\tilde{c}_1(v)), c_1^*(v)) \leq \varepsilon + \varepsilon_1 \leq 2\varepsilon_1$.

We can now apply Lemma 2.5.7. Thus if ε_2 is sufficiently small then there exists a Borel function $g: \mathcal{F}_{k-1}(X) \to Z_k$ such that $\rho_s(c) = \sigma_{k-1+i}(g \circ c)$ for all $c \in C^{k-i+1}(\mathcal{F}_{k-1}(X))$. Let $m: \mathcal{F}_{k-1}(X) \to \mathcal{T}^*$ be the measurable function $x \mapsto s(x) - g(x)$. We have that m is a morphism, indeed for any $c \in C^n(\mathcal{F}_{k-1}(X))$, on one hand we clearly have that the \mathcal{T}^* -valued map $m \circ c$ is a lift of c, and on the other hand for some (any)

P. CANDELA

lift $c' \in C^n(\mathfrak{T}^*)$ of c we have $\sigma_{k-i+1}(m \circ c - c') = \sigma_{k-i+1}(s \circ c - c') - \sigma_{k-i+1}(g \circ c) = \rho_s(c) - \rho_s(c) = 0$, so by [8, Definition 3.3.13 (ii)] we have indeed $m \circ c \in C^n(\mathfrak{T}^*)$. Now [8, Proposition 3.3.39] gives us an i-translation β on X that is a lift of α , and is defined as follows: for each $x \in X$ we have $\beta(x) = \phi(x)$, where ϕ is the local translation corresponding to the equivalence class $m \circ \pi_{k-1}(x) \in \mathfrak{T}^*$. It only remains to check that β is Borel measurable, as then, being a morphism, it must be continuous by Theorem 2.4.6. By [8, Definition 3.3.34 and Proposition 3.3.36], the nilspace \mathfrak{T} is a continuous Z-bundle over \mathfrak{T}^* , for the polish group $Z = \{(z,z) : z \in Z_k\} \leq Z_k \times Z_k$. Therefore, arguing as in the end of the proof of Lemma 2.4.5, we obtain a Borel cross section $\lambda : \mathfrak{T}^* \to \mathfrak{T}$, and so the map $\lambda \circ m \circ \pi_{k-1} : X \to \mathfrak{T}$ is Borel. Moreover, letting $\lambda \circ m \circ \pi_{k-1}(x) = (y_0, y_1)$, we have $\beta(x) = y_1 + (x - y_0)$, so β is indeed Borel. \square

We shall need a useful criterion for a translation to lie in the kernel of h. Note that if $\alpha \in \ker(h)$ then, since by [8, Lemma 3.2.31] the restriction of α to each fibre of π_{k-1} is a local translation, we deduce that the map $x \mapsto \alpha(x) - x$ is a constant element of Z_k for every x in a given fibre. Thus α induces a well-defined map $\alpha' : X_{k-1} \to Z_k$, $\pi_{k-1}(x) \mapsto \alpha(x) - x$. Conversely, given $\alpha' : X_{k-1} \to Z_k$ we can define $\alpha : X \to X$, $x \mapsto x + \alpha'(\pi_{k-1}(x))$, and we can ask when is α in $\Theta(X)$.

Lemma 2.9.5. Let i < k. Then $\alpha \in \ker(h) \cap \Theta_i(X)$ if and only if $\alpha' \in \hom(X_{k-1}, \mathcal{D}_{k-i}(Z_k))$.

Proof. By [8, Lemma 3.2.32] we know that $\alpha \in \Theta_i(X)$ if and only if for every $c \in C^n(X)$ we have $\langle c, \alpha \circ c \rangle_i \in C^{n+i}(X)$. If $\alpha \in \ker(h) \cap \Theta_i(X)$ then we can take the difference $\alpha \circ c - c = \alpha' \circ c'$ for $c' = \pi_{k-1} \circ c$, and by definition of degree-k extensions this difference must be a cube in $\mathcal{D}_k(Z_k)$. But we also have $\langle c, \alpha \circ c \rangle_i - \langle c, c \rangle_i \in C^{n+i}(\mathcal{D}_k(Z_k))$, and this equals $\langle 0, \alpha' \circ c' \rangle_i$. By [8, Lemma 3.3.37], we have $\langle 0, \alpha' \circ c' \rangle_i \in C^{n+i}(\mathcal{D}_k(Z_k))$ if and only if $\alpha' \circ c' \in C^n(\mathcal{D}_{k-i}(Z_k))$. The converse follows similarly.

The next main tool that we need, Lemma 2.9.7 below, is a type of rigidity result for morphisms between abelian torsors of the form $\mathcal{D}_k(\mathbf{Z})$. The proof will use the following basic fact.

Lemma 2.9.6. Let Z' be a CFR abelian group, and let d_2 be the metric from Definition 2.5.4. Then there exists $\eta > 0$ such that for every compact abelian group Z and every non-constant continuous affine homomorphism $\phi : Z \to Z'$, there exist $x, y \in Z$ such that $d_2(\phi(x), \phi(y)) \ge \eta$.

Proof. We can assume without loss that ϕ is a homomorphism (not just an affine one). If Z' is a finite group then the claim is verified by any two points x, y such that $\phi(x) \neq \phi(y)$. We may therefore suppose that $Z' = F \times \mathbb{T}^n$ with n > 0 and F a finite abelian group, and that for some $i \in [n]$ the projection $\pi_i : \mathbb{T}^n \to \mathbb{T}, x \mapsto x_i$ satisfies that $\pi_i \circ \phi(Z)$ is a nontrivial subgroup of \mathbb{T} . Identifying \mathbb{T}^n with $[0,1)^n$, there is then $y \in Z$ such that $|\pi_i \circ \phi(y) - 0|_{\mathbb{T}} \geq 1/4$ and so with $x = 0_Z$ we have $d_2(\phi(x), \phi(y)) \geq 1/4$.

Lemma 2.9.7. Let $k, \ell \in \mathbb{N}$, let Z, Z' be compact abelian groups, and suppose that Z' has finite rank. Then there exists $\varepsilon = \varepsilon(\ell, Z') > 0$ such that if $\phi \in \text{hom}(\mathbb{D}_k(Z), \mathbb{D}_\ell(Z'))$ satisfies $d_2(\phi(x), \phi(y)) \leq \varepsilon$ for every $x, y \in Z$ then ϕ is a constant function.

Proof. Since hom $(\mathcal{D}_k(\mathbf{Z}), \mathcal{D}_\ell(\mathbf{Z}')) \subseteq \text{hom}(\mathcal{D}_1(\mathbf{Z}), \mathcal{D}_\ell(\mathbf{Z}'))$ we can assume that k = 1. Let ϕ be an arbitrary non-constant morphism from $\mathcal{D}_1(\mathbf{Z})$ to $\mathcal{D}_\ell(\mathbf{Z}')$. As mentioned in [8, Example 2.2.13] (recall also [8, Definition 2.2.30]), the map ϕ is a polynomial map of degree at most ℓ . Thus, if for every $t \in \mathbf{Z}$ and every $f: \mathbf{Z} \to \mathbf{Z}'$ we denote by $\Delta_t f$ the function $x \mapsto f(x+t) - f(x)$, then there is $i < \ell$ and elements

 $t_1, t_2, \ldots, t_i \in \mathbb{Z}$ such that $\phi' = \Delta_{t_1} \circ \Delta_{t_2} \circ \cdots \circ \Delta_{t_i} \phi$ is non-constant but $\Delta_t \phi'$ is constant for every $t \in \mathbb{Z}$. It follows that ϕ' is a non-constant continuous affine homomorphism from \mathbb{Z} to \mathbb{Z}' , so by Lemma 2.9.6 there are $x, y \in \mathbb{Z}$ with $d_2(\phi'(x), \phi'(y)) \geq \eta(\mathbb{Z}') > 0$. This cannot hold if $\max_{x,y} d_2(\phi(x), \phi(y))$ is too small, so this quantity must have a positive lower bound (depending only on \mathbb{Z}' and ℓ).

Lemma 2.9.7 has the following consequence.

Corollary 2.9.8. Let $\ell \in \mathbb{N}$, let X be a k-step compact nilspace, and let Z be a CFR abelian group. Then there exists $\varepsilon = \varepsilon(\ell, \mathbb{Z}) > 0$ such that if $\phi \in \text{hom}(X, \mathcal{D}_{\ell}(\mathbb{Z}))$ satisfies $d_2(\phi(x), \phi(y)) \leq \varepsilon$ for every $x, y \in X$ then ϕ is constant. In particular, for any $x_0 \in X$, $z_0 \in Z$, the set $Y = \{\phi \in \text{hom}(X, \mathcal{D}_{\ell}(\mathbb{Z})) : \phi(x_0) = z_0\}$, equipped with the metric $d_{\infty}(\phi, \phi') = \max_{x \in X} d_2(\phi(x), \phi'(x))$, is discrete.

Proof. Suppose that $d_2(\phi(x), \phi(y)) < \varepsilon$ for every $x, y \in X$, where $\varepsilon = \varepsilon(\ell, Z)$ is the constant from Lemma 2.9.7. We prove by induction on k that ϕ is constant.

If k=1 then X is abelian and the result follows from Lemma 2.9.7. Suppose that the statement holds for k-1. Then Lemma 2.9.7 tells us that ϕ is constant on the \sim_{k-1} classes of X. This means that ϕ can be regarded as a function on $\mathcal{F}_{k-1}(X)$, and so by our assumption ϕ is constant.

To see the last sentence in the corollary, note that if $\phi, \phi' \in Y$ satisfy $d_{\infty}(\phi, \phi') = d_{\infty}(\phi - \phi', 0) \le \varepsilon/2$, then by the triangle inequality $\phi - \phi'$ satisfies the premise of the previous sentence in the corollary, and then by that sentence and the definition of Y we deduce that $\phi - \phi'$ must be the constant 0.

With these results we can now give a useful description of the kernel of h.

Lemma 2.9.9. Let X be a CFR k-step nilspace, and let $h: \Theta(X) \to \Theta(X_{k-1})$ be the homomorphism from Lemma 2.9.3. Then $\ker(h)$ is isomorphic as a topological group to $F \times Z_k$ for some discrete group F. In particular $\ker(h)$ is a Lie group.

Proof. Let $x \in X$ be arbitrary and let F be the stabilizer $\operatorname{Stab}_{\ker(h)}(x) = \{\alpha \in \ker(h) : \alpha(x) = x\}$. Let $\tau : Z_k \to \ker(h)$ be the map sending $z \in Z_k$ to the translation $\tau_z : x \mapsto x + z$, and note that τ is an isomorphism of topological groups between Z_k and $\tau(Z_k) \leq \ker(h)$. Now given any $\alpha \in \ker(h)$, letting $z = \alpha(x) - x$ we have that $\alpha - \tau_z \in F$, so $\ker(h) = F \cdot \tau(Z_k)$. Moreover, since every $\alpha \in F$ must in fact stabilize every point in the fibre of x (as the trivial local translation from this fibre to itself), it follows that $F \lhd \ker(h)$. We also have $\tau(Z_k) \lhd \ker(h)$. (In fact $\tau(Z_k) \subset Z(\ker(h))$.) Using Lemma 2.9.5 and Corollary 2.9.8 we see that F is discrete. It follows that $\ker(h) \cong F \times Z_k$ as claimed. Since Z_k is a Lie group, the proof is complete.

If G is a topological group then we denote the connected component of id_G by G^0 . This is a closed normal subgroup of G. It is also standard that G is a Lie group if and only if G^0 is a Lie group.

Theorem 2.9.10. Let X be a CFR k-step nilspace and let $i \in [k]$. Then the following statements hold.

- (i) $\Theta_i(X)$ is a Lie group,
- (ii) $h(\Theta_i(X)^0) = \Theta_i(X_{k-1})^0$.

Proof. We prove the statements by induction on k. The case k = 0 being trivial, let k > 0 and suppose that the statements hold for k - 1.

To see statement (i), note first that since h is a continuous homomorphism and $\Theta_i(X)$ is closed, we have that $h(\Theta_i(X))$ is a closed subgroup of $\Theta_i(X_{k-1})$. By induction the latter is a Lie group, so by Cartan's theorem $h(\Theta_i(X))$ is a Lie group [41, Theorem 2.12.6]. By Lemma 2.9.9 we have that $\ker(h)$ is a Lie group, and so $H := \Theta_i(X) \cap \ker(h)$ is a Lie group (being a closed subgroup of $\ker(h)$). On the other hand $h(\Theta_i(X))$ is isomorphic to $\Theta_i(X)/H$ as a topological group (see [3, Theorem 1.2.6]). Therefore $\Theta_i(X)/H$ must be a Lie group (this can be seen using [13, Theorem 3.1]). We can now deduce that $\Theta_i(X)$ is itself a Lie group, by [25, Lemma A.3].

Next we show that

$$\Theta_i(X_{k-1})^0 \subseteq h(\Theta_i(X)). \tag{2.22}$$

To see this we use that $\Theta_i(X_{k-1})$ is a Lie group and so every element $\alpha \in \Theta_i(X_{k-1})^0$ is connected to the identity by a continuous path $p:[0,1] \to \Theta_i(X_{k-1})$ with $p(0)=\operatorname{id}$ and $p(1)=\alpha$. For $n\in\mathbb{N}$ let $\alpha_i=p((i-1)/n)^{-1}p(i/n)$ for $i\in[n]$ and let $\alpha=\prod_{i=1}^n\alpha_i$. Lemma 2.9.4 implies that if n is sufficiently large then for every α_i there is $\beta_i\in\Theta_i(X)$ with $h(\beta_i)=\alpha_i$. Let $\beta=\prod_{i=1}^n\beta_i$. Since h is a homomorphism we have $h(\beta)=\alpha$.

To obtain statement (ii), we first show that

$$h(\Theta_i(X)^0) = h(\Theta_i(X))^0. \tag{2.23}$$

On one hand we have $h(\Theta_i(X)^0) \subset h(\Theta_i(X))^0$ since h is continuous. On the other hand $h(\Theta_i(X))^0$ is open in $h(\Theta_i(X))$ and h is an open map $\Theta_i(X) \to h(\Theta_i(X))$ (again by [3, Theorem 1.2.6]), so $h(\Theta_i(X)^0)$ is open in $h(\Theta_i(X))$. Since it is also closed, we have that both $h(\Theta_i(X)^0)$ and its complement in $h(\Theta_i(X))^0$ are open, so this complement must be empty, since $h(\Theta_i(X))^0$ is connected, and (2.23) follows. Now in statement (ii), the inclusion $h(\Theta_i(X)^0) \subset \Theta_i(X_{k-1})^0$ follows by continuity of h, and for the opposite inclusion, note that by (2.22) we have $\Theta_i(X_{k-1})^0 \subseteq h(\Theta_i(X))^0$, so by (2.23) we have $\Theta_i(X_{k-1})^0 \subseteq h(\Theta_i(X))^0$.

Remark 2.9.11. It follows from statement (ii) above that $\Theta_i(X)^0$ is a topological extension of $\Theta_i(X_{k-1})^0$ by $\ker(h) \cap \Theta_i(X)^0$, that is we have the short exact sequence

$$0 \longrightarrow \ker(h) \cap \Theta_i(X)^0 \stackrel{l}{\longrightarrow} \Theta_i(X)^0 \stackrel{h}{\longrightarrow} \Theta_i(X_{k-1})^0 \longrightarrow 0.$$

This gives another way to obtain statement (i) by induction given (ii), using that a topological extension of a Lie group by a Lie group is again a Lie group (see [40, §2.6]).

Corollary 2.9.12. The Lie group $\Theta(X)^0$ acts transitively on the connected components of X.

Proof. We argue by induction on k. For k>0 suppose that $\Theta(X_{k-1})^0$ acts transitively on the connected components of X_{k-1} . Let τ be the isomorphism $Z_k \to \tau(Z_k) \le \Theta(X)$ in the proof of Lemma 2.9.9. Consider the group $\tau(Z_k) \cdot \Theta(X)^0 = \{\tau(z) \ \alpha : z \in Z_k, \ \alpha \in \Theta(X)^0\} \le \Theta(X)$, and let M be any connected component of X. We claim that $\tau(Z_k) \cdot \Theta(X)^0$ acts transitively on M. Indeed, for any $x, y \in M$, by Theorem 2.9.10 (ii) and the induction hypothesis, there exists $\alpha \in \Theta(X)^0$ such that $\pi_{k-1} \circ \alpha(x) = \pi_{k-1}(y)$. Combining this with the transitivity of Z_k on the fibres of π_{k-1} we deduce that there exists

 $\alpha' \in \tau(Z_k) \cdot \Theta(X)^0$ such that $\alpha'(x) = y$, which proves our claim. It follows that $\Theta(X)$ acts transitively on M, and we can then deduce that $\Theta(X)^0$ also acts transitively on M by standard results (e.g. [23, Chapter II, Theorem 3.2 and Proposition 4.3]).

Example 2.9.13. The main results of this section so far can be illustrated with the Heisenberg nilmanifold H/Γ , where $H = \begin{pmatrix} 1 & \mathbb{R} & \mathbb{R} \\ 1 & \mathbb{R} \end{pmatrix}$, $\Gamma = \begin{pmatrix} 1 & \mathbb{Z} & \mathbb{Z} \\ 1 & \mathbb{Z} \end{pmatrix}$ (recall [8, Example 2.3.2]). Let X denote H/Γ equipped with the cubes associated with the lower central series on H. We then have that X is a 2-step compact nilspace, with 1-step factor $X_1 = \mathbb{T}^2$ equipped with the standard degree-1 cubes, and we compute that

$$\Theta_1(X) = \left(\begin{smallmatrix} 1 & \mathbb{R} & \mathbb{T} \\ & 1 & \mathbb{R} \\ & & 1 \end{smallmatrix}\right), \ \ \Theta_2(X) = \left(\begin{smallmatrix} 1 & 0 & \mathbb{T} \\ & 1 & 0 \\ & & 1 \end{smallmatrix}\right) \cong \mathbb{T}, \ \ \Theta(X_1) = \mathbb{T}^2.$$

We also check that $\ker(h) = \begin{pmatrix} 1 & \mathbb{Z} & \mathbb{T} \\ 1 & \mathbb{Z} \end{pmatrix} \cong \mathbb{Z}^2 \times \mathbb{T}$. Note that the structure groups of X are tori $(Z_2 = \mathbb{T}, Z_1 = \mathbb{T}^2)$. Thus X is also an example of the type of connected CFR nilspace to which we turn next.

2.9.1 Toral nilspaces

Recall that a CFR abelian group Z has torsion-free dual group if and only if $Z \cong \mathbb{T}^n$ for some $n \geq 0$.

Definition 2.9.14. We say that a *k*-step nilspace X is *torsion-free* if every structure group of X has torsion-free dual group. We call X a *toral nilspace* if it is CFR and torsion-free.

Thus a *k*-step compact nilspace is toral if and only if its structure groups are all tori. One of the reasons for treating these nilspaces in particular is that they occur naturally in the study of the regularizations of ultralimits of functions on abelian groups, which can be used to prove inverse theorems for the Gowers uniformity norms (see the characteristic-0 case of [39, Definition 1.3]).

Lemma 2.9.15. Every toral nilspace is connected.

Proof. Recall the general fact that a fibre bundle is connected if the base space and the fibre are connected (see [38, Part I, §2.12]). Thus if Z has finite rank and is connected then a continuous Z-bundle over a connected space is connected. Now given a k-step toral nilspace X, with factors X_i , $i \in [k]$, a simple induction on i using this fact shows that X is connected.

Remark 2.9.16. The converse of Lemma 2.9.15 is false. Consider for example the connected 2-step nilspace X consisting of the circle group \mathbb{T} with cube sets $C^n(G_{\bullet})$ for the degree-2 filtration G_{\bullet} with $G_0 = G_1 = \mathbb{T}$ and $G_2 = \{0, 1/2\} \cong \mathbb{Z}/2\mathbb{Z}$ (viewing \mathbb{T} as [0,1) with addition mod 1). A straightforward calculation shows that $x \sim_1 y$ in X if and only if $x - y \in G_2$, whence the structure group Z_2 is the disconnected group G_2 .

We can now establish the main result of this section.

Theorem 2.9.17. Let X be a k-step toral nilspace, let $G = \Theta(X)^0$, let G_{\bullet} denote the degree-k filtration $(\Theta_i(X)^0)_{i\geq 0}$, and for an arbitrary fixed $x\in X$ let $\Gamma=\operatorname{Stab}_G(x)$. Then X is isomorphic as a compact nilspace to the nilmanifold G/Γ with cube sets $C^n(X)=(C^n(G_{\bullet})\cdot\Gamma^{\{0,1\}^n})/\Gamma^{\{0,1\}^n}$, $n\geq 0$.

Note that for the nilmanifold obtained here, every group in the filtration G_{\bullet} is connected. (The fact that G_{\bullet} here is indeed a filtration follows from the fact that if H_1, H_2 are closed subgroups of a topological group and one of these subgroups is connected, then $[H_1, H_2]$ is connected.)

Proof. We argue by induction on k. For k = 1, by Lemma 2.1.4 we have that X is the principal homogeneous space of a torus with cubes being the projections of the standard cubes on the torus, so the statement holds. Supposing then that the statement holds for k - 1, we fix $x \in X$ and let $\Gamma = \operatorname{Stab}_G(x)$.

We first claim that Γ is discrete. To see this, note that $h(\Gamma)$ is a subgroup of the stabilizer of $\pi_{k-1}(x)$ in $\Theta(X_{k-1})$, so by induction it is discrete. Then using that $h^{-1}(h(\Gamma))$ is a union of cosets of $\ker(h)$, we have that it suffices to show that $\Gamma \cap \ker(h)$ is discrete. But this follows from Lemma 2.9.9, since no non-trivial element of $\tau(Z_k)$ stabilizes x.

By Corollary 2.9.12 the Lie group $\Theta(X)^0$ acts transitively on the connected space X, and it follows that X is homeomorphic to the coset space G/Γ (see [23, Chapter II, Theorem 3.2]). Therefore Γ is cocompact.

It now only remains to determine the cubes on X. Recall from [8, Definition 3.2.38] that two cubes $c_1, c_2 \in C^n(X)$ are said to be translation equivalent if there is an element $c \in C^n(G_{\bullet})$ such that $c_2(v) = c(v) \cdot c_1(v)$. We claim that for every cube $c \in C^n(X)$ there is a cube $c' \in C^n(X)$ that is translation equivalent to the constant x cube and such that $\pi_{k-1} \circ c = \pi_{k-1} \circ c'$. Indeed, given $c \in C^n(X)$, we have $\pi_{k-1} \circ c \in C^n(X_{k-1})$, and by induction the latter cube is translation equivalent to the cube with constant value $x' = \pi_{k-1}(x)$, i.e. it is of the form $\tilde{c} \cdot x'$ for some cube \tilde{c} on the group $\Theta(X_{k-1})^0$ with the filtration $(\Theta_i(X_{k-1})^0)_{i\geq 0}$. Now by the unique factorization result for these cubes [8, Lemma 2.2.5], we have $\tilde{c} = \tilde{g}_0^{F_0} \cdots \tilde{g}_{2^n-1}^{F_{2^n-1}}$ where $\tilde{g}_j \in \Theta_{\operatorname{codim}(F_j)}(X_{k-1})^0$. Theorem 2.9.10 (ii) then tells us that for each $j \in [0,2^n)$ there is $g_j \in \Theta_{\operatorname{codim}(F_j)}(X)^0$ such that $h(g_j) = \tilde{g}_j$. Let c^* be the cube in $C^n(\Theta(X)^0)$ defined by $c^* = g_0^{F_0} \cdots g_{2^n-1}^{F_{2^n-1}}$. Let $c' = c^* \cdot x$. This is in $C^n(X)$, and is translation equivalent to the constant x cube. Moreover, by construction we have

$$\pi_{k-1} \circ \mathbf{c}' = \pi_{k-1}(\mathbf{c}^* \cdot \mathbf{x}) = \left(\prod_j h(g_j)^{F_j}\right) \cdot \mathbf{x}' = \left(\prod_j \tilde{g}_j^{F_j}\right) \cdot \mathbf{x}' = \tilde{\mathbf{c}} \cdot \mathbf{x}' = \pi_{k-1} \circ \mathbf{c}.$$

This proves our claim.

It follows from [8, Theorem 3.2.19] and the definition of degree-k bundles (in particular [8, (3.5)]) that $c - c' \in C^n(\mathcal{D}_k(Z_k))$. But then using translations from $\tau(Z_k)$ we can correct c' further to obtain c, thus showing that c is itself a translation cube with translations from $\Theta(X)^0$. (Such a correction procedure has been used in previous arguments, see for instance the proof of [8, Lemma 3.2.25].)

We have thus shown that $C^n(X) \subset (C^n(G_{\bullet}) \cdot \Gamma^{\{0,1\}^n}) / \Gamma^{\{0,1\}^n}$. The opposite inclusion is clear, by definition of the groups $\Theta_i(X)$.

Remark 2.9.18. In Theorem 2.9.17 we can replace the groups G_i by *simply connected* Lie groups \tilde{G}_i while conserving the nilspace isomorphism with X. More precisely, let \tilde{G} be the universal covering group of the connected Lie group G, with covering homomorphism $p: \tilde{G} \to G$ (see [35, §47, Theorem 61]). Let $\tilde{\Gamma} = p^{-1}(\Gamma)$, and for each $i \geq 0$, let \tilde{G}_i be the identity component of the closed (hence Lie) subgroup $p^{-1}(G_i)$ of \tilde{G} . Then $\tilde{G}_{\bullet} = (\tilde{G}_i)_{i \geq 0}$ is a degree-k filtration of closed connected (hence simply-connected [11, Corollary 1.2.2]) Lie subgroups of \tilde{G} . It can be checked in a straightforward way that $\tilde{G}/\tilde{\Gamma}$ with the cubes determined by \tilde{G}_{\bullet} is isomorphic to X as a compact nilspace.

We know by a theorem of Palais and Stewart that X is the total space of a principal torus bundle over a torus if and only if it is a 2-step nilmanifold (see [34, Theorem 1]). Theorem 2.9.17 can be viewed as a generalization of this result for arbitrary step, in the category of compact nilspaces (in particular, making use of the assumed cube structure).

One may wonder whether, by adding some simple assumption in Theorem 2.9.17, the description of the filtered nilmanifold can then be made even more precise. For instance, having seen Example 2.9.13, we may seek some simple additional condition under which Theorem 2.9.17 yields the Heisenberg nilmanifold H/Γ . In this direction we have the following result.

Proposition 2.9.19. *Let* X *be a 2-step toral nilspace of dimension 3. Then* X *is isomorphic to one of the following compact nilspaces:*

- (i) The torus $\mathbb{R}^3/\mathbb{Z}^3$ with cube structure determined by a filtration of degree at most 2 on \mathbb{R}^3 .
- (ii) The Heisenberg nilmanifold H/Γ with cube structure determined by the lower central series on H.

Here the dimension is that of X when viewed as a manifold (a view which we noted after Lemma 2.5.3).

Proof. Applying Theorem 2.9.17 and Remark 2.9.18, we have that X is isomorphic as a compact nilspace to G/Γ for a simply-connected Lie group G with a filtration G_{\bullet} of degree at most 2 consisting of closed (therefore Lie) simply-connected groups $G_i \leq G$, and a discrete cocompact subgroup $\Gamma \leq G$. The assumption that X has dimension 3 implies that G is 3-dimensional.

If X is a 1-step nilspace then X is isomorphic to the 3-torus $\mathbb{R}^3/\mathbb{Z}^3$ in case (i), with the standard degree-1 cube structure.

If X is not 1-step then, if G is abelian, we have again that X is isomorphic to $\mathbb{R}^3/\mathbb{Z}^3$ but now the cube structure is given by a degree-2 filtration on \mathbb{R}^3 (for example it could be the filtration determining the degree-2 structure on $\mathbb{R}^3/\mathbb{Z}^3$ defined in [8, Definition 2.2.30]). The remaining possibility is that G is non-abelian 2-step nilpotent. The simply-connected Lie group G is uniquely determined by its Lie algebra, which is a real non-abelian, 2-step nilpotent Lie algebra of dimension 3. It follows from the Bianchi classification [4] that the only such Lie algebra is the 3-dimensional Heisenberg Lie algebra, and so G is the Heisenberg group. Finally, we have that the only degree-2 filtration of closed connected subgroups of G is the lower central series.

To close this section, we prove additional pleasant properties of toral nilspaces. In particular, property (iii) below gives a simple relation between the translation groups and the structure groups.

Proposition 2.9.20. Let X be a k-step toral nilspace. For each $j \in [k]$ let $h_{j-1,j}$ denote the continuous surjective homomorphism $\Theta(X_j)^0 \to \Theta(X_{j-1})^0$ defined as in Lemma 2.9.3, and for each $i \in [k]$ let $h_i = h_{i,i+1} \circ \cdots \circ h_{k-1,k} : \Theta(X)^0 \to \Theta(X_i)^0$. Fix an arbitrary element $x \in X$, and for each $i \in [k]$ let $\Gamma_i = \Theta_i(X)^0 \cap \operatorname{Stab}_{\Theta(X)^0}(x)$. Then for each i, the following statements hold:

- (i) $\Theta_i(X)^0$ acts transitively on each fibre of $\pi_{i-1}: X \to X_{i-1}$.
- (ii) $\ker(h_i) = \Gamma_i \cdot \Theta_{i+1}(X)^0$.
- $\label{eq:continuity} \mbox{\it (iii)} \ \Theta_i(X)^0 \, / \, \left(\Gamma_i \cdot \Theta_{i+1}(X)^0 \right) \, \cong \, Z_i.$

P. CANDELA

Proof. We prove (i) by downward induction on i. For the case i = k, since each fibre of π_{k-1} is a homogeneous space of Z_k , it suffices to show that $\Theta_k(X)^0 = \tau(Z_k)$. This is immediate from the following equality, which actually holds for a general nilspace and is established in [8, Lemma 3.2.37]:

$$\Theta_k(X) = \tau(Z_k). \tag{2.24}$$

Suppose now that $i \leq k-1$ and that the statement holds for i+1, and note that any fibre of $\pi_{i-1} = \pi_{i-1,i} \circ \pi_i$ has the form $\pi_i^{-1}(\pi_{i-1,i}^{-1}(y))$ for some $y \in X_{i-1}$. Given w,w' in such a fibre, we have $\pi_i(w), \pi_i(w') \in \pi_{i-1,i}^{-1}(y)$. Since $h_i(\Theta_i(X)^0) = \Theta_i(X_i)^0 = \tau(Z_i)$, we have that this acts transitively on $\pi_{i-1,i}^{-1}(y)$ and so there is $\alpha \in \Theta_i(X)^0$ such that $\pi_i(\alpha(w)) = \pi_i(w')$. Then, since $\Theta_i(X)^0$ contains $\Theta_{i+1}(X)^0$, by induction it acts transitively on each π_i -fibre and so there exists $\alpha' \in \Theta_i(X)^0$ such that $\alpha'\alpha(w) = w'$. This proves statement (i).

To see statement (ii), note that $\Theta_{i+1}(X)^0 \leq \ker(h_i)$, since h_i sends $\Theta_{i+1}(X)^0$ into the trivial group $\Theta_{i+1}(X_i)^0$. We also have that $\Gamma_i \subset \ker(h_i)$, because by (2.24) we have $h_i(\Gamma_i) \subset \tau(Z_i)$ but also each element of $h_i(\Gamma_i)$ fixes $\pi_i(x)$, so must be the trivial translation in $\tau(Z_i)$. It follows that $\Gamma_i \cdot \Theta_{i+1}(X)^0 \subset \ker(h_i)$. To see the opposite inclusion, note that for $\alpha \in \ker(h_i)$ we have from the definitions that $\pi_i(\alpha(x)) = h_i(\alpha)(\pi_i(x)) = \pi_i(x)$. Therefore, by statement (i) there exists $\alpha' \in \Theta_{i+1}(X)^0$ such that $\alpha'\alpha(x) = x$, so $\alpha' \alpha \in \Gamma_i$, and so $\alpha \in \Theta_{i+1}(X)^0 \cdot \Gamma_i$.

Statement (iii) now follows from (ii), the first isomorphism theorem for h_i , and (2.24) for k = i.

Bibliography

- [1] A. Adem, J. F. Davis, *Topics in transformation groups*, Handbook of geometric topology, 1–54, North-Holland, Amsterdam, 2002. 33
- [2] R. D. Anderson, *Hilbert space is homeomorphic to the countable infinite product of lines*, Bull. Amer. Math. Soc. **72** (1966), 515–519. 23
- [3] H. Becker, A. S. Kechris, *The descriptive theory of polish group actions*, London Mathematical Society Lecture Note Series, 232. Cambridge University Press, Cambridge, 1996. 46, 50
- [4] L. Bianchi, On the three-dimensional spaces which admit a continuous group of motions, Gen. Relativity Gravitation **33** (2001), no. 12, 2171–2253. 53
- [5] P. Billingsley, Convergence of probability measures, Second edition. John Wiley & Sons, Inc., New York, 1999. 13
- [6] N. Bourbaki, *Elements of mathematics. General topology. Part 1*, Hermann, Paris; Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont. 1966. 8, 11, 28, 39
- [7] N. Bourbaki, *Integration. I. Chapters 1–6*, translated from the 1959, 1965 and 1967 French originals by Sterling K. Berberian. Elements of Mathematics (Berlin). Springer-Verlag, Berlin, 2004. 12
- [8] P. Candela, *Notes on nilspaces: algebraic aspects*, Discrete Analysis, 2017, Paper No. 15, 59 pp. 3, 5, 6, 7, 8, 9, 10, 11, 13, 14, 15, 16, 17, 18, 19, 20, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 42, 44, 45, 46, 47, 48, 51, 52, 53, 54
- [9] O. A. Camarena, B. Szegedy, *Nilspaces, nilmanifolds and their morphisms*, preprint. arXiv:1009.3825. 3, 4, 37
- [10] A. Censor, D. Grandini, *Borel and continuous systems of measures*, Rocky Mountain J. Math. **44** (2014), no. 4, 1073–1110. 12, 13, 15
- [11] L. Corwin, F. P. Greenleaf, *Representations of nilpotent Lie groups and their applications, Part 1: Basic theory and examples*, Cambridge studies in advanced mathematics 18, C.U.P. 1990. 52
- [12] A. Deitmar, S. Echterhoff, *Principles of harmonic analysis*, Universitext. Springer, New York, 2009. 15, 16
- [13] A. M. Gleason, Groups without small subgroups, Ann. of Math. (2) 56 (1952), 193–212. 50
- [14] A. M. Gleason, *Spaces with a compact Lie group of transformations*, Proc. Amer. Math. Soc. **1** (1950), 35–43. 33
- [15] W. T. Gowers, A new proof of Szemerédi's theorem, GAFA 11 (2001), 465–588. 4

P. CANDELA

- [16] W. T. Gowers, *Generalizations of Fourier analysis, and how to apply them*, Bull. Amer. Math. Soc. **54** (2017), no. 1, 1–44. 4
- [17] B. Green, T. Tao, *An arithmetic regularity lemma, an associated counting lemma, and applications*, An irregular mind, 261-334, Bolyai Soc. Math. Stud., 21, Janos Bolyai Math. Soc., Budapest, 2010. 4, 5
- [18] B. Green, T. Tao, Linear equations in primes, Ann. of Math. (2) 171 (2010), no. 3, 1753–1850. 4, 5
- [19] B. Green, T. Tao, T. Ziegler, An inverse theorem for the Gowers $U^{s+1}[N]$ -norm, Ann. of Math. 176 (2012), no. 2, 1231–1372. 4
- [20] Y. Gutman, F. Manners, P. P. Varjú, *The structure theory of nilspaces I*, preprint. arXiv:1605.08945.
- [21] Y. Gutman, F. Manners, P. P. Varjú, *The structure theory of nilspaces II: Representation as nilmanifolds*, preprint. arXiv:1605.08948. 4
- [22] Y. Gutman, F. Manners, P. P. Varjú, *The structure theory of nilspaces III: Inverse limit representations and topological dynamics*, preprint. arXiv:1605.08950. 4
- [23] S. Helgason, *Differential geometry, Lie groups, and symmetric spaces*, Pure and Applied Mathematics, 80. Academic Press, Inc., New York-London, 1978. 51, 52
- [24] K. H. Hofmann, S. A. Morris, *The structure of compact groups. A primer for the student a handbook for the expert*, Second revised and augmented edition. de Gruyter Studies in Mathematics, 25. Walter de Gruyter & Co., Berlin, 2006. 38
- [25] B. Host, B. Kra, *Nonconventional ergodic averages and nilmanifolds*, Ann. of Math. (2) **161** (2005), no. 1, 397–488. 3, 4, 50
- [26] B. Host, B. Kra, *Parallelepipeds, nilpotent groups, and Gowers norms*, Bull. Soc. Math. France **136** (2008), no. 3, 405–437. 3
- [27] D. Husemoller, *Fibre bundles*, third edition, Graduate Texts in Mathematics, 20, Springer-Verlag, New York, 1994. 36
- [28] A. S. Kechris, *Classical descriptive set theory*, Graduate Texts in Mathematics, 156. Springer-Verlag, New York, 1995. 3, 21, 23, 31, 32, 36
- [29] A. Kleppner, *Measurable homomorphisms of locally compact groups*, Proc. Amer. Math. Soc. **106** (1989), no. 2, 391–395. 32
- [30] A. Leibman, *Rational sub-nilmanifolds of a compact nilmanifold*, Ergodic Theory Dynam. Systems **26** (2006), no. 3, 787–798. 5
- [31] A. I. Mal'cev, *On a class of homogeneous spaces*, Amer. Math. Soc. Translation **1951**, (1951). no. 39, 33 pp. 3

- [32] J. R. Munkres, Topology, Second Edition, Prentice Hall, 2000. 8, 12, 21, 24, 33, 46
- [33] R. S. Palais, The classification of G-spaces, Mem. Amer. Math. Soc. No. 36, 1960. 11
- [34] R. S. Palais, T. E. Stewart, *Torus bundles over a torus*, Proc. Amer. Math. Soc. **12** (1961), no. 1, 26–29. 53
- [35] L. Pontrjagin, *Topological groups*, Translated from the Russian by Emma Lehmer, Princeton Mathematical Series, v. 2, Princeton University Press, Princeton, 1939. 52
- [36] B. D. Ripley, *The disintegration of invariant measures*, Math. Proc. Cambridge Philos. Soc. **79** (1976), no. 2, 337–341. 13
- [37] W. Rudin, *Fourier analysis on groups*, Interscience Tracts in Pure and Applied Mathematics, no. 12, Interscience Publishers, New York-London 1962. 12, 13, 21
- [38] N. Steenrod, *The topology of fibre bundles*, Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1999. 51
- [39] B. Szegedy, On higher order Fourier analysis, preprint. arXiv:1203.2260. 4, 51
- [40] T. Tao, *Hilbert's fifth problem and related topics*, Graduate Studies in Mathematics, 153. American Mathematical Society, Providence, RI, 2014. 50
- [41] V. S. Varadarajan, *Lie groups, Lie algebras, and their representations*, Graduate Texts in Mathematics, 102, Springer-Verlag, New York, 1984. 50

AUTHOR

Pablo Candela
Universidad Autónoma de Madrid
Madrid, Spain
pablo [dot] candela [at] uam [dot] es
http://verso.mat.uam.es/~pablo.candela/index.html