Chapter 1

Introduction

Our aim in this memoir is to present some new techniques to establish the Universal Coefficient Theorem in C^*-algebra K-theory, and some new necessary and sufficient conditions for the Universal Coefficient Theorem to hold for all nuclear C^*-algebras.

Unless otherwise stated, anything in this introduction called A or B is a *separable* C^*-algebra.

1.1 The universal coefficient theorem

A C^*-algebra A satisfies the *Universal Coefficient Theorem* (UCT) of Rosenberg and Schochet [55] if for any C^*-algebra B, there is a canonical short exact sequence

$$0 \to \mathrm{Ext}(K_*(A), K_*(B)) \to KK(A, B) \to \mathrm{Hom}(K_*(A), K_*(B)) \to 0.$$

Equivalently (see [55, p. 456] or [60, Proposition 5.2]), A satisfies the UCT if it is KK-equivalent to a commutative C^*-algebra.

The UCT is known to hold for a large class of C^*-algebras. The fundamental examples are the C^*-algebras in the *bootstrap class* $\mathcal{N}$. This is the smallest collection of separable, nuclear C^*-algebras that contains all type I C^*-algebras, and that is closed under the following operations: extensions; stable isomorphisms; inductive limits; and crossed products by $\mathbb{R}$ and $\mathbb{Z}$. Rosenberg and Schochet [55] showed that any C^*-algebra in $\mathcal{N}$ satisfies the UCT. Another important class of examples was established by Tu in [64, Proposition 10.7]; building on the work of Higson and Kasparov [35] on the Baum–Connes conjecture for a-T-menable groups, Tu showed that the groupoid[1] C^*-algebra of any a-T-menable groupoid satisfies the UCT. In particular, Tu's work applies to the groupoid C^*-algebras of amenable groupoids.

There has been other significant work giving sufficient conditions for the UCT to hold, and in some cases also necessary conditions as well as the work mentioned already, one has for example [60, Proposition 5.2], [53, Corollary 8.4.6], [21], [43, Remark 2.17], [6, Theorem 4.17], [4], and [5]. Nonetheless, the bootstrap class and the class of C^*-algebras of a-T-menable groupoids, which are defined in terms of *global* properties of the C^*-algebras involved, remain the most important classes of C^*-algebras known to satisfy the UCT.

[1]To be more precise, we need standard assumptions so that the groupoid C^*-algebra is defined and separable. Here, appropriate assumptions are that the groupoid is locally compact, Hausdorff, and second countable, and that it admits a Haar system.

On the other hand, Skandalis [60, p. 571] has shown[2] that there are C^*-algebras that do not satisfy the UCT. Skandalis's examples are quite concrete; they are reduced group C^*-algebras of countably infinite hyperbolic groups with property (T), and in particular are exact [44, Section 6.E]. Looking to more exotic examples, failures of exactness can also be used to produce non-UCT C^*-algebras; see for example [14, Remark 4.3].

Despite these counterexamples, there are no known *nuclear* C^*-algebras that do not satisfy the UCT. Whether or not the UCT holds for all nuclear C^*-algebras is a particularly important open problem. One reason for this is the spectacular recent progress (see for example [11, 23, 24, 26, 27, 42, 50, 63]) in the Elliott program [22] to classify simple, separable, nuclear C^*-algebras by K-theoretic invariants. Establishing the range of validity of the UCT is now the only barrier to getting the "best possible" classification result in this setting.

On the other hand, work inspired by the Elliott program has led to recent, and again spectacular, success in the general structure theory of nuclear C^*-algebras, including the recent solution of a large part of the Toms–Winter conjecture [12, 13]. Our motivation in the current paper is to try to bridge the gap between properties that are relevant in this structure theory – in particular the theory of nuclear dimension [70] introduced by Winter and Zacharias – and properties that imply the UCT. In particular, our aim is to give *local* conditions that imply the UCT, in contrast to the global conditions from the work of Rosenberg and Schochet [55] and Tu [64] mentioned above.

1.2 Decompositions and the main theorem

We now introduce our sufficient condition for the UCT. For the statement below, if X is a metric space, S is a subset of X, $x \in X$, and $\varepsilon > 0$ we write "$x \in_\varepsilon S$" if there exists $s \in S$ with $d(x, s) < \varepsilon$.

Definition 1.1. Let $\mathcal{C}$ be a class of unital C^*-algebras. A unital C^*-algebra[3] A *decomposes over* $\mathcal{C}$ if for every finite subset X of the unit ball of A and every $\varepsilon > 0$ there exist C^*-subalgebras C, D, and E of A that are in the class $\mathcal{C}$ and contain 1_A, and a positive contraction $h \in E$ such that

(i) $\|[h, x]\| < \varepsilon$ for all $x \in X$;

(ii) $hx \in_\varepsilon C$, $(1 - h)x \in_\varepsilon D$, and $h(1 - h)x \in_\varepsilon E$ for all $x \in X$;

(iii) for all e in the unit ball of E, $e \in_\varepsilon C$ and $e \in_\varepsilon D$.

[2] See also the exposition in [34, Sections 6.1 and 6.2].

[3] Not necessarily separable. For applications to the UCT, only the separable case is relevant, but the definition admits interesting examples in the non-separable case, and it seems plausible there will be other applications.

One should think of C and D as being approximately (unitizations of) ideals in A such that $C + D = A$, and E being approximately equal to (the unitization of) $C \cap D$. We will discuss examples later.

Here, is our main theorem, which was inspired by our earlier work on the Künneth formula (in collaboration with Oyono-Oyono) [48, 67], and by our earlier work on finite dynamical complexity and finite decomposition complexity (in collaboration with Guentner and Tessera) [29, 31]. See Corollary 7.5 below for the proof.

Theorem 1.2. *If A is a separable, unital C^*-algebra that decomposes over the class of separable, nuclear C^*-algebras that satisfy the UCT, then A is nuclear and satisfies the UCT.*

One can thus think of decomposability as an addition to the closure operations that are used in the definition of the bootstrap class $\mathcal{N}$.

1.3 C^*-algebras with finite complexity

Following the precedent established by [30] in coarse geometry, the notion of decomposability suggests a complexity hierarchy on C^*-algebras.

Definition 1.3. Let $\mathcal{D}$ denote a class of unital C^*-algebras. For an ordinal number α,

(i) if $\alpha = 0$, let $\mathcal{D}_0$ be the class of C^*-algebras D that are locally[4] in $\mathcal{D}$;

(ii) if $\alpha > 0$, let $\mathcal{D}_\alpha$ be the class of C^*-algebras that decompose over C^*-algebras in $\bigcup_{\beta<\alpha} \mathcal{D}_\beta$.

A unital C^*-algebra D has *finite complexity relative to* $\mathcal{D}$ if it is in $\mathcal{D}_\alpha$ for some α. If $\mathcal{D}$ is the class of finite-dimensional C^*-algebras, we just say that D has *finite complexity*.

If a unital C^*-algebra D has finite complexity relative to $\mathcal{D}$, the *complexity rank of D relative to* $\mathcal{D}$ is the smallest α such that D is in $\mathcal{D}_\alpha$. If $\mathcal{D}$ is the class of finite-dimensional C^*-algebras, we just say the *complexity rank of D* with no additional qualifiers.

The following result is equivalent to Theorem 1.2 above. However, we think the reframing in terms of complexity is quite suggestive.

Theorem 1.4. *Let $\mathcal{C}$ be a class of separable, unital, nuclear C^*-algebras that satisfy the UCT. Then, the class of separable, unital C^*-algebras that have finite complexity relative to $\mathcal{C}$ consists of nuclear C^*-algebras that satisfy the UCT.*

---

[4] A C^*-algebra is *locally* in a class $\mathcal{D}$ if for any finite subset X of D and any $\varepsilon > 0$ there is a C^*-subalgebra C of D that is in $\mathcal{D}$, and such that $x \in_\varepsilon C$ for all $x \in X$.

In particular, every separable C^*-algebra of finite complexity is nuclear and satisfies the UCT.

Examples 1.5. We can now give some non-trivial examples of C^*-algebras that decompose over natural, simpler, classes.

(i) In Proposition A.1, we show that for
$$2 \leq n < \infty,$$
the Cuntz algebra $\mathcal{O}_n$ has complexity rank one.

(ii) In [31], Guentner and the authors introduced "finite dynamical complexity" for groupoids, which also comes with a notion of complexity rank. In Proposition A.8 we show that if G is a locally compact, Hausdorff, étale, principal, ample groupoid with compact base space, then the complexity rank of $C_r^*(G)$ is bounded above by that of G. The class of groupoids with finite dynamical complexity is quite large; see Examples A.9 and A.11 below.

Combining part (ii) above with Theorem 1.4 gives a new proof of the UCT for the groupoid C^*-algebras of a large class of groupoids. However, we cannot claim any genuinely new examples; this is because the groupoids involved are all amenable, so the UCT for their C^*-algebras also follows from Tu's theorem [64] (see Remark A.13 below for more details).

1.4 Kirchberg algebras

Generalizing the Cuntz algebras from (i) above, recall that a *Kirchberg algebra* is a separable, nuclear C^*-algebra A such that for any non-zero $a \in A$, there are $b, c \in A$ such that $bac = 1_A$. Kirchberg algebras are closely connected to the UCT problem for nuclear C^*-algebras thanks to the following theorem of Kirchberg; see [53, Corollary 8.4.6] or [43, Remark 2.17].

Theorem 1.6 (Kirchberg). *To establish the UCT for all separable, nuclear C^*-algebras, it suffices to establish the UCT for any Kirchberg algebra with zero K-theory.* ∎

Theorems 1.4 and 1.6 imply that if any Kirchberg algebra with zero K-theory has finite complexity, then the UCT holds for all separable, nuclear C^*-algebras. Conversely, if the UCT holds for all separable, nuclear C^*-algebras, then from the Kirchberg–Phillips classification theorem [42, 50] (see also [53, Corollary 8.4.2] for the precise statement we want here), any unital Kirchberg algebra with zero K-theory will be isomorphic to the Cuntz algebra $\mathcal{O}_2$, and so will have complexity rank one by Examples 1.5 (i). We summarize this discussion in the theorem below.

Theorem 1.7. *The following are equivalent:*

(i) *Any Kirchberg algebra with zero K-theory has complexity rank one.*

(ii) *All separable nuclear C^*-algebras satisfy the UCT.* ■

Generalizing Examples 1.5 (i) above Jaime and the first author show in [37] that a Kirchberg algebra *that satisfies the UCT* has complexity rank one if and only if its K_1 group is torsion free, and that moreover any UCT Kirchberg algebra has complexity rank at most two. From Theorem 1.7, if one could prove this without the UCT assumption, then the UCT for all separable nuclear C^*-algebras would follow.

The paper [37] also discusses several other connections between complexity rank, real rank zero, and nuclear dimension. We will not go into this any more deeply here; suffice to say that these other connections inspired us to make the following conjectures.

Conjecture 1.8. *Any separable unital C^*-algebra with real rank zero and finite nuclear dimension has finite complexity.*

Conjecture 1.9. *Any separable unital C^*-algebra with finite nuclear dimension has finite complexity relative to the class of subhomogeneous*[5] *C^*-algebras.*

Thanks to Theorem 1.7 and the fact that all Kirchberg algebras have nuclear dimension one (see [9, Theorem G]) and real rank zero (see [72]), either of these conjectures implies the UCT for all separable, nuclear C^*-algebras. There are many other related conjectures one could reasonably make that imply the UCT for all nuclear C^*-algebras. About the strongest such conjecture would be that any separable, nuclear C^*-algebra with real rank zero has finite complexity[6]. One of the weakest is that any Kirchberg algebra with zero K-theory has finite complexity.

1.5 A local reformulation of the UCT

We now discuss the methods that go into the proof of Theorem 1.2.

---

[5] Recall that a C^*-algebra C is *subhomogeneous* if there is $N \in \mathbb{N}$ and a compact Hausdorff space X such that C is a C^*-subalgebra of $M_N(C(X))$; see for example [8, Section IV.1.4] for background.

[6] It would also be natural to drop the real rank zero assumption, and then only ask for finite complexity relative to the subhomogeneous C^*-algebras, or even just relative to the type I C^*-algebras.

In our earlier work [68], we introduced *controlled KK-theory* groups $KK_\varepsilon(X, B)$ associated to a C^*-algebra B, a finite subset X of a C^*-algebra A and a constant $\varepsilon > 0$. Very roughly (we give more details below), one defines these by representing A in "general position" inside the stable multiplier algebra $M(B \otimes \mathcal{K})$ of B. The group $KK_\varepsilon(X, B)$ then consists of the "part of the K-theory of B that commutes with X, up to ε".

To be more precise about this, assume that A and B are C^*-algebras, and assume for simplicity[7] that A is unital. Let $\pi : A \to M(B \otimes \mathcal{K})$ be a faithful, unital, and strongly unitally absorbing[8] representation. Fixing such a representation, identify A with a diagonal subalgebra of $M_2(M(B \otimes \mathcal{K}))$ via the representation $\pi \oplus \pi$. For a finite subset X of the unit ball of A and $\varepsilon > 0$, define $\mathcal{P}_\varepsilon(X, B)$ to be the set of projections in $M_2(M(B \otimes \mathcal{K}))$ such that p-$\left(\begin{smallmatrix} 1 & 0 \\ 0 & 0 \end{smallmatrix}\right)$ is in $M_2(B \otimes \mathcal{K})$, and such that $\|[p, x]\| < \varepsilon$ for all $x \in X$. The associated *controlled KK-theory group*[9] is then defined to be the set

$$KK^0_\varepsilon(X, B) := \pi_0(\mathcal{P}_\varepsilon(X, B))$$

of path components in $\mathcal{P}_\varepsilon(X, B)$. One can show that this group is determined up to canonical isomorphism by the subset inclusion $X \subseteq A$, by B, and by ε; it does not depend on the choice of representation.

Note that if $X = \varnothing$, then $KK^0_\varepsilon(\varnothing, B)$ is canonically isomorphic to the usual K-theory group $K_0(B)$ (for any ε); this is what we mean when we say $KK_\varepsilon(X, B)$ consists of the "part of the K-theory of B that commutes with X, up to ε".

Now, if $0 < \delta \leq \varepsilon$ and if $Y \supseteq X$ are finite subsets of A_1, then there is an inclusion $\mathcal{P}_\delta(Y, B) \subseteq \mathcal{P}_\varepsilon(X, B)$ that induces a "forget control map"

$$KK_\delta(Y, B) \to KK_\varepsilon(X, B).$$

In [68, Theorem 1.1], we showed that there is a short exact "Milnor sequence" relating the inverse system built from these forget control maps to the usual KK-group $KK(A, B)$; see Theorem 2.13 below for details. This sequence is an analogue of the Milnor sequence appearing in Schochet's work [56, 57]; however, unlike Schochet's version, it is local in nature, and does not require the UCT.

Our first goal in this memoir is to use the Milnor sequence to establish the following "local reformulation" of the UCT.

---

[7]The theory also works for C^*-algebras that are not unital, but the definitions are a little more complicated.

[8]Roughly, a strongly unitally absorbing representation is one that satisfies the conclusion of Voiculescu's theorem for all representations of A on Hilbert B-modules; for the current discussion, it is just important that such a representation always exists. See Definition 2.5 below for details.

[9]It is canonically a group, with the operation given by Cuntz sum in an appropriate sense.

Theorem 1.10. *Let A be a unital C^*-algebra. Then, the following are equivalent:*

(i) *A satisfies the UCT.*

(ii) *Let B be a separable C^*-algebra such that $K_*(B) = 0$, and let $\pi : A \to M(SB \otimes \mathcal{K})$ be a strongly unitally absorbing representation into the stable multiplier algebra of the suspension of B. Then, for any finite subset X of A and any $\varepsilon > 0$ there exists a finite subset Y of A containing X and $\delta \leq \varepsilon$ such that the canonical forget control map*

$$KK_\delta(Y, SB) \to KK_\varepsilon(X, SB)$$

for the suspension of B is zero.

This is a key ingredient in our main results, but we hope it will prove to be useful in its own right. Note in particular that there are no assumptions on A other than that it is separable and unital[10].

There is a technical variation of Theorem 1.10 that applies to nuclear C^*-algebras, and that plays an important role in our arguments. The key point is one of order of quantifiers; condition (ii) from Theorem 1.10 starts with quantifiers of the form

$$\text{“}\forall B\ \forall \pi\ \forall X\ \forall \varepsilon\ \exists Y\ \exists \delta \ldots\text{”}.$$

If A is nuclear, the same statement is true with the order of quantifiers replaced with

$$\text{“}\forall \varepsilon\ \exists \delta\ \forall B\ \forall \pi\ \forall X\ \exists Y \ldots\text{”},$$

i.e., δ depends *only* on ε and not on any of the other choices involved. To establish this, we adapt an averaging argument due to Christensen, Sinclair, Smith, White, and Winter [17, Section 3], which is in turn based on Haagerup's theorem that nuclear C^*-algebras are amenable [33].

1.6 Strategy for the proof of the main theorem

Assume that A is a nuclear, unital C^*-algebra that decomposes with respect to the class of nuclear UCT C^*-algebras as in the statement of Theorem 1.2. Assume moreover that $K_*(B) = 0$. Thanks to Theorem 1.10 above, to establish the UCT for A it suffices to show that for any finite subset X of the unit ball A_1 of A, and any $\varepsilon > 0$ there exist $Y \supseteq X$ and $\delta \leq \varepsilon$ such that the canonical forget control map

$$KK^0_\delta(Y, SB) \to KK^0_\varepsilon(X, SB)$$

is zero.

[10] Unitality is not really necessary – we do not do it in this memoir, but similar techniques establish the result above for non-unital separable C^*-algebras, with appropriately reformulated controlled KK-groups.

Our approach to this is inspired directly by our earlier work with several collaborators; this includes the work on the Künneth formula of Oyono-Oyono and the second author [48], and separately by the first author [67]; the work of Guentner and the authors on the Baum–Connes conjecture for transformation groupoids with finite dynamical complexity [31]; and the work of Guentner, Tessera, and the second author on the stable Borel conjecture for groups of finite decomposition complexity [29]. These other papers all use controlled K-theory as opposed to KK-theory; the seminal result along these lines is the second author's work on the Novikov conjecture for groups with finite asymptotic dimension [71].

In the current context, we use decomposability and a Mayer–Vietoris argument. Let $\gamma > 0$ be a very small constant, which is in particular smaller than ε. Then, any suitably small[11] $\delta > 0$ will have the following property. Let h and C, D, and E be nuclear UCT algebras as in the definition of decomposability for the given set X and parameter δ. Let Y_C, Y_D and Y_E be finite subsets of the unit balls C_1, D_1, and E_1 respectively that contain $hX \cup \{h\}$, $(1-h)X \cup \{h\}$ and $h(1-h)X \cup \{h\}$ respectively up to δ-error, and so that Y_C and Y_D both contain Y_E up to δ-error. Let

$$Y = Y_C \cup Y_D \cup Y_E \cup X.$$

Then, one can construct a diagram[12] of the form

$$\begin{array}{ccc} & KK_\delta^0(Y, SB) \xrightarrow{\kappa_C \oplus \kappa_D} KK_{2\delta}^0(Y_C, SB) \oplus KK_{2\delta}^0(Y_D, SB)\ , \\ & \downarrow & \\ KK_\gamma(Y_E, S^2B) \xrightarrow{\partial} & KK_\varepsilon^0(X, SB) & \end{array} \tag{1.1}$$

where the vertical arrow is the canonical forget control map. This diagram has the "exactness" property that if $[p]$ goes to zero under the map

$$\kappa_C \oplus \kappa_D : KK_\delta^0(Y, B) \to KK_{2\delta}^0(Y_C, SB) \oplus KK_{2\delta}^0(Y_D, SB) \tag{1.2}$$

then the image of $[p]$ under the forget control map $KK_\delta^0(Y, SB) \to KK_\varepsilon^0(X, SB)$ is in the image of the map

$$\partial : KK_\gamma(Y_E, S^2B) \to KK_\varepsilon^0(X, SB). \tag{1.3}$$

[11] The size of γ depends linearly on ε and the size of δ depends linearly on γ; the constants involved are very large.

[12] The form of this diagram is not new; the basic idea is modeled on [29, Diagram (5.8)] from the work of the Guentner, Tessera, and the second author on the stable Borel conjecture for groups with finite decomposition complexity. See also [31, Proposition 7.6] from work of the Guentner and the authors in a more closely related context.

However, as $K_*(B) = 0$, if γ and δ are small enough, one can use Theorem 2.15 (in the stronger form for nuclear C^*-algebras) to choose Y_C, Y_D, and Y_E large enough so that the maps in lines (1.2) and (1.3) are zero. This completes the proof.

In the detailed exposition below we structure the proof to give it as "local" a flavor as possible, partly as we suspect that the ideas might be useful in other contexts. The two main "local"(ish) technical results are recorded as Propositions 7.1 and 7.2 below.

The argument above is directly inspired by the classical Mayer–Vietoris principle. Indeed, assume that C and D are nuclear *ideals* in A with intersection E, and such that

$$A = C + D.$$

Then, there is[13] an exact Mayer–Vietoris sequence

$$\cdots \to KK^0(E, SB) \to KK^0(A, B) \to KK^0(C, B) \oplus KK^0(D, B) \to \cdots .$$

In particular, if the groups at the left and right are zero, then the group in the middle is also zero. Our analysis of the diagram in line (1.1) is based on a concrete construction of this classical Mayer–Vietoris sequence that can be adapted to our controlled setting. The idea has its roots in algebraic K-theory, going back at least as far as [46, Chapter 2]. Having said this, there is significant work to be done adapting these classical ideas to the analytic superstructure that we built in [68], and the resulting formulas and arguments end up being quite different.

Remark 1.11. It would be very interesting to remove the nuclearity hypothesis from Theorem 1.2, or at least to replace it with something weaker such as exactness. Let us explain how nuclearity is used in the proof of Theorem 1.2, in the hope that some reader will see a way around it.

The first use of nuclearity is to show that any nuclear, unital C^*-algebra admits strongly unitally absorbing representations whose restriction to any nuclear, unital C^*-subalgebra is also strongly unitally absorbing; see Corollary 2.7 below. The proof of this is based on Kasparov's version of Voiculescu's theorem for Hilbert modules [40, Section 7]. It seems plausible from the discussion in Remark 2.8 below that some form of nuclearity is necessary for this to hold, but we do not know this.

The second place nuclearity is used is via an averaging argument due to Christensen, Sinclair, Smith, White, and Winter [17, Section 3]; this is applicable to nuclear C^*-algebras thanks to Haagerup's theorem that nuclear C^*-algebras are always amenable [33]. This lets us prove a stronger version of Theorem 1.10; see Corollary 2.22 below. We do not know if this result holds without nuclearity; see Remark 2.19 for a more detailed discussion.

[13] It is not in the literature as far as we can tell. For nuclear C^*-algebras, it can be derived from the usual long exact sequence in KK-theory using, for example, the argument of [69, Proposition 2.7.15].

1.7 Notation and conventions

For a subset S of a metric space X, $x \in X$ and $\varepsilon > 0$, we write "$x \in_\varepsilon S$" if there is $s \in S$ with $d(x, s) < \varepsilon$. For elements x, y of a metric space X, we write "$x \approx_\varepsilon y$" if $d(x, y) < \varepsilon$.

We write ℓ^2 for $\ell^2(\mathbb{N})$. Throughout, the letters A and B are reserved for *separable* C^*-algebras. The letter C will refer to a possibly non-separable C^*-algebra. The unit ball of C (or a more general normed space) is denoted by C_1, its unitization is C^+, its multiplier algebra is $M(C)$, its suspension is SC, and its n-fold suspension is $S^n C$. We write M_n or $M_n(\mathbb{C})$ for the $n \times n$ matrices, and $M_n(C)$ for the $n \times n$ matrices over a C^*-algebra C.

Our conventions on Hilbert modules follow those of Lance [45]. We will write $H_B := \ell^2 \otimes B$ for the standard Hilbert B-module, and $\mathcal{L}_B$, respectively $\mathcal{K}_B$, as shorthand for the C^*-algebra $\mathcal{L}(H_B)$ of adjointable operators on H_B, respectively the C^*-algebra $\mathcal{K}(H_B)$ of compact operators on H_B. We will typically identify $\mathcal{L}_B$ with the "diagonal subalgebra" $1_{M_n} \otimes \mathcal{L}_B$ of $M_n \otimes \mathcal{L}_B = M_n(\mathcal{L}_B)$. Thus, we might write "$[x, y]$" for the commutator of $x \in \mathcal{L}_B$ and $y \in M_n(\mathcal{L}_B)$, when it would be more strictly correct to write something like "$[1_{M_n} \otimes x, y]$".

The symbol "$\otimes$" always denotes a completed tensor product: either the external tensor product of Hilbert modules (see [45, Chapter 4] for background on this), or the minimal tensor product of C^*-algebras (see for example [10, Chapter 3]).

We will sometimes write 0_n and 1_n for the zero matrix and identity matrix of size n when this seems helpful to avoid confusion, although we will generally omit the subscripts to avoid clutter. If $n \leq m$, we will also use $1_n \in M_m(\mathbb{C})$ for the rank n projection with n ones in the top-left part of the diagonal and zeros elsewhere. Given an $n \times n$ matrix a and an $m \times m$ matrix b, $a \oplus b$ denotes the "block sum" $(n + m) \times (n + m)$ matrix defined by

$$a \oplus b := \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}.$$

Finally, $K_*(A) := K_0(A) \oplus K_1(A)$ denotes the graded K-theory group of a C^*-algebra, and $KK^*(A, B) := KK^0(A, B) \oplus KK^1(A, B)$ the graded KK-theory group. We will typically just write $KK(A, B)$ instead of $KK^0(A, B)$.

1.8 Outline of the paper

Chapter 2 gives our reformulation of the UCT in terms of a concrete vanishing condition for controlled KK-theory. The key ingredients for this are the Milnor sequence from [68, Theorem 1.1], and some ideas around the Mittag–Leffler condition from the theory of inverse limits (see for example [66, Section 3.5]). We also show that a stronger vanishing result holds for nuclear, UCT C^*-algebras using an averaging argument of Christensen, Sinclair, Smith, White, and Winter [17, Section 3]; the averaging argument is in turn based on Haagerup's theorem [33] that nuclearity implies amenability.

Chapter 3 discusses our controlled KK^0-groups. We introduced these in [68], but we need a technical variation here. This is essentially because in [68] we were setting up general theory, and for this it is easier to work with projections in a fixed C^*-algebra. In this memoir we are doing computations with concrete algebraic formulas, where it is more convenient to work with general idempotents, and to allow taking matrix algebras. We will, however, use both versions in this memoir, as we need to relate our work back here to the general theory of [68]. We also introduce controlled KK^1-groups in a concrete formulation using invertible operators; in our earlier work [68] we (implicitly) defined controlled KK^1-groups using suspensions, but here we also need the more concrete version.

Chapter 4 collects together some technical facts. These are all analogues for controlled KK-theory of well-known results from K-theory; for example, we prove "controlled versions" of the statements that homotopic idempotents are similar, and that similar idempotents are homotopic (up to increasing matrix sizes). Some arguments in this chapter are adapted from the work of Oyono-Oyono and the second author [47] on controlled K-theory.

Chapter 5 revisits the vanishing conditions of Chapter 2. Using the techniques of Chapter 4, we reformulate these results in the more flexible setting allowed by Chapter 3. This gives us the vanishing conditions that are the first main technical ingredient needed for Theorem 1.2.

Chapter 6 establishes the second main technical ingredient needed for Theorem 1.2. Here, we construct a "Mayer–Vietoris boundary map" for controlled KK-theory, and prove that it has an exactness property. The construction is an analogue of the usual index map of operator K-theory (see for example [54, Chapter 9]), although concrete formulas for the Mayer–Vietoris boundary map unfortunately seem to be missing from the C^*-algebra literature. The formulas we use are instead inspired by classical formulas from algebraic K-theory [46, Chapter 2], adapted to reflect our analytic setting.

Finally, in the main body of the paper, Chapter 7 puts everything together and gives the proofs of Theorem 1.2 and Theorem 1.4. We also include technical "local" vanishing results that we hope to elucidate the structure of the proof, and might be useful in other contexts.

The paper concludes with Appendix A, which gives examples of C^*-algebras with finite complexity. We first use a technique of Winter and Zacharias [70, Section 7] to show that the Cuntz algebras $\mathcal{O}_n$ with $2 \leq n < \infty$ have complexity rank one. We then use our joint work with Guentner on dynamic complexity [31] to show that ample, principal, étale groupoids with finite dynamical complexity and compact base space have C^*-algebras of finite complexity; we also get a similar result without the ampleness assumption if we allow C^*-algebras with finite complexity relative to subhomogeneous C^*-algebras.