FUSION OVER SUBLANGUAGES

ASSAF HASSON AND MARTIN HILS

Abstract. Generalising Hrushovski's fusion technique we construct the free fusion of two strongly minimal theories T_1, T_2 intersecting in a totally categorical sub-theory T_0. We show that if, e.g., T_0 is the theory of infinite vector spaces over a finite field then the fusion theory T_ω exists, is complete and ω-stable of rank ω. We give a detailed geometrical analysis of T_ω, proving that if both T_1, T_2 are 1-based then, T_ω can be collapsed into a strongly minimal theory, if some additional technical conditions hold—all trivially satisfied if T_0 is the theory of infinite vector spaces over a finite field $\mathbb{F}_q$.

§1. **Introduction.** In [17] the fusion technique for the construction of strongly minimal sets is introduced. In this paper Hrushovski notes that " ... it seems likely that ... if T_1, T_2 are strongly minimal sets with DMP and $T_1 \cap T_2$ is the theory of infinite vector spaces over a finite field then $T_1 \cup T_2$ has a strongly minimal expansion". In the present paper we suggest a first step towards proving this observation.

Given two strongly minimal theories T_1, T_2 (with DMP) the fusion construction can be roughly divided into three main steps:

(1) Define a class $\mathscr{C}$ of finite (or more generally finitely generated) $\mathscr{L}(T_1) \cup \mathscr{L}(T_2)$-structures and a notion $\leq$ of a *strong substructure*, such that $(\mathscr{C}, \leq)$ (is countable and) has the Joint Embedding Property and the Amalgamation Property.

(2) Show that the Fraïssé limit of $(\mathscr{C}, \leq)$ is a saturated model of a first order theory T. In the present context we will be mostly interested in fusions such that T is ω-stable of rank ω (we denote it T_ω), in which case T_ω will usually be coordinatised by strongly minimal locally finite types together with a unique (up to domination equivalence) generic type of rank ω.

(3) Collapse all the strongly minimal dimensions of T into finite sets, to obtain a new strongly minimal theory $\hat{T}$. This stage of the construction can be viewed as a construction and axiomatisation of smooth approximations to a (saturated) model of T_ω with respect to the requirement that in the approximation every strongly minimal set of T_ω has finitely many solutions.

In the original construction, T_1 and T_2 are fused to obtain 'as little interaction as possible' between them in the fusion theory T_ω (as in (2) above). In particular,

the language of the fusion is the disjoint union of $\mathscr{L}(T_1)$ with $\mathscr{L}(T_2)$ over equality. In this work we investigate the possibility of changing the language of the fusion, requiring the interpretation of T_1 and T_2 in a model of the fused theory T to share a non-trivial common reduct.

Towards this end we show that if $T_1 \cap T_2 = T_0$ is totally categorical then with a few slight technical changes using extensively the modularity of T_0 the first stage of Hrushovski's construction—the free amalgamation—can be worked out (Sections 2–4). We then show that if, e.g., T_0 is the theory of infinite vector spaces over a finite field, then the second stage of the construction can be carried out as well. In that case the Fraïssé limit is a saturated model of its first order theory, T_ω, an ω-stable theory of rank ω with a unique generic type (of rank ω in general). All the remaining regular types are strongly minimal. Furthermore, we show that every strongly minimal set definable in T_ω has precisely the structure inherited from $\mathscr{L}_0$ (possibly after an expansion by finitely many constants). Consequently, every strongly minimal set of T is locally modular and locally finite (Sections 5–6).

Generalising the last stage of Hrushovski's construction—the collapse of the infinite rank structure into a strongly minimal one—proved to be more delicate than the generalisation of the first two steps. The first difficulty is due to the emergence of affine sets—locally modular non-modular strongly minimal sets—in T_ω, with the resulting distinction (which does not exist for $\mathscr{L}_0 = \{=\}$) between orthogonality and almost orthogonality. This calls for a more delicate axiomatisation of the collapse, similar in spirit to the construction of envelopes in [8]. A second important difficulty which does not occur in the original case concerns definability of orthogonality in the theory T_ω. In the case of $\mathscr{L}_0 = \{=\}$, orthogonality of any two sets (definable without quantifiers in the fused theory) could be checked directly by comparing their respective $\mathscr{L}_1$- and $\mathscr{L}_2$-parts over their respective canonical bases. But if we discard this assumption on $\mathscr{L}_0$, this need no longer be true, and parameters may, a priori, have an important role to play in this analysis. A third important difficulty relates to a very delicate technical point—the analysis of the possible interactions of orthogonal strongly minimal types over a parameter set over which only one of them is based. Rather miraculously (and quite implicitly) Hrushovski, in his original fusion paper, shows that for $\mathscr{L}_0 = \{=\}$ this sort of interaction between orthogonal strongly minimal types can only be very limited. His arguments do not, however, translate to the case where $\mathscr{L}_0$ has a non-trivial geometry.

In the present work we were not able to give a complete proof of the collapse of T_ω in the general case, but we give a detailed 'site survey' of the fused theory T_ω, preparing the ground for an eventual proof of the collapse. Our work towards the general collapse can be partitioned in two:

- A thorough geometrical analysis of the fused theory T_ω resulting in a concrete presentation of all the strongly minimal dimensions (i.e., non-orthogonality classes of strongly minimal types) of T_ω, together with a good understanding of all the possible occurrences of non-orthogonality between strongly minimal sets in T_ω (Section 7).
- A suggested strategy for dealing with affine sets.

The above analysis already proves the collapse e.g., under the additional (very strong) assumption that T_1 and T_2 are both 1-based expansions of the theory of

an infinite vector space over some finite field $\mathbb{F}_q$ (the *abelian fusion context* which is treated in Section 8). The main merit of carrying out the construction in this special case lies not so much in the new structures it produces (as those will have a modular geometry) but more in that the solution suggested here for dealing with affine spaces is not restricted to the 1-based case, and can be translated almost unaltered to the general setting. The collapse in the 1-based case as well as the general strategy for a general collapse are perhaps best stated in terms of *envelopes*, as developed by the first author in [12].

At the request of the referee, two appendices are added to the main text.

In Appendix A we give a quick survey of the more general context in which we believe one should look in order to perform the collapse under less restrictive assumptions than those in Section 8. We point out how to use the results of Section 7 in order to solve the first of the problems discussed above (definability of orthogonality) and the precise statement of the third problem which we were to date unable to solve.

A problem closely related to the collapse of Hrushovski's fusion is that of collapsing Poizat's bicoloured fields. In the case of black-and-white fields—algebraically closed fields with a distinguished subset—due to its similarity with the collapse of the fusion over $\mathscr{L}_0 = \{=\}$, the collapse could be carried out. However, the red-and-white fields—algebraically closed fields with a distinguished subgroup of its additive group (see [21])—proved much harder to collapse, and indeed those fields have not yet been collapsed. In Appendix B we point out the geometric and structural similarities between the fusion over a common totally categorical and projective (i.e., modular non-trivial) theory and the red-and-white fields of positive characteristic. It follows that the same difficulties which arise when trying to collapse the fusion, also arise in the red-and-white field context, and we formulate our collapsing strategy for the former in a sufficiently general way so that it applies to the latter, too.

Finally, let us remark that large parts of this work (especially sections 2–5) were discovered independently by the two authors.

§**2. Preliminaries and notation.** First we indicate the context in which we work and fix some notation. We consider complete strongly minimal theories T_1 and T_2, in countable languages $\mathscr{L}_1$ and $\mathscr{L}_2$ respectively, having a common reduct $T_0 := T_1{\restriction}\mathscr{L}_0 = T_2{\restriction}\mathscr{L}_0$, where $\mathscr{L}_0 := \mathscr{L}_1 \cap \mathscr{L}_2$. Of course T_0 is strongly minimal, too.

To simplify the exposition, we assume that the theories T_i have quantifier elimination (in the corresponding, not necessarily relational languages $\mathscr{L}_i$, for $i = 0, 1, 2$). We denote by acl_i the algebraic closure in the sense of T_i. and require that $\mathrm{acl}_i(\emptyset)$ is infinite for $i \in \{1, 2\}$ (so equal to the prime model of T_i). This can always be achieved by Morleyising and adding some constants to the language, if necessary. Note that the latter assumption has weak elimination of imaginaries in T_i ($i = 1, 2$) as a consequence, and that these theories are not $\aleph_0$-categorical. In order to avoid trivial cases, let us further assume that the expansions $T_0 \subseteq T_i$ be *essential*, in the sense that in a monster model $\mathfrak{C}$ of T_i there exist $\mathscr{L}_i$-definable (over parameters) sets which are not $\mathscr{L}_0$-definable. Most of the time we suppose in addition:

(Geom) For both $i = 1, 2$, acl_i does not coincide with acl_0 (not even over some set of parameters), i.e., the expansions are not geometry preserving.

We always assume T_0 to be modular. Here are other restrictive (and crucial) conditions we often require:

- T_0 is $\aleph_0$-categorical.
- For $i = 1$ *or* $i = 2$ the expansion $T_0 \subseteq T_i$ *preserves multiplicities*, i.e., if $\overline{a} \in \mathrm{acl}_0(\overline{b})$, then $\mathrm{mult}_0(\overline{a}/\overline{b}) = \mathrm{mult}_i(\overline{a}/\overline{b})$.

From 3.4 on, T_0 is assumed to be $\aleph_0$-categorical throughout, as is condition (Geom). The preservation of multiplicities (or more precisely *good control*, a consequence thereof) is assumed from Section 5 on.

REMARK 2.1. (1) *If* $\mathrm{acl}_0 = \mathrm{dcl}_0$, *then any s.m. expansion of T_0 preserves multiplicities.*

(2) *If a strongly minimal expansion $T_0 \subseteq T_1$ is essential and preserves multiplicities, then this expansion is not geometry preserving.*

(3) *If T_0 is $\aleph_0$-categorical, the modularity assumption is no restriction, since T_0 is always locally modular in this case, so modular after adding some constant.*

PROOF. (1) is easy, and (2) is shown as follows. First, it is an (easy) exercise that if a s.m. expansion preserves multiplicities, then it preserves also Morley degrees, i.e., if $\varphi(\overline{x}, \overline{b})$ is definable in $\mathscr{L}_0$, then $\mathrm{MD}_{T_0}(\varphi(\overline{x}, \overline{b})) = \mathrm{MD}_{T_1}(\varphi(\overline{x}, \overline{b}))$ (note that for Morley rank, this equality always holds in s.m. expansions). Now choose an $\mathscr{L}_1$-definable (with parameters $\overline{b}$) set X which is not $\mathscr{L}_0$-definable and of minimal MRD_{T_1} (in lexicographic order) with this property. A generic solution $\overline{a}$ of X then has $\mathrm{d}_1(\overline{a}/\overline{b})) < \mathrm{d}_0(\overline{a}/\overline{b})$. If $\overline{a}_1$ is an $\mathscr{L}_1$-basis of $\overline{a}$ over $\overline{b}$, one then has $\overline{a} \subseteq \mathrm{acl}_1(\overline{b}\overline{a}_1)$ and $\overline{a} \not\subseteq \mathrm{acl}_0(\overline{b}\overline{a}_1)$, thus witnessing that the geometry is not preserved in the expansion.

(3) is a theorem of Zil'ber.   ⊣

We now proceed as in [17], although in notation and techniques, our exposition is closer to [20]. For $\mathscr{L} := \mathscr{L}_1 \cup \mathscr{L}_2$ define $\tilde{\mathscr{C}}$ as the class of all $\mathscr{L}$-structures which are models of T_1 and T_2. For $A \subseteq \mathfrak{M} \in \tilde{\mathscr{C}}$ define $\langle A \rangle$ as the smallest superset of A in M which is algebraically closed in the sense of T_1 and T_2. Equivalently, $\langle \cdot \rangle$ denotes the transitive closure of the operators acl_1 and acl_2. We call $\mathfrak{M}$ *finitely generated* (in the sense of $\langle \cdot \rangle$), if $M = \langle \overline{b} \rangle$ for some finite tuple $\overline{b} \in M$. Let $\mathscr{C}$ denote the class of all finitely generated structures in $\tilde{\mathscr{C}}$. For notational convenience we write AB for $A \cup B$, and $A \subseteq_\omega B$ means that A is a finite subset of B.

DEFINITION 2.2. Let $\mathfrak{M} \in \tilde{\mathscr{C}}$, $A \subseteq_\omega M$ and $B \subseteq M$.

(1) $\delta(A) := \mathrm{d}_1(A) + \mathrm{d}_2(A) - \mathrm{d}_0(A)$, the *predimension* of A, where d_i denotes the Morley rank in the sense of T_i.

(2) $\delta(A/B) := \mathrm{d}_1(A/B) + \mathrm{d}_2(A/B) - \mathrm{d}_0(A/B)$.

(3) $\tilde{\mathscr{C}}_0 := \{\mathfrak{M} \in \tilde{\mathscr{C}} \mid \delta(A) \geq 0 \,\forall A \subseteq_\omega M\}$, $\mathscr{C}_0 := \tilde{\mathscr{C}}_0 \cap \mathscr{C}$. Elements of $\tilde{\mathscr{C}}_0$ are called *fusions*.

(4) If $\mathfrak{M} \in \tilde{\mathscr{C}}_0$, $\mathrm{d}_{\mathfrak{M}}(A) := \min\{\delta(\tilde{A}) \mid A \subseteq \tilde{A} \subseteq_\omega M\}$, the *dimension* of A in $\mathfrak{M}$. Similarly, we define the relative dimension $\mathrm{d}_{\mathfrak{M}}(A/B) := \min\{\mathrm{d}_{\mathfrak{M}}(AB_0) - \mathrm{d}_{\mathfrak{M}}(B_0) \mid B_0 \subseteq_\omega B\}$.

(5) For $\mathfrak{M} \in \mathscr{C}_0$, we put $\delta(\mathfrak{M}) := \min\{\delta(A) \mid A \subseteq_\omega M \text{ and } \langle A \rangle = M\}$. Similarly, if $\mathfrak{M}$ is finitely generated over $A \subseteq M$, one can define $\delta(\mathfrak{M}/A)$.

(6) For $C \subseteq B$, both acl_0-closed subsets of a given structure in $\tilde{\mathscr{C}}$, one puts $C \leq B$ (C is *selfsufficient* or *strong* in B) if and only if for all finite tuples $\overline{b}$ from B one has $\delta(\overline{b}/C) \geq 0$.

Fusions will normally be denoted by K, L etc. whereas k, l etc. are mostly reserved for finitely generated fusions (i.e., structures in $\mathscr{C}_0$). If the ambient fusion K is clear from the context we just write d instead of d_K. Note that if B happens to be finite, then $\delta(A/B) = \delta(AB) - \delta(B)$.

It is convenient to extend the notion of a strong subset to arbitrary subsets $C \subseteq B$ of some fusion. We put $C \leq B$ if and only if $\text{acl}_0(C) \leq \text{acl}_0(B)$. This use is justified, since the definition does not depend on the particular embedding of B into an ambient fusion. By definition, we have $\tilde{\mathscr{C}}_0 = \{\mathfrak{M} \in \tilde{\mathscr{C}} \mid \emptyset \leq M\}$. Note also that $\delta(A) = \delta(\text{acl}_0(A))$.

REMARK 2.3. *The requirement of referring to acl_0-closed sets in the definition of $\leq$ is necessary since we want to obtain a transitive notion. If, for arbitrary $C \subseteq B$ one merely requires $\delta(\overline{b}/C)$ to be nonnegative for all finite tuples $\overline{b}$ in B, then already in the easiest cases (e.g., vector spaces) transitivity can fail.*

REMARK 2.4. *If T_0 is $\aleph_0$-categorical, $\tilde{\mathscr{C}}_0$ is an elementary class.*

PROOF. Let $\varphi_i(\overline{x})$ be $\mathscr{L}_i$-formulas, for $i = 1, 2$, with $\text{MR}(\varphi_i(\overline{x})) = m_i$. For such a pair we include the following (which is definable since T_0 is $\aleph_0$-categorical) as an axiom: $\forall \overline{x}\{[\varphi_1(\overline{x}) \wedge \varphi_2(\overline{x})] \rightarrow [d_0(\overline{x}) \leq m_1 + m_2]\}$. ⊣

DEFINITION 2.5. Let $K \in \tilde{\mathscr{C}}_0$ be a fusion and $A \subseteq K$. We say that *A controls K* if $\langle A \rangle = K$ and $A \leq K$. If $B \subseteq K$ is another subset, *A controls K over B* if AB controls K.

We observe that for finite A we have $\delta(\langle A \rangle) \leq \delta(A)$, and equality holds if and only if $\langle A \rangle$ is controlled by A.

§3. **Free amalgam and some δ-arithmetics.** We first gather some easy facts about the predimension δ which we will use constantly, sometimes without explicit reference. Comparing our lemma with [17, Lemma 1], one sees that we must switch to acl_0-closed sets most of the time, since modularity of T_0 has to be used. The proof is only a slight variation of the one given in [17]. Nevertheless, we prefer to give it in detail for the sake of completeness.

LEMMA 3.1. *Let $K \in \tilde{\mathscr{C}}$, and suppose that all the sets and tuples that appear are contained in K.*

(1) *(submodularity) $\delta(\overline{c}/\text{acl}_0(A\overline{c}) \cap \text{acl}_0(AB)) \geq \delta(\overline{c}/AB)$.*

(2) *(transitivity) If $A \leq B \leq C$, then $A \leq C$.*

(3) *(continuity) If $(A_i)_{i \in I}$ is a directed system of subsets of C such that $A_i \leq C$ for all i, then $\bigcup_{i \in I} A_i \leq C$.*

(4) *Let A_1, A_2 be acl_0-closed subsets of B, both selfsufficient in B. Then $A_1 \cap A_2 \leq B$.*

Now suppose that $K \in \tilde{\mathscr{C}}_0$. Computing d with respect to this K, we then have:

(5) *For any $A \subseteq K$ there is a minimal superset $\overline{A} \supseteq A$ (depending on K) such that $\overline{A} \leq K$ and $\overline{A}$ is acl_0-closed. If A is finite, $\overline{A}$ is the acl_0-closure of a finite set (so in particular is finite if T_0 is $\aleph_0$-categorical). Moreover, one then has $d(A) = \delta(\overline{A})$.*

(6) *Let $A \subseteq B \subseteq C$ be finite. Then we have $d(C/A) = d(C/B) + d(B/A)$, $d(B/A) \leq d(C/A)$ and $d(C/A) \geq d(C/B)$.*

(7) *$d(a/B) \in \{0, 1\}$ for any singleton a, and the geometric closure operator $\text{cl}^{geom}(B) := \{a \in K \mid d(a/B) = 0\}$ defines a pregeometry, i.e., is monotone, transitive and satisfies the Steinitz exchange property.*

PROOF. (1) follows from the fact that T_0 is modular and T_1, T_2 are submodular. To show (2), by our definition of selfsufficiency we can clearly assume that A, B and C are acl_0-closed. Now, $\delta(\overline{c}/A) = \delta(\overline{c}/\mathrm{acl}_0(A\overline{c}) \cap B) + \delta(\mathrm{acl}_0(A\overline{c}) \cap B/A)$ for any $\overline{c} \in C$. The term to the left is at least $\delta(\overline{c}/B)$ by (1), and so both terms are nonnegative by hypothesis. (3) is easy.

In (4), by transitivity it is sufficient to prove that $A_1 \cap A_2 \leq A_1$, which is true by submodularity. To prove (5), we first reduce it to the case of a finite set A. Note that by (4) the operator $A \mapsto \overline{A}$ is well-defined on finite sets and monotone where it is defined. By (3) it is easy to see that $\overline{A} = \bigcup_{A_0 \subseteq_\omega A} \overline{A_0}$ has the desired property. If A is finite, there are (finite) supersets of A which are selfsufficient in K, since $K \in \tilde{\mathscr{C}}_0$. Thus, (5) follows from (4).

CLAIM. If A and B are finite, then $\mathrm{d}(B/A) = \mathrm{d}(AB) - \mathrm{d}(A)$.

The claim gives (6) and the fact that the minimum occurring in the definition of relative dimension is a limit. To prove the claim we take $A' \subseteq A$ and compute:

$$\mathrm{d}(AB) \leq \delta(\overline{A'B} \cup \overline{A}) \leq \delta(\overline{A'B}) + \delta(\overline{A}) - \delta(\overline{A'B} \cap \overline{A})$$
$$\leq \delta(\overline{A'B}) + \delta(\overline{A}) - \mathrm{d}(A') = \mathrm{d}(A'B) + \mathrm{d}(A) - \mathrm{d}(A').$$

The inequalities follow from (in this order): $AB \subseteq \overline{A'B} \cup \overline{A}$, submodularity, $A' \subseteq \overline{A'B} \cap \overline{A}$ and (5).

We finally prove (7). It is immediate to reduce to finite B, since cl^{geom} is obviously continuous. Monotonicity and transitivity of cl^{geom} follow from (6). Since $\mathrm{d}(Ba) \leq \delta(\overline{B}a) \leq \delta(\overline{B}) + \delta(a/\overline{B}) \leq \mathrm{d}(B) + 1$, one has $\mathrm{d}(a/B) \in \{0, 1\}$. By (6), $\mathrm{d}(a/B) + \mathrm{d}(c/Ba) = \mathrm{d}(ac/B) = \mathrm{d}(c/B) + \mathrm{d}(a/Bc)$. Now, if $\mathrm{d}(a/B) = 1$ and $\mathrm{d}(c/Ba) = 0$, necessarily $\mathrm{d}(a/Bc) = 0$. This gives Steinitz exchange. ⊣

We now introduce yet another closure operator for which will be reserved the term selfsufficient closure.

DEFINITION 3.2. Let K be a fusion and $A \subseteq K$. Then we put $\mathrm{cl}_K(A) := \langle \overline{A} \rangle$ (for $\overline{A}$ as in (5) above), the *selfsufficient closure* of A (in K).

The selfsufficient closure of A equals the least subfusion of K which contains A and which is strong in K. Most of the time we write cl instead of cl_K. Note that $\mathrm{acl}_0(A) \subseteq \overline{A} \subseteq \mathrm{cl}(A) \subseteq \mathrm{cl}^{geom}(A)$.

Observe that since cl^{geom} gives rise to a pregeometry, there is a notion of dimension attached to it. Clearly, this dimension equals the one we already defined (on finite sets), and from now on, $\mathrm{d}(A/B)$ will denote this dimension for arbitrary A and B.

REMARK 3.3. *If $K \leq L$ are fusions and $A \subseteq K$ then $\mathrm{d}_K(A) = \mathrm{d}_L(A)$, $\mathrm{cl}_K(A) = \mathrm{cl}_L(A)$, and $\overline{A}$ calculated in K or in L amounts to the same.*

PROOF. By transitivity of $\leq$, $\overline{A}$ calculated in K and L amounts to the same. Moreover, we already know by the preceding lemma that for any finite A, $\mathrm{d}_M(A) = \delta(\overline{A})$, if $\overline{A}$ is calculated in M. Thus, using the proof of (5) of Lemma 3.1 $\mathrm{d}_K(A) = \delta(\overline{A}) = \mathrm{d}_L(A)$ and $\mathrm{cl}_K(A) = \langle \overline{A} \rangle = \mathrm{cl}_L(A)$ follows too. ⊣

CONVENTION 3.4. From now on it will be assumed throughout the rest of the paper that T_0 is $\aleph_0$-categorical and that the expansions $T_0 \subseteq T_i$ are not geometry preserving (assumption (Geom)).

DEFINITION 3.5. The triple (T_0, T_1, T_2) is said to have *good control* if whenever $K \in \tilde{\mathscr{C}}$ is controlled by A, then (the $\mathscr{L}$-isomorphism type of) K is determined by $\mathrm{qftp}_{\mathscr{L}}(A)$.

LEMMA 3.6. *If one of the two expansions $T_0 \subseteq T_i$ preserves multiplicities, then (T_0, T_1, T_2) has good control. In particular, if $\mathrm{dcl}_0 = \mathrm{acl}_0$, any context (T_0, T_1, T_2) has good control.*

PROOF. W.l.o.g. we can assume that $T_0 \subseteq T_1$ preserves multiplicities. By abuse of natural language, in the course of this proof maps which preserve the quantifier free $\mathscr{L}$-type will be called $\mathscr{L}$-isomorphisms. Let $\imath \colon A \cong A'$ be an $\mathscr{L}$-isomorphism, with A controlling $K \in \tilde{\mathscr{C}}$, A' controlling K'. By induction, it suffices to show that $\imath$ extends to an $\mathscr{L}$-isomorphism $\tilde{\imath} \colon \mathrm{acl}_i(A) \cong \mathrm{acl}_i(A')$, for $i = 1, 2$.

First, we show that $\mathrm{acl}_0(A) \cong_{\mathscr{L}} \mathrm{acl}_0(A')$. Just choose any $\mathscr{L}_2$-elementary map $\tilde{\imath} \colon \mathrm{acl}_0(A) \cong \mathrm{acl}_0(A')$ extending $\imath$ (T_2 has QE hence $\imath$ extends to an $\mathscr{L}_2$-isomorphism $\mathrm{acl}_2(A) \cong \mathrm{acl}_2(A')$ and so in particular to $\mathrm{acl}_0(A)$). Trivially, $\tilde{\imath}$ is $\mathscr{L}_0$-elementary, too. As $T_0 \subseteq T_1$ preserves multiplicities, any $\sigma \in \mathrm{Aut}_{\mathscr{L}_0}(\mathrm{acl}_0(A)/A)$ is $\mathscr{L}_1$-elementary. Thus, $\tilde{\imath}$ is an $\mathscr{L}_1$-elementary map, too and so an $\mathscr{L}$-isomorphism, so we may assume that $\tilde{\imath} = \imath$.

We now treat the case of acl_1. Let $\imath_1 \colon \mathrm{acl}_1(A) \cong \mathrm{acl}_1(A')$ be an auxiliary $\mathscr{L}_1$-isomorphism extending $\imath$, and choose B, an $\mathscr{L}_0$-basis of $\mathrm{acl}_1(A)$ over A, and put $B' := \imath(B) \subseteq \mathrm{acl}_1(A')$. Since $A \le \mathrm{acl}_1(A)$, for any $\overline{b} \subseteq_\omega B$ one has $\mathrm{d}_2(\overline{b}/A) = \mathrm{d}_0(\overline{b}/A)$, so B is $\mathscr{L}_2$-independent over A (similarly B' over A'). Thus, $\imath\!\restriction_{AB} \colon AB \cong A'B'$ is an $\mathscr{L}$-isomorphism extending $\imath$.

The fact that $\imath$ extends to an $\mathscr{L}$-isomorphism $\mathrm{acl}_2(A) \cong \mathrm{acl}_2(A')$ follows in a similar way (although this case is easier since any $\mathscr{L}_2$-isomorphism extending $\imath$ would do), and we conclude by induction.

Finally, note that if $\mathrm{dcl}_0 = \mathrm{acl}_0$, then any s.m. expansion T_i of T_0 preserves multiplicities. $\dashv$

COROLLARY 3.7. *Assuming good control, the class $\mathscr{C}_0$ is countable, as is (the number of isomorphism classes in) $\{l \in \mathscr{C}_0 \mid k \le l\}$, for every $k \in \mathscr{C}_0$.*

PROOF. First, note that every $k \in \mathscr{C}_0$ is controlled by a finite tuple $\overline{b}$. Just take a finite tuple $\overline{b}_0$ generating k and choose $\overline{b} \supseteq \overline{b}_0$ with minimal predimension. Such a $\overline{b}$ controls k. In T_1 and in T_2 there exist only countably many types of finite tuples over $\emptyset$, whence the countability of $\mathscr{C}_0$ follows from Lemma 3.6. Before proceeding to prove the second part of the corollary we introduce a useful definition:

DEFINITION 3.8. Let $k, l \in \tilde{\mathscr{C}}_0$ with $k \le l$. A pair of finite tuples $(\overline{a}/\overline{b})$ is called a *controlling pair* for the extension $k \le l$ if: $\overline{b}$ controls k and $\overline{a}\overline{b}$ controls l (so $\overline{a}$ controls l over k, too) and $\mathrm{tp}_{\mathscr{L}_i}(\overline{a}/k)$ is a non-forking extension of $\mathrm{tp}_{\mathscr{L}_i}(\overline{a}/\overline{b})$ over k. A controlling pair $(\overline{a}/\overline{b})$ is *strong*, if $\mathrm{tp}_{\mathscr{L}_i}(\overline{a}/\overline{b})$ is stationary for $i = 1, 2$.

Now fix $k \in \mathscr{C}_0$. Note that, since $k \models T_i$ for $i \in \{1, 2\}$, for any $k \le l \in \tilde{\mathscr{C}}_0$ strong controlling pairs exist. If $(\overline{a}/\overline{b})$ is such a pair, $\mathrm{qftp}_{\mathscr{L}}(\overline{a}\overline{b})$ completely determines $k \le l$, in the sense that fixing $a'b' \models \mathrm{qftp}_{\mathscr{L}}(\overline{a}\overline{b})$, any isomorphism of $k' := \langle \overline{b}' \rangle$ with k extends to an isomorphism of $l' := \langle \overline{a}'\overline{b}' \rangle$ with l. Hence, $\{k \le l \mid l \in \mathscr{C}_0\}$ is countable. $\dashv$

The following example appears in [17] and shows that one cannot hope to find strongly minimal expansions in an arbitrary context of a fusion over a sublanguage.

It even provides an example where no $M \in \tilde{\mathscr{C}}$ can be superstable with a unique type of maximal rank.

EXAMPLE 3.9. Let $G_1 := \mathbb{Z}/4$, $G_2 := \mathbb{Z}/2 \times \mathbb{Z}/2$ and consider $T_i :=$ theory of the free action of G_i on an infinite set in the language $\mathscr{L}_i$ with function symbols for elements of G_i. $T_0 :=$ theory of an equivalence relation E, all classes of which contain 4 elements, where $\mathscr{L}_0 = \{E\}$. T_0 is obtained as a reduct of T_i by interpreting E as the orbits of G_i. Now suppose $M \models T_1 \cup T_2$ is strongly minimal (or more generally has a unique type of maximal rank $< \infty$). Let $a \in G_1$ be an element of order 4. For generic x we have $a \cdot x = c \cdot x$ for some id $\neq c \in G_2$ with $c^2 = $ id. Since $a \cdot x$ is generic, too, this leads to the following contradiction: $x \neq a^2 \cdot x = c \cdot (a \cdot x) = c^2 \cdot x = x$.

What makes this example work is the lack of good control in the fusion context (T_0, T_1, T_2). Strictly speaking, (T_0, T_1, T_2) does not fit in our framework (hypothesis (Geom) is not satisfied, and $\mathrm{acl}_i(\emptyset)$ is finite). This problem disappears once we work with larger groups $G_i' \supsetneq G_i$ and name some constants.

Construction of the free amalgam in $\tilde{\mathscr{C}}$. Let $K \subseteq L, M$ be three elements of $\tilde{\mathscr{C}}$. To such a triple one wants to associate (canonically, if we have good control) a *free amalgam* $N := M \otimes_K L$. It turns out that the right candidate is an $N \in \tilde{\mathscr{C}}$, $N \supseteq M, L$ such that:

(α) $M \underset{K}{\overset{i}{\downarrow}} L$, for $i = 1, 2$ (where $\overset{i}{\downarrow}$ means independence in the sense of T_i) and
(β) ML controls N.

If we assume good control, (α) and (β) together determine N completely, as $K \models T_i$ and so $\mathscr{L}_i$-types are stationary over K. Recall that for $i \in \{0, 1, 2\}$ the free amalgam $\otimes^i$ is defined (just define $M \otimes_K^i L$ as the T_i-prime model over ML where $M \underset{K}{\overset{i}{\downarrow}} L$). To see that free amalgams always exist (in $\tilde{\mathscr{C}}$) it will suffice to show the following lemma (apply it for $p_i := \mathrm{tp}_{\mathscr{L}_i}(ML/K)$, where $M \underset{K}{\overset{i}{\downarrow}} L$).

LEMMA 3.10. *Let $K \in \tilde{\mathscr{C}}$. Suppose that for $i = 1, 2$ complete (infinitary) $\mathscr{L}_i$-types $p_i(x_I)$ over K are given, such that $p_0 := p_1 {\restriction}_{\mathscr{L}_0} = p_2 {\restriction}_{\mathscr{L}_0}$. Then there is $L \in \tilde{\mathscr{C}}$ and $A \subseteq L$, $A \models p_1 \cup p_2$ such that L is controlled by A over K.*

PROOF. In this proof, for convenience we suppose that $\mathscr{L}$ is relational. By Robinson's consistency lemma, there is $A \models p_1 \cup p_2$. Clearly, $KA \models T_1^\forall \cup T_2^\forall$. We now show that for any $B \models T_1^\forall \cup T_2^\forall$ there is an $L \in \tilde{\mathscr{C}}$ controlled by B.

By transitivity and continuity of selfsufficiency, using Zorn's lemma, we can find a set $\tilde{B}$ with $B \leq \tilde{B}$, $\tilde{B} \models T_1^\forall \cup T_2^\forall$ and $\tilde{B} \subseteq \langle B \rangle$ such that $\tilde{B}$ is maximal with these properties. It is easy to see that $\tilde{B}$ has to be acl_0-closed. Now suppose $\tilde{B}$ were not acl_1-closed, say. Choose $B' = \mathrm{acl}_0(\tilde{B}b') \supseteq \tilde{B}$, where $B' \models T_1^\forall$ and $b' \in \mathrm{acl}_1(\tilde{B}) \setminus \tilde{B}$ is a singleton. We make B' into a $T_2^\forall$-model (over $\tilde{B}$) in such a way that b' is $\mathscr{L}_2$-generic over $\tilde{B}$. Since then b' is $\mathscr{L}_0$-generic (over $\tilde{B}$) in both $\mathscr{L}_1$ and $\mathscr{L}_2$, the two constructions coincide in $\mathscr{L}_0$. Over $\tilde{B}$, every tuple $\overline{a}$ from B' which does not entirely lie in $\tilde{B}$ is interalgebraic in $\mathscr{L}_0$ with b', and thus $\delta(\overline{a}/\tilde{B}) = \delta(b'/\tilde{B}) = 0 + 1 - 1 = 0$. So $\tilde{B} \leq B'$ follows. This contradicts the maximality of $\tilde{B}$. $\dashv$

NOTATION. Let K^* be some big fusion. For $B \subseteq A, C \subseteq K^*$ we put $A \underset{B}{\overset{\mathrm{d}}{\downarrow}} C$ if $\mathrm{d}(A/B) = \mathrm{d}(A/BC)$, i.e., if A and C are d-independent over B.

LEMMA 3.11. *Let K_1, K_2 be two fusions strongly embedded in K^*. Put $K_0 := K_1 \cap K_2$ and $L := \langle K_1 K_2 \rangle$. The following are equivalent*:

(1) $K_1 \underset{K_0}{\overset{\mathrm{d}}{\downarrow}} K_2$.

(2) L *is isomorphic to a free amalgam $K_1 \otimes_{K_0} K_2$ and L is strong in K^*.*

(3) $K_1 \underset{K_0}{\overset{i}{\downarrow}} K_2$ $(i = 1, 2)$ *and $K_1 K_2 \leq K^*$.*

PROOF. $(2) \iff (3)$ is easy. In order to show the remaining part, we first observe that both (1) and (2) are finitary statements over K_0 in the sense that their content is true if and only if it is true for all strong subfusions of K_i finitely generated over K_0. Thus, we may assume that K_i/K_0 is finitely generated for $i = 1, 2$. We now show $(3) \Rightarrow (1)$. From the hypotheses in (3) we deduce

$$\mathrm{d}(K_1/K_2) = \delta(K_1/K_2) = \delta(K_1/K_0) = \mathrm{d}(K_1/K_0),$$

whence (1) follows (as $K_1 K_2 \leq K^*$ implying $d(K_1 K_2/K_0) = \delta(K_1 K_2/K_0)$).

For the other direction, by $K_1 \underset{K_0}{\overset{\mathrm{d}}{\downarrow}} K_2$, in the following we have equality throughout

$$\mathrm{d}(K_1/K_2) \leq \delta(K_1/K_2) \leq \delta(K_1/K_0)\mathrm{d}(K_1/K_0).$$

Here, the second inequality is just submodularity. Going from left to right, this means that $K_1 K_2 \leq K^*$ and $K_1 \underset{K_0}{\overset{i}{\downarrow}} K_2$ for $i = 1, 2$ (since $K_1 \underset{K_0}{\overset{0}{\downarrow}} K_2$ by modularity of T_0). ⊣

FACT 3.12. *Let T' be strongly minimal and modular. Then the lattice of* acl-*closed sets is modular. This means the following*: *If C and $A \subseteq B$ are* acl-*closed sets, then* $\mathrm{acl}(A(B \cap C)) = B \cap \mathrm{acl}(AC)$.

PROOF. Of course this fact is well-known since the 70's and the source of the term 'modular'. Nevertheless, we give a proof.

It is clear how to reduce to $\mathrm{d}'(BC) < \infty$, so let us suppose this (here, d' denotes the dimension in the sense of T'). Since $\mathrm{acl}(A(B \cap C)) \subseteq B \cap \mathrm{acl}(AC)$ is obvious, it suffices to show that $\mathrm{d}'(B \cap C/A) = \mathrm{d}'(B \cap \mathrm{acl}(AC)/A)$ holds. By modularity of T' we have $C \underset{A \cap C}{\downarrow} A$ and $C \underset{B \cap C}{\downarrow} B \cap \mathrm{acl}(AC)$. Thus,

$$\mathrm{d}'(B \cap C/A) = \mathrm{d}'(B \cap C/A \cap C) = \mathrm{d}'(C/A \cap C) - \mathrm{d}'(C/B \cap C)$$
$$= \mathrm{d}'(C/A) - \mathrm{d}'(C/B \cap \mathrm{acl}(AC)) = \mathrm{d}'(B \cap \mathrm{acl}(AC)/A). \quad ⊣$$

LEMMA 3.13 (asymmetric amalgamation). *Let $K, L, M \in \tilde{\mathscr{C}}_0$, $K \leq L$, $K \subseteq M$. Then M is selfsufficient in any free amalgam $N := L \otimes_K M$, and N is in $\tilde{\mathscr{C}}_0$.*

PROOF. The construction of a (or the) free amalgam gives $LM \leq N$. Thus, in order to prove $M \leq N$ it suffices to show $M \leq LM$. We recall that this requires $\delta(A/M) \geq 0$ for every $A \subseteq_\omega \mathrm{acl}_0(LM)$. So suppose we are given such a set A. Put $B := \mathrm{acl}_0(AM)$. By Fact 3.12, $B = \mathrm{acl}_0((B \cap L)M)$. Thus, A is inter-$\mathscr{L}_0$-algebraic over M with some $A' \subseteq_\omega L$. So $\delta(A/M) = \delta(A'/M)$ which is equal to $\delta(A'/K)$, since $M \underset{K}{\overset{i}{\downarrow}} L$ for $i \in \{0, 1, 2\}$ (by construction). But $\delta(A'/K)$ is nonnegative by $K \leq L$. Finally, $N \in \tilde{\mathscr{C}}_0$, since $\emptyset \leq M$ and $M \leq N$. ⊣

DEFINITION 3.14. $\mathscr{M} \in \tilde{\mathscr{C}}_0$ is *rich* if for every $k \leq l$ in $\mathscr{C}_0$, $k \leq \mathscr{M}$ there exists a selfsufficient k-embedding of l in $\mathscr{M}$.

PROPOSITION 3.15. *Assuming good control, there is a rich countable $\mathcal{M} \in \tilde{\mathscr{C}}_0$, and it is unique up to isomorphism (we call it the* generic *model). Moreover, every two rich structures in $\tilde{\mathscr{C}}_0$ are $\mathscr{L}_{\infty,\omega}$-equivalent.*

PROOF. The class $(\mathscr{C}_0, \leq)$ has the amalgamation property by Lemma 3.13. As $\langle \emptyset \rangle \in \mathscr{C}_0$ embeds strongly in every member of this class, joint embedding is guaranteed, too. The necessary countability results are shown in Corollary 3.7, so that, by the usual Fraïssé method, a countable rich structure in $\tilde{\mathscr{C}}_0$ can be constructed. Finally, the $\mathscr{L}_{\infty,\omega}$-equivalence of two rich structures is true almost by definition, and so uniqueness of the generic model follows, since it is rich and countable. ⊣

§4. Decomposition of finitely generated extensions.

DEFINITION 4.1. Let $K \leq L$ be an extension in $\tilde{\mathscr{C}}_0$ (we suppose that $L \leq K^*$ for some ambient K^*). The extension is called

- *finitely generated* if $L = \mathrm{cl}(K\overline{a})$ for some finite $\overline{a} \in L$,
- *generic* if $L = \langle Ka \rangle$ for a singleton a with $\mathrm{d}(a/K) = 1$,
- *parasitic* if it is finitely generated and $\delta(L/K) = 0$,
- *primitive* if it is parasitic, proper and there is no $K' \in \tilde{\mathscr{C}}_0$ such that $K \subsetneqq K' \subsetneqq L$ and $K' \leq L$.

LEMMA 4.2. *Let $K \leq K' \leq K^*$, and $K \leq L \leq K^*$ with L primitive over K. Then either $L \subseteq K'$ or $L' := \langle LK' \rangle$ equals a free amalgam $K' \otimes_K L$ (and $L' \leq K^*$).*

PROOF. Clearly, $L \underset{L \cap K'}{\overset{\mathrm{d}}{\downarrow}} K'$, Now, if $L \cap K' \supsetneqq K$, then $L \subseteq K'$, since L/K is primitive. Else, we conclude by Lemma 3.11. ⊣

Technically it is convenient to consider a concept of primitiveness for extensions of acl_0-closed subsets of some fusion. In order to avoid confusions the term *primitive* will only be used for extensions of fusions.

DEFINITION 4.3. Let $B \subseteq A \subseteq K \in \tilde{\mathscr{C}}_0$. The extension A/B is called *prime*, if both A and B are acl_0-closed, $\mathrm{d}_0(A/B)$ is finite and ≥ 2, $\delta(A/B) = 0$ and $\delta(A'/B) > 0$ for every acl_0-closed A' with $B \subsetneqq A' \subsetneqq A$. $\mathrm{d}_0(A/B)$ is called the *length* of the extension.

REMARK. *By our definition we exclude the "prime extensions of length 1", corresponding to $A := \mathrm{acl}_0(B\alpha)$, where α is in exactly one of $\mathrm{acl}_i(B)$, $i = 1, 2$.*

We recall that for $A \subseteq K \in \tilde{\mathscr{C}}_0$ we defined $\overline{A}$ as well as $\mathrm{cl}(A)$, the selfsufficient closure of A (see 3.1(5) and 3.2, respectively).

LEMMA 4.4. *Let A/B be a prime extension of length n (inside $K \in \tilde{\mathscr{C}}_0$) and let $B \subseteq B' \subseteq K$ with B' acl_0-closed. Put $A' := \mathrm{acl}_0(AB')$. One then has:*

(1) *If $B' \underset{B}{\overset{i}{\downarrow}} A$ for $i = 1, 2$, then A'/B' is prime of length n. In particular $B' \underset{B}{\overset{0}{\downarrow}} A$.*

(2) *$\delta(A/B') = \delta(A'/B') \leq 0$, where equality holds if and only if either $A \subseteq B'$ or A'/B' is prime (of length n).*

(3) *If $B' = \overline{B'}$, then either $A \subseteq B'$ or A'/B' is prime of length n.*

PROOF. For any singleton $a \in A \setminus B$ one has $\mathrm{d}_1(a/B) = \mathrm{d}_2(a/B) = 1$ by the fact that A/B is prime and $n \geq 2$ (by the definition of a prime extension). So, if $B' \underset{B}{\overset{1}{\downarrow}} A$, one has $1 = \mathrm{d}_1(a/B') \leq \mathrm{d}_0(a/B')$, thus $A \cap B' = B$. From modularity of T_0 one deduces $B' \underset{B}{\overset{0}{\downarrow}} A$. We can now prove (1) using Fact 3.12.

In (2), submodularity gives $\delta(A'/B') = \delta(A/B') \leq \delta(A/A \cap B')$. By definition of a prime extension, $\delta(A/A \cap B')$ is negative unless $A \cap B'$ is equal to A or B, in which case it is 0. Thus, if equality holds, either $A \subseteq B'$ or $A \underset{B}{\overset{0}{\downarrow}} B'$ (by modularity of T_0), so $d_0(A/B') = d_0(A/B) = n$. In this latter case, one deduces that for $i = 1, 2$ $d_i(A/B') = d_i(A/B)$, too, whence $B' \underset{B}{\overset{i}{\downarrow}} A$. The rest of (2) follows by (1). Note that the hypothesis of (3) forces $\delta(A/B') \geq 0$, and so (2) applies.   ⊣

For convenience we introduce some notation which will be used in several proofs proceeding by induction. The operators $\langle \cdot \rangle_1^n, \langle \cdot \rangle_2^n$ and $\langle \cdot \rangle^n$ ($n \in \mathbb{N}$) give rise to different filtrations of $\langle \cdot \rangle$.

DEFINITION 4.5. For $X \subseteq K$ and $n \in \mathbb{N}$ we now define $\langle X \rangle_1^n$, $\langle X \rangle_2^n$ and $\langle X \rangle^n$. First, put $\langle X \rangle_1^0 := \langle X \rangle_2^0 := \langle X \rangle^0 := \text{acl}_0(X)$. Inductively set $\langle X \rangle_i^{m+1} := \text{acl}_i(\langle X \rangle^m)$, and $\langle X \rangle^{m+1} := \text{acl}_0(\langle X \rangle_1^{m+1} \cup \langle X \rangle_2^{m+1})$.

LEMMA 4.6. *Let $B \subseteq A \subseteq K \in \tilde{\mathscr{C}}_0$ with A/B prime. If $A \subseteq \text{cl}(D)$ for some $D \supseteq B$, then $A \subseteq \overline{D}$.*

PROOF. Using Lemma 4.4(3) we easily reduce to the case where $B = \overline{B} \subseteq D = \overline{D}$. Now, let m be minimal with $A \subseteq \langle D \rangle_i^m$ for some i. We may assume that this is the case for $i = 1$. If $m \geq 1$, Lemma 4.4 forces A'/B' to be prime, where $B' := \langle D \rangle_2^{m-1}$ and $A' := \text{acl}_0(AB')$, as $B' = \overline{B'}$. By primality, no $a \in A' \setminus B'$ is $\mathscr{L}_i$-algebraic over B', for $i = 1$ or 2. This contradicts the fact that (by definition) $\langle D \rangle_1^m = \text{acl}_1(B')$. Thus, $m = 0$ and the lemma is proved.   ⊣

PROPOSITION 4.7. *For any parasitic extension $K \leq L$ the following holds:*

(1) *There is a finite decomposition $K = K_0 \leq K_1 \leq \cdots \leq K_n = L$ such that $K_i \leq K_{i+1}$ is primitive.*

(2) *The decomposition is essentially unique, i.e., for any other decomposition $K = K_0' \leq K_1' \leq \cdots \leq K_{n'}' = L$ into primitive extensions we have $n = n'$. Moreover, assuming good control and setting $L_i := L \otimes_{K_{i-1}} K_i$ as well as $L_i' := L \otimes_{K_{i-1}'} K_i'$ (for $1 \leq i \leq n$), then, up to permutation, $(L_1, \ldots, L_n)$ is isomorphic to $(L_1', \ldots, L_n')$.*

PROOF. In order to prove the existence part (1), we choose first a finite tuple $\overline{a}$ controlling L over K. Now, choose A_1 with $K \subsetneq A_1 \subseteq \text{acl}_0(K\overline{a})$ and $A_1 = \overline{A_1}$ such that A_1 is minimal with these properties. Since $d_0(\overline{a}/K)$ is finite, such an A_1 always exists. By minimality it is easy to see that A_1/K is prime. Put $K_1 := \langle A_1 \rangle$.

It is easy to check that L is controlled by $\overline{a}$ over K_1. Since $d_0(\overline{a}/K) > d_0(\overline{a}/K_1)$, part (1) of the proposition follows by induction on $d_0(\overline{a}/K)$ once we prove:

LEMMA 4.8. *Let $K \leq A_1 \leq K^*$ with $K \in \tilde{\mathscr{C}}_0$ and A_1/K prime. Then, $K_1 := \langle A_1 \rangle$ is a primitive extension of K.*

PROOF OF THE LEMMA. We have to show that $\text{cl}(K\alpha) = K_1$ for every $\alpha \in K_1 \setminus K$. For $\alpha \in A_1$ this is the case by primality of A_1/K, since then $\overline{K\alpha} = A_1$. If $\alpha \in \text{acl}_1(A_1) \setminus A_1$ we have $\overline{K\alpha} \subseteq \text{acl}_0(A_1\alpha) = \overline{A_1\alpha}$. Now, by modularity of T_0,

$$(*) \quad d_0(A_1/K) + 1 = d_0(A_1\alpha/K) = d_0(A_1/K) + d_0(\overline{K\alpha}/K) - d_0(\overline{K\alpha} \cap A_1/K).$$

Since $0 = d(\alpha/K) < \delta(\alpha/K) = 1$, the dimension $d_0(\overline{K\alpha}/K)$ is at least 2, from which we deduce by $(*)$ that there is $\alpha' \in \overline{K\alpha} \cap (A_1 \setminus K)$. So, by the first case $\text{cl}(K\alpha) = K_1$.

A similar argument works for $\alpha \in \mathrm{acl}_2(A_1) \setminus A_1$. Now, every element α of K_1 can be attained by adding a finite number of elements $\alpha_1, \ldots, \alpha_m = \alpha$ with $\alpha_{j+1} \in \mathrm{acl}_i(A_1\alpha_1 \ldots \alpha_j)$. So, working with $A_1' := \mathrm{acl}_0(A_1\alpha_1 \ldots \alpha_{m-1})$ instead of A_1, the same argument shows that $\overline{K\alpha} \cap (A_1' \setminus K)$ is nonempty, and thus we are done by induction over $\mathrm{d}_0(A_1'/A_1)$. $\qquad \dashv$

For the moment we are not yet able to prove (2). This will be done in Section 6. $\quad \dashv$

COROLLARY 4.9. *Let* $K \leq L$ *be finitely generated* (*with* L *selfsufficient in* K^*), $\mathrm{d}(L/K) = d$. *Then there is a decomposition* $K = K_0 \leq K_1 \leq \cdots \leq K_{d+n} = L$, *where* K_i *over* K_{i-1} *is generic for* $i \leq d$ *and primitive for* $i > d$.

PROOF. First pick a d-basis $a_1, \ldots, a_d$ of L over K and then put $K_i := \langle K_{i-1}a_i \rangle$. Now $K_d \leq L$ is parasitic, so one concludes by Proposition 4.7. $\qquad \dashv$

We end the section with two lemmas.

LEMMA 4.10. *Let* $k \subseteq A$, $B \subseteq K$, *where we assume that* $A = \overline{A}$, $B = \overline{B}$, $\delta(B/A) = 0$ *and* $[\mathrm{acl}_1(B) \cup \mathrm{acl}_2(B)] \cap A = k$. *Then* $\langle B \rangle \cap A = k$.

PROOF. Set $B' := \mathrm{acl}_1(B)$. We have $B' \underset{k}{\overset{0}{\downarrow}} A$ by assumption, so $B' \underset{B}{\overset{0}{\downarrow}} A$. Because $\delta(B/A) = 0$ this implies $\delta(B'/AB) \leq 0$. Hence, $B' \underset{B}{\overset{2}{\downarrow}} A$ (as $A = \overline{A}$) and thus $\mathrm{acl}_2(B') \cap A \subseteq \mathrm{acl}_2(B) \cap A = k$ follows. Since $B = \overline{B}$ we know that $B \leq B'$ so that $\delta(B'/B) = 0$ hence also $\delta(B'/A) = 0$ and $B' = \overline{B'}$. By symmetry the same is true for $B'' := \mathrm{acl}_2(B)$, and the claim follows by induction. $\qquad \dashv$

LEMMA 4.11. *Let* $K \leq L$ *be primitive. Then there is a unique minimal* $A = \overline{A} \supset K$ *controlling* L/K. A/K *is prime, and we put the* length *of* L/K *equal to the length of* A/K.

PROOF. Note that if $K \subseteq A'$, then L is controlled by A' if and only if $A' \supsetneq K$ and $\delta(A'/K) = 0$. Consider $A = \overline{A}$ and $B = \overline{B}$, where $K \subseteq A, B$ and both A and B control L. Suppose that A is minimal such. By Lemma 4.10 there is an element $b \in \mathrm{acl}_1(B) \cap A$, $b \notin K$. Since $A = \overline{Kb}$ we deduce that $A \subseteq \mathrm{acl}_0(Bb) = \overline{Bb}$. By modularity of T_0 and the fact that $d_0(A/K)$ and $d_0(B/K)$ both equal at least 2, we have $A \cap B \supsetneq K$, whence $A \subseteq B$ by minimality of A. Finiteness of $\mathrm{d}_0(A/K)$ is clear. $\qquad \dashv$

§5. **Axiomatisation.** *From now on we assume good control throughout the paper.* In this section we axiomatise the theory of the generic model which we denote by T_ω. It turns out that being rich is model-theoretically significant, since the rich models are exactly the ω-saturated models of T_ω. This is shown in Proposition 5.3. Before we start, we would like to stress that at this stage (the free fusion) no DMP assumption about the theories T_i is required.

DEFINITION 5.1. Let $k, l \in \mathscr{C}_0$ with $k \leq l$ primitive. A controlling pair (A, B) for l/k is called *prime*, if A/B is a prime extension.

We recall that by definition of a controlling pair (see 3.8) B and A are finite, B controls k, A controls l and $k \underset{B}{\overset{i}{\downarrow}} A$ for $i = 1, 2$. Thus, by Lemma 4.4, $A_k := \mathrm{acl}_0(Ak)$ is a prime extension of k. Therefore, it coincides with the minimal controlling set found in Lemma 4.11.

LEMMA 5.2. *Let* $k, l \in \mathscr{C}_0$ *with* $k \leq l$ *primitive, and* $B_0 \subseteq_\omega k$. *Then there exists* $B \supseteq B_0$ *and* $A \supseteq B$ *such that* (A, B) *is a prime controlling pair for* l/k.

PROOF. Take $\tilde{A} \supseteq k$ as given by Lemma 4.11 and choose an $\mathscr{L}_0$-basis $\overline{a}$ of $\tilde{A}/k$. Choose $B \subseteq_\omega k$ controlling k such that $B_0 \subseteq B = \overline{B}$ and $\overline{a} \mathop{\underset{B}{\perp}}^i k$, for $i = 1, 2$. Now, put $A := \mathrm{acl}_0(B\overline{a})$. It is routine to check that A/B works. $\dashv$

We now consider A/B, a prime controlling pair of length n controlling some $k \leq l$, where we assume that $\overline{a}$ enumerates A and $\overline{b}$ enumerates B. For all finite $\overline{c} \subseteq k$ such that $\mathrm{tp}_{\mathscr{L}_i}(\overline{a}/\overline{b}\overline{c})$ is stationary for $i \in \{0, 1, 2\}$ consider $\mathscr{L}_i$-formulas $(i = 1, 2)$ $\varphi^i_{l/k}(\overline{x}, \overline{u})$ (depending only on A/B) and $\tilde{\varphi}^i_{l/k}(\overline{x}, \overline{u}, \overline{z})$ (depending on A/B and $\overline{c}$) with the following properties (choose $\tilde{\varphi}^0_{l/k}(\overline{x}, \overline{u}, \overline{z})$ isolating $\mathrm{tp}_{\mathscr{L}_0}(\overline{a}\overline{b}\overline{c})$):

(0) $\tilde{\varphi}^i_{l/k}(\overline{x}, \overline{u}, \overline{z}) \vdash \varphi^i_{l/k}(\overline{x}, \overline{u}) \wedge \tilde{\varphi}^0_{l/k}(\overline{x}, \overline{u}, \overline{z})$.

(1) For $i \in \{1, 2\}$, $\tilde{\varphi}^i_{l/k}(\overline{x}, \overline{b}, \overline{c})$ is $\mathscr{L}_i$-stationary, $\models \tilde{\varphi}^i_{l/k}(\overline{a}, \overline{b}, \overline{c})$, and $\overline{a}$ is $\mathscr{L}_i$-generic in $\varphi^i_{l/k}(\overline{x}, \overline{b})$ over k.

(2) If $\overline{b}', \overline{c}'$ are such that $\tilde{\varphi}^i_{l/k}(\overline{x}, \overline{b}', \overline{c}')$ is not empty, then $\mathrm{MR}_i(\tilde{\varphi}^i_{l/k}(\overline{x}, \overline{b}', \overline{c}')) = \mathrm{MR}_i(\varphi^i_{l/k}(\overline{x}, \overline{b}')) = \mathrm{MR}_i(\varphi^i_{l/k}(\overline{x}, \overline{b}))$.

(3) For all partitions $\overline{x} = \overline{x}_1\overline{x}_2$ with $0 < m := \mathrm{d}_0(\overline{a}_1/\overline{a}_2\overline{b}) < n$ there are $k_1, k_2 \in \mathbb{N}$ such that $k_1 + k_2 < m$ and $\mathrm{MR}_i(\varphi^i_{l/k}(\overline{x}_1, \overline{a}'_2, \overline{b}')) \leq k_i$ for all $\overline{a}'_2, \overline{b}'$.

Finding $\tilde{\varphi}^0_{l/k}$ is easy because T_0 is $\aleph_0$-categorical, so (0) is a triviality. For (1) we use the fact that A/B is a prime controlling pair, (2) and (3) are possible by definability of MR in $\mathscr{L}_i$ (for (3) we use the total categoricity of T_0 again). Hence, one can always find $\varphi^1_{l/k}, \tilde{\varphi}^1_{l/k}, \varphi^2_{l/k}, \tilde{\varphi}^2_{l/k}$ such that (0)–(3) hold. Note that the choice of the above formulas depends not only on l/k but also on $\overline{c}$, a fact that is not reflected by the notation in the hope of keeping it readable.

For every primitive l/k we now choose a prime controlling pair A/B. Moreover, for $\overline{c}_1 \supseteq \overline{c}$ as above we choose the formulas in a coherent way, i.e., so that $\tilde{\varphi}^i_{l/k, \overline{c}_1}$ implies $\tilde{\varphi}^i_{l/k}$, where $\tilde{\varphi}^i_{l/k, \overline{c}_1}$ are the formulas we choose for A/B and $\overline{c}_1$. Putting $\theta_{l/k, \overline{c}}(\overline{u}, \overline{z}) := \{(\overline{u}, \overline{z}) \mid \models \exists \overline{x}\tilde{\varphi}^1_{l/k}(\overline{x}, \overline{u}, \overline{z}) \wedge \exists \overline{x}\tilde{\varphi}^2_{l/k}(\overline{x}, \overline{u}, \overline{z})\}$, the axiom corresponding to A/B and $\overline{c}$ is as follows (note that for every primitive extension we chose one prime controlling pair A/B, but due to the varying $\overline{c} \in k$ there are infinitely many axioms concerning A/B):

$$\psi_{k/l, c} := \forall \overline{u}\, \forall \overline{z}\{\theta_{l/k, \overline{c}}(\overline{u}, \overline{z}) \rightarrow \exists \overline{x}[\tilde{\varphi}^1_{l/k}(\overline{x}, \overline{u}, \overline{z}) \wedge \tilde{\varphi}^2_{l/k}(\overline{x}, \overline{u}, \overline{z})]\}.$$

The above axioms should be seen as a first order approximation of the statement "for every $k \leq M$ (where M is some rich model) and every extension $k \leq l$ where k and l are finitely generated, there exists an embedding $f : l \rightarrow M$ such that $f{\restriction}k = \mathrm{id}$ and $f(l) \leq M$".

Decomposing parasitic into primitive extensions and using the fact (which we prove in Lemma 5.4) that the generic type can be approximated by parasitic types it follows that if one can approximate the existence of such embeddings in the case of primitive extensions $l \geq k$, one can approximate this for arbitrary $l \geq k$ as well. By Lemma 5.2 and the definition of controlling, it is enough to ensure that the types of controlling pairs are realised as strong subsets. The formulas $\tilde{\varphi}^i_{l/k}$ are introduced in order to ensure that for any finite set of parameters $\overline{c} \subseteq k$ there exists an approximate solution to $\mathrm{qftp}(A/k)$ which is $\mathscr{L}_0$-independent from $\overline{c}$ over B. As we will see, this guarantees that in a saturated model M, l can be strongly k-embedded into M.

Let T'_ω be the theory axiomatised by $\mathrm{Th}(\tilde{\mathscr{C}}_0) \cup \{\psi_{l/k,\overline{c}}\colon l/k$ primitive, $\overline{c} \subseteq_\omega k \leq l \leq M \in \tilde{\mathscr{C}}_0$ generic$\}$.

PROPOSITION 5.3. *The theories T'_ω and T_ω coincide, and the rich structures in $\tilde{\mathscr{C}}_0$ are precisely the ω-saturated models of T'_ω. In particular, the generic model is saturated.*

PROOF. Let $K \in \tilde{\mathscr{C}}_0$ be rich. We first show that $K \models T'_\omega$. Let $\overline{b}'\overline{c}' \in K$, with $K \models \theta_{l/k,\overline{c}}(\overline{b}',\overline{c}')$ for some primitive l/k and $\overline{c} \subseteq k$. We have to find a solution of $\tilde{\varphi}^1_{l/k}(\overline{x},\overline{b}',\overline{c}') \wedge \tilde{\varphi}^2_{l/k}(\overline{x},\overline{b}',\overline{c}')$ in K. Set $k' := \mathrm{cl}(\overline{b}'\overline{c}') \leq K$, so k' is in $\mathscr{C}_0$. We now construct a certain primitive extension l' of k' such that l' contains a solution of $\tilde{\varphi}^1_{l/k}(\overline{x},\overline{b}',\overline{c}') \wedge \tilde{\varphi}^2_{l/k}(\overline{x},\overline{b}',\overline{c}')$. Once such an extension l' is found we are done by the fact that K is rich. Just take the image of the solution under the k'-embedding of l' in K.

We build l' as follows. For $i = 1,2$, let p'_i be some generic $\mathscr{L}_i$-type in $\tilde{\varphi}_i(\overline{x},\overline{b}',\overline{c}')$ over k', and choose $\overline{a}'_i \models p'_i$ (in some $k'_i \succeq_{\mathscr{L}_i} k'$). So by (2) we must have $\overline{a}'_i \underset{\overline{b}'}{\overset{i}{\downarrow}} k'$. As in the proof of Lemma 4.4 we know that $\overline{a}'_i \underset{\overline{b}'}{\overset{0}{\downarrow}} k'$. On the other hand, by the definition of the $\tilde{\varphi}^i_{l/k}$, we have $\mathrm{tp}_{\mathscr{L}_0}(\overline{a}'_1/\overline{b}'\overline{c}') = \mathrm{tp}_{\mathscr{L}_0}(\overline{a}'_2/\overline{b}'\overline{c}')$, and thus by stationarity of this $\mathscr{L}_0$-type, $p'_1 \restriction \mathscr{L}_0 = p'_2 \restriction \mathscr{L}_0$. By Lemma 3.10 there is an extension $l' \geq k'$ which is controlled by some $\overline{a}' \models p'_1 \cup p'_2$ over k'. From Lemma 4.4(1) we deduce that $k' \leq A_1 := \mathrm{acl}_0(k'\overline{a}')$ is prime, whence l' is a primitive extension by Lemma 4.8. This shows in particular that T'_ω is consistent, since rich models exist.

We prove next that ω-saturated models of T'_ω are rich. Once this is established the whole proposition is proved, i.e., it follows that $T'_\omega = T_\omega$ and that any rich model is ω-saturated, since any two rich structures are $\mathscr{L}_{\infty,\omega}$-equivalent, ω-saturation is invariant under $\mathscr{L}_{\infty,\omega}$-equivalences and every structure has ω-saturated elementary extensions.

Suppose $K \models T'_\omega$ is ω-saturated and $k \leq K$ is finitely generated. For any $k \leq l \in \mathscr{C}_0$ we have to find a strong k-embedding into K. We first treat the case of a primitive extension $k \leq l$. Take the prime controlling pair A/B we chose for (the isomorphism type of) this extension in the axiomatisation, and choose some $\overline{c}$ over which the $\mathscr{L}_i$-types of $\overline{a}/\overline{b}$ are stationary for $i \in \{0,1,2\}$. Now look at the $\mathscr{L}_i$-formulas $\varphi^i_{l/k}(\overline{x},\overline{u}), \tilde{\varphi}^i_{l/k}(\overline{x},\overline{u},\overline{z})$ which we chose satisfying (0)–(3). Let p_i be the (by (1)) unique $\mathscr{L}_i$-generic type in $\tilde{\varphi}^i_{l/k}(\overline{x},\overline{b},\overline{c})$. Note that any $\overline{a}_0 \models p_1 \cup p_2$ would give rise to a k-embedding of l into K. As a consequence of the fact that $k \leq K$ we get that

$$\tilde{\varphi}^1_{l/k}(\overline{x},\overline{b},\overline{c}) \wedge \tilde{\varphi}^2_{l/k}(\overline{x},\overline{b},\overline{c}) \wedge \overline{x} \underset{\overline{b}\overline{c}}{\overset{0}{\downarrow}} k \vdash p_1 \cup p_2$$

but our axioms explicitly state that any finite approximation of this type is realised in any model, and by saturation the type itself is realised in K.

Finally, we treat the case of a generic extension $k \leq l = \langle ka \rangle$, with $\mathrm{d}(a/k) = 1$, a being a singleton. We need a lemma:

LEMMA 5.4. *The generic type can be approximated by parasitic types, i.e., if a is generic over k, then for every formula $\varphi(x)$ with parameters in k satisfied by a there is an $a' \in K$ such that $\delta(a'/k) = 0$ and $K \models \varphi(a')$.*

PROOF OF THE LEMMA. Assume that k is controlled by $\overline{b}$. Let us first observe that a is generic over k if and only if, for every $n \in \mathbb{N}$, a satisfies $q_n(x)$. Here $q_n(x)$

denotes the following partial type over $\overline{b}$:

$$\forall y_1 \ldots y_n \delta(x\overline{y}/\overline{b}) \geq 1.$$

To prove the lemma it is thus sufficient to find singletons $a_n \in K$ such that $d_0(\overline{A_n}/\overline{b}) \geq n$, where A_n denotes the finite set $\overline{b}a_n$. We will construct a parasitic extension $k \leq l_n = \text{cl}(ka_n)$ with a_n as required.

Recall the definition of $\langle \cdot \rangle^n$ in 4.5. The construction of l_n goes as follows. First, consider $A := \langle ka \rangle^n$, where a is generic over k. By assumption (Geom) clearly A is neither acl_1-closed nor acl_2-closed, so we can choose $c_i \in \text{acl}_i(A) \setminus A$. We now apply Lemma 3.10 to $p_i := \text{tp}_i(\text{acl}_0(Ac_i)/k)$ over k, which provides an extension l_n/k which is easily seen to be parasitic. If $a_n \in B_n \models p_1 \cup p_2$ is the element corresponding to $a \in A$, then $\text{cl}(ka_n) = l_n$ and we see that $\tilde{A}_n := \overline{ka_n} \subseteq \langle ka_n \rangle^{n+1}$ and $\tilde{A}_n \nsubseteq \langle ka_n \rangle^n$ (by construction). Therefore, $d_0(\overline{A_n}/\overline{b}) \geq d_0(\tilde{A}_n/k) \geq n + 1$. Note that the last inequality follows from the fact that if $\tilde{A}_n \cap \langle ka_n \rangle^m = \tilde{A}_n \cap \langle ka_n \rangle^{m+1}$ for some m, then $\tilde{A}_n \subseteq \langle ka_n \rangle^m$. ⊣

Combining the preceding lemma with the first part of the proof and the existence of a finite decomposition of l_n/k into primitive extensions, we conclude, by ω-saturation of K, that if l/k is generic, then l can be strongly k-embedded into K. This concludes the proof of the Proposition. ⊣

COROLLARY 5.5 (Quantifier elimination for strong subsets). *Let A and A' be (not necessarily finite) tuples in $K \models T_\omega$. Then $\text{tp}(A) = \text{tp}(A')$ if and only if $\overline{A} \cong_{\mathscr{L}} \overline{A'}$. Thus, T_ω is near model complete (i.e., has QE to the level of boolean combinations of existential formulas).* ⊣

PROOF. Once the corollary is shown for finite tuples, it follows for infinite tuples. We may thus work with finite A and A'. Suppose that $\overline{A} \cong_{\mathscr{L}} \overline{A'}$. Due to good control, this extends to an isomorphism of $\text{cl}(A)$ and $\text{cl}(A')$. By Proposition 5.3, saturated models of T_ω are rich. Thus, by a back-and-forth argument, we deduce that $\text{tp}(A) = \text{tp}(A')$. The other implication is clear.

As usual, near model completeness follows. ⊣

LEMMA 5.6. *Selfsufficient closure equals algebraic closure in the sense of T_ω, i.e., for all A one has $\text{cl}(A) = \text{acl}_{T_\omega}(A)$.*

PROOF. The inclusion $\text{cl}(A) \subseteq \text{acl}_{T_\omega}(A)$ being obvious, it suffices to show that $K := \text{cl}(A)$ is algebraically closed. Let $K \leq K^* \models T_\omega$ and $\alpha \in K^* \setminus K$, so we get $K \leq \text{cl}(K\alpha) =: L \leq K^*$. For $n \in \mathbb{N}$ let $L_1, \ldots, L_n$ be isomorphic copies of L over K. Provided K^* is sufficiently saturated, the free amalgam $L_1 \otimes_K \cdots \otimes_K L_n$ can be strongly embedded into K^*. Thus, $\alpha \notin \text{acl}_{T_\omega}(K)$. ⊣

LEMMA 5.7. (1) *Let $\overline{a}, \overline{b} \in K \models T_\omega$, with $\overline{b} \leq K$ and $\delta(\overline{a}/\overline{b}) = 0$. Then $\text{tp}(\overline{a}/\overline{b})$ is isolated.*

(2) *Let $k \in \mathscr{C}_0$ be a finitely generated fusion, embedded strongly into the ambient model. If l/k is a primitive extension and A/B is a prime controlling pair for l/k, then $\text{tp}(A/k)$ is isolated.*

PROOF. Choose $\mathscr{L}_i$-formulas $\varphi_i(\overline{x}, \overline{b})$ such that $\text{tp}_{\mathscr{L}_i}(\overline{a}/\overline{b})$ is the only generic type in $\varphi_i(\overline{x}, \overline{b})$ over $\overline{b}$ and let $\varphi_0(\overline{x}, \overline{b})$ isolate $\text{tp}_{\mathscr{L}_0}(\overline{a}/\overline{b})$. The fact that $\wedge_{i=0}^2 \varphi_i(\overline{x}, \overline{b})$ isolates $\text{tp}(\overline{a}/\overline{b})$ follows from Corollary 5.5, as $\overline{b} \leq K$.

To show (2), we choose $\mathscr{L}_i$-formulas $\varphi_i(\overline{x}, \overline{c}_i)$ such that $\mathrm{tp}_{\mathscr{L}_i}(A/k)$ is the only generic type in $\varphi_i(\overline{x}, \overline{c}_i)$ over k. W.l.o.g. we can suppose that the $\overline{c}_i$ contain B and $\psi(\overline{x}, \overline{c}) := \varphi_1(\overline{x}, \overline{c}_1) \wedge \varphi_2(\overline{x}, \overline{c}_2)$ implies $\mathrm{tp}(\overline{x}/B) = \mathrm{tp}(A/B)$. We conclude, combining Lemma 4.4 with Lemma 4.6, that $\mathrm{tp}(A/k)$ is isolated by $\psi(\overline{x}, \overline{c})$.   ⊣

§6. ω-**stability and ranks.** From now on we will be working in a monster model $K^* \models T_\omega$.

LEMMA 6.1. T_ω is stable.

PROOF. Let $K \leq K^*$ and $\overline{\alpha} \in K^*$. To determine $\mathrm{tp}(\overline{\alpha}/K)$ we have to look at $L := \mathrm{cl}(K\overline{\alpha})$, which is equal to $\langle K\overline{a} \rangle$ for some finite $\overline{a} \in K^*$ controlling L over K. Now choose k finitely generated and selfsufficient in K such that $\mathrm{d}(\overline{a}/K) = \delta(\overline{a}/K) = \delta(\overline{a}/k) = \mathrm{d}(\overline{a}/k)$, and let $l := \langle k\overline{a} \rangle$ (which is strong in K^*). By Lemma 3.11 we know that l and K are in free position over their intersection k', with $L \cong l \otimes_{k'} K$. Now $\mathrm{tp}(\overline{a}/k')$ together with the choice of $k' \subseteq K$ determine completely $\mathrm{tp}(\overline{\alpha}/K)$, whence $|S(K)| \leq 2^{\aleph_0} \cdot |K|^{\aleph_0}$.   ⊣

LEMMA 6.2. *Let* $K \leq L \leq K^*$ *and* $K' \leq K^*$ *with* $K' \supseteq K$. *If we put* $L' := \langle K'L \rangle$, *the following are equivalent*:

(1) $L \underset{K}{\mathop{\smile\hskip-0.8em\smile}} K'$,

(2) $L' \cong L \otimes_K K'$ *and* L' *is strong in* K^*.

PROOF. Note that $p := \mathrm{tp}(L/K)$ can always be extended to a type over K' as described in (2). We can clearly suppose that $K' \models T_\omega$. Then, since this extension is invariant under all automorphisms (of K' over K) by the uniqueness of the free amalgam, it must coincide with the unique non-forking extension of p to K'.   ⊣

REMARK 6.3. (1) *Alternatively, one can prove stability in showing that* (2) *in Lemma 6.2 gives rise to a notion of independence* (*including 'boundedness of free extensions'*). *This gives stability of the theory, and the notion of independence has to coincide with the non-forking relation, see e.g.,* [1, VII.1].

(2) *From Lemma 6.2 together with Lemma 5.6 one deduces that in* T_ω *types over* (*real*) *algebraically closed sets are stationary.*   ⊣

Observe that in view of this characterisation of non-forking in T_ω, the content of Lemma 4.2 is exactly that $\mathrm{U}(L/K) = 1$ if $K \leq L$ is primitive. By the Lascar inequalities this determines the U-rank of every parasitic extension, too. There are enough isolating formulas, in the sense that if $K \leq L$ is parasitic (L being controlled by $\overline{a}$ over K) then we find $\varphi(\overline{x}, \overline{b}) \in p := \mathrm{tp}(\overline{a}/K)$ isolating p from all other types q with $\mathrm{U}(q) \geq \mathrm{U}(p)$. Thus, Morley rank exists and equals U-rank on parasitic types . The only 1-type we missed so far is the generic type. But any of its forking extensions is necessarily parasitic which means that we may conclude:

PROPOSITION 6.4. T_ω *is* ω-*stable. The Lascar rank of a parasitic extension equals its Morley rank and is given by the length of a decomposition into primitive extensions* (*so in particular is finite*). *Lascar rank and Morley rank of the generic are* $\leq \omega$.   ⊣

Note that this partially settles the second part of Proposition 4.7. We now show the "uniqueness" result. Combining Lemma 3.11 with Lemma 6.2 one gets

LEMMA 6.5. *Let* L *be a parasitic extension of* $K = \mathrm{cl}(K)$ *and* $M = \mathrm{cl}(M)$ *an arbitrary extension of* K. *Then* $L \underset{K}{\mathop{\smile\hskip-0.8em\smile}} M$ *iff* $L \cap M = K$ *iff* $L \underset{K}{\overset{0}{\mathop{\smile\hskip-0.8em\smile}}} M$.   ⊣

In the context of Proposition 4.7 (and with notation there) we know that since $K_1 \not\perp_K L$, there is some minimal i_0 such that $K_1 \not\perp_{K'_{i_0-1}} K'_{i_0}$. Thus, $K_1 \otimes_K K'_{i_0-1} \simeq K'_{i_0}$, whence $L_1 \simeq L'_1$, too. Since the length of the decomposition for L/K_1 is $n-1$, we can inductively assume that the uniqueness result holds for $K_1 \leq \cdots \leq K_n = L$ and $K_1 \leq K''_2 \leq \cdots \leq K''_n = L$, where K''_i equals $K_1 \otimes_K K'_{i-1}$ for $i < i_0$ and K'_i for $i \geq i_0$. This finishes the proof of 4.7.

COROLLARY 6.6. *Parasitic types (and more generally types orthogonal to the generic type) are 1-based. In particular, the pregeometry of a primitive type is locally modular. If T_0 has a trivial pregeometry, then the pregeometry of every primitive type is trivial.*

PROOF. The characterisation of non-forking for parasitic extensions in Lemma 6.5 proves 1-basedness for parasitic types. Now consider a type p which is orthogonal to the generic type. Then, p is analysable in parasitic types (even in primitive types). It is known that types which are analysable in 1-based types are 1-based. Thus, p is 1-based, too.

Local modularity of primitive types is a consequence. The last statement of the corollary (about trivial pregeometries) will follow from Remark 6.9. ⊣

EXAMPLE 6.7. Let $T_1 = T_2 = ACF_p$, T_0 being the theory of $\mathbb{F}_p$-vector spaces. Over some fusion K we take a, a' $\mathbb{F}_p$-independent over K. We further demand that, for some $b \in K$, we have $a = b *_1 a'$ as well as $a = b *_2 a'$, where $*_i$ denotes multiplication in the sense of T_i. By Lemma 3.10, this defines a primitive extension L of K which is controlled by aa' over K. Since $\mathrm{tp}(aa'/K)$ is a group generic, its pregeometry cannot be trivial.

The following proposition, important in itself, provides an alternative proof (via Zil'ber's characterisation of the geometry of $\aleph_0$-categorical strongly minimal sets) for the fact that every primitive type is locally modular.

PROPOSITION 6.8. *Let A/B be a prime controlling pair for the primitive extension l/k, where k is strong. Then $\mathrm{tp}(A/k)$ contains a totally categorical strongly minimal formula. There is even such a formula which isolates $\mathrm{tp}(A/k)$.*

PROOF. By Lemma 5.7 we can choose $B \subseteq C = \overline{C} \subseteq_\omega k$, $\overline{c}$ an enumeration of C, and an $\mathscr{L}(C)$-formula $\psi(\overline{x}, \overline{c})$ isolating $\mathrm{tp}(A/k)$. We show that $\psi(\overline{x}, \overline{c})$ defines a locally finite strongly minimal set, which will settle everything. Strong minimality is clear by the choice of the formula, and local finiteness is a consequence of the following claim:

CLAIM. Let $M \models T_\omega$ and $D := \psi(M, \overline{c})$. For $E \subseteq D$ we then have $\mathrm{acl}_S(E) = \mathrm{acl}_0(EC) \cap D$, where $S := \mathrm{Th}(D)$, in the language consisting of the traces for all C-definable sets.

To prove this claim, we can clearly assume that E is finite. Note that $\delta(E/C) = 0$ (≥ 0 is obvious, and ≤ 0 follows from Lemma 4.4 and induction on the cardinality of E). One has $\mathrm{acl}_S(E) = \mathrm{acl}_{T_\omega}(EC) \cap D$. As $\mathrm{acl}_{T_\omega} = \mathrm{cl}$, for any $e \in \mathrm{acl}_S(E)$, the fact that $e \in \mathrm{acl}_0(EC)$ follows from Lemma 4.6. ⊣

In fact, the proof of the preceding claim gives a bit more:

REMARK 6.9. *Let $A_1, \ldots, A_n, A'$ be prime over $C = \overline{C}$. Then $A' \subseteq \mathrm{cl}(\cup_{i=1}^n A_i)$ iff $A' \subseteq \mathrm{acl}_0(\cup_{i=1}^n A_i)$.* ⊣

DEFINITION 6.10. An expansion $T_0 \subseteq T_1$ of strongly minimal (complete) theories is called *relatively trivial* if for all $\overline{a} \in M \models T_1$ we have $\mathrm{acl}_1(\overline{a}) = \mathrm{acl}_1^u(\mathrm{acl}_0(\overline{a}))$, where $\mathrm{acl}^u(A) := \bigcup_{x \in A} \mathrm{acl}(x)$ denotes the unitary algebraic closure.

Following [19], a pregeometry is called (*locally*) *projective* if it is (locally) modular non-trivial.

REMARK 6.11. (1) *A relatively trivial expansion remains so if a set of parameters is added.*
(2) *If T_0 is a trivial strongly minimal theory, then the expansion T_1 is relatively trivial if and only if it has a trivial pregeometry.*
(3) *A relatively trivial expansion of a (locally) projective theory is geometry preserving.*

PROOF. Only the last point deserves an argument. By (1), we can assume T_0 to be projective. Clearly, a relatively trivial expansion of a modular pregeometry is modular, so T_1 is projective, too. There is a strongly minimal abelian group A interpretable in $M \models T_0$, A not almost orthogonal to M (over $\emptyset$). W.l.o.g. we can suppose that $M = A$. Recall that the geometry of a modular strongly minimal group is determined by its field of definable quasi-endomorphisms. Thus, if the expansion is not geometry preserving, necessarily the inclusion $K_0 \subseteq K_1$ must be strict, where K_i denotes the field of quasi-endomorphisms of A definable in T_i. Now choose $\alpha \in K_1 \setminus K_0$. For b_1, b_2 independent generics we claim that $b_1 + \alpha b_2 \in \mathrm{acl}_1(b_1 b_2) \setminus \mathrm{acl}_1^u(\mathrm{acl}_0(b_1 b_2))$ (attention: this is an abuse of notation, since this computation takes place in A/H for a certain finite group $H \subseteq \mathrm{acl}_1(\emptyset)$). Suppose, maybe factoring out some larger finite group, we had $b_1 + \alpha b_2 = \alpha'(\beta_1 b_1 + \beta_2 b_2)$ with $\alpha' \in K_1$, $\beta_i \in K_0$. Then, $\alpha' \beta_2 = \alpha$ and $\alpha' \beta_1 = id$ (since this is true generically). Thus, $\beta_1^{-1} = \alpha'$ and finally $\alpha = \beta_1^{-1} \beta_2$, contradicting the choice of α. $\dashv$

We now supply a partial converse to the first part of the preceding remark.

LEMMA 6.12. *Assume T_0 modular, and let $A \subseteq M^* \models T_1$ be a set of parameters.*

(1) *If $T_0(A) \subseteq T_1(A)$ is relatively trivial, so is $T_0 \subseteq T_1$.*
(2) *If $T_0 \subseteq T_1$ is not relatively trivial, there exists a natural number n such that, whenever $\mathrm{d}_1(A/\emptyset) \geq n$, we have $\mathrm{acl}_1^u(\mathrm{acl}_0(Aa_1a_2)) \subsetneq \mathrm{acl}_1(Aa_1a_2)$ for every $\mathscr{L}_1$-generic (over A) pair a_1, a_2.*

PROOF. To prove (1), suppose that $T_0 \subseteq T_1$ is not relatively trivial, and choose an $\mathscr{L}_1$-independent generic tuple $\overline{b}$ and an element a such that $a \in \mathrm{acl}_1(\overline{b}) \setminus \mathrm{acl}_1^u(\mathrm{acl}_0(\overline{b}))$. In addition, suppose that $a\overline{b} \underset{}{\overset{1}{\smile}} A$, so in particular $a \notin \mathrm{acl}_1(A)$. Since the expansion $T_0(A) \subseteq T_1(A)$ is relatively trivial, there is $b' \in \mathrm{acl}_0(A\overline{b})$ such that $a \in \mathrm{acl}_1(Ab')$. Now, since T_0 is modular (and b' is not in $\mathrm{acl}_1(A)$) there is some $b'' \in \mathrm{acl}_0(\overline{b})$ which is $\mathscr{L}_0$-interalgebraic with b' over $\mathrm{acl}_1(A)$. Thus, $a \in \mathrm{acl}_1(Ab'')$. But we have $a\overline{b}b'' \underset{}{\overset{1}{\smile}} A$ and therefore $a \underset{b''}{\overset{1}{\smile}} A$, from which we deduce that $a \in \mathrm{acl}_1(b'')$. This is a contradiction, so (1) is proved.

The proof of (2) is just a variation of this. If the length of the tuple $\overline{b}$ used in the proof is m, then $n := m - 2$ will do. We only have to be careful in choosing $\mathrm{tp}_1(a/\overline{b})$. In fact we should minimise the length of $\overline{b}$ (inside $\mathrm{acl}_0(\overline{b})$) first, and then work with a generic (over $\emptyset$) tuple $c_1, \ldots, c_{m-2} \in A$, thus replacing $\overline{b}$ by $\overline{c} a_1 a_2$. $\dashv$

The following proposition shows that the free fusion is not necessarily of 'infinite rank', so one should maybe stick to the distinction *non-collapsed* vs. *collapsed* instead of *infinite rank* vs. *finite rank*. Moreover—as the referee pointed out to us—it clarifies the last section of [4], where conditions are studied which make possible exact rank computations.

PROPOSITION 6.13 (Rank Dichotomy). *Let* $\mathfrak{g}$ *denote the generic type.*
(1) *If* T_1 *and* T_2 *are trivial then* $\mathrm{U}(\mathfrak{g}) = 1$ *and* $\mathrm{MR}(\mathfrak{g}) = 2$.
(2) *If one of the* T_i *is not trivial, then* $\mathrm{U}(\mathfrak{g}) = \mathrm{MR}(\mathfrak{g}) = \omega$.

PROOF. The proof of part (1) is rather easy, and we do not give all the details. In fact, the triviality assumptions have $\langle\cdot\rangle = \mathrm{cl}$ as a consequence. In addition, in this context the operator $\langle\cdot\rangle$ satisfies the Steinitz exchange property and is trivial (in the sense that $\langle A\rangle = \cup_{a\in A}\langle a\rangle$ for all A). From this one establishes $\mathrm{U}(p) = 1$ for all non-algebraic 1-types, especially for the generic type $\mathfrak{g}$. The general context guarantees that the number of different non-algebraic 1-types is infinite (this is shown in the proof of Lemma 5.4, where $\mathfrak{g}$ is approximated by parasitic types). Hence we may conclude that $\mathrm{MR}(\mathfrak{g}) = 2$.

In order to show (2), w.l.o.g. we may assume that T_1 is not trivial, so the expansion $T_0 \subseteq T_1$ is not relatively trivial by Remark 6.11. We first prove the following

CLAIM. If $k \leq l = \langle ka\rangle$ is generic, there exists $a' \in \mathrm{acl}_1(\mathrm{acl}_2(A))$ such that $\mathrm{acl}_1(ka') \cap \mathrm{acl}_2(ka) = k$.

Since we always assume that $\mathrm{acl}_2(\emptyset)$ is infinite, in particular $d_1(k)$ is infinite for every fusion $k \in \mathscr{C}_0$. By condition (Geom), necessarily $d_1(\mathrm{acl}_2(ka)/k) \geq 2$, so (as $d_1(k/\emptyset)$ is infinite) we can apply the second part of Lemma 6.12 and find $a' \in \mathrm{acl}_1(\mathrm{acl}_2(ka)) \setminus \mathrm{acl}_1^u(\mathrm{acl}_2(ka))$. Thus, $\mathrm{acl}_1(ka') \cap \mathrm{acl}_2(ka) = k$ follows and the claim is established. Note that, since a' is not in $\mathrm{acl}_2(ka)$, it follows that $\mathrm{acl}_2(ka') \cap \mathrm{acl}_2(ka) = k$, too.

In order to see that $l' := \langle ka'\rangle \subsetneq l$ we apply Lemma 4.10 to $A := \mathrm{acl}_2(ka)$ and $B := \mathrm{acl}_0(ka')$. Note that l'/k is generic, too, and so $\mathrm{U}(l'/k) = \mathrm{U}(l/k)$. On the other hand, $\mathrm{U}(l/l') = n$ for some $0 < n \in \mathbb{N}$, as l/l' is parasitic and $l' \subsetneq l$.

Now, by the Lascar inequalities, $\mathrm{U}(l/l') + \mathrm{U}(l'/k) = n + \mathrm{U}(\mathfrak{g}) \leq \mathrm{U}(l/k) = \mathrm{U}(\mathfrak{g})$, where $n > 0$. Since $\mathrm{U}(\mathfrak{g}) \leq \omega$ by Proposition 6.4, necessarily $\mathrm{U}(\mathfrak{g}) = \omega$, as $n + m > m$ for every finite m. $\mathrm{MR}(\mathfrak{g}) = \omega$ follows. ⊣

REMARK 6.14. *Let* $K \leq L \leq M$ *with* $K \leq M$ *finitely generated. Then,* $K \leq L$ *is finitely generated, too.*

PROOF. Picking a finite d-basis of L/K we reduce to the case $\delta(L/K) = 0$. By Proposition 6.4, we have $\mathrm{U}(M/K) = \omega \cdot m + n$, for some $n, m \in \mathbb{N}$. The length of chains of primitive extensions in M starting with K is thus bounded by n, due to the Lascar inequalities. ⊣

REMARK 6.15. *If the hypothesis of* good control *is dropped, there is still a reasonable theory of the free fusion. This is studied in* [14], *where among other things the following is shown:* $(\aleph_1)$*-rich structures in* $\tilde{\mathscr{C}}_0$ *are model theoretically meaningful, i.e., sufficiently saturated models of their common* $\mathscr{L}$*-theory* T_ω*. This theory is no longer complete. Its completions are given by fixing* $\mathrm{qftp}_{\mathscr{L}}(\langle\emptyset\rangle)$*, more generally* $\mathrm{tp}(A) = \mathrm{tp}(A')$ *iff* $\mathrm{cl}(A) \cong_{\mathscr{L}} \mathrm{cl}(A')$*. Every completion is supersimple, and one has* $A \mathop{\smile\hskip-0.9em|\,}_B C$ *if and only if* $\mathrm{cl}(AB) \cap \mathrm{cl}(BC) = \mathrm{cl}(B)$ *and* $\mathrm{cl}(AB) \mathop{\smile\hskip-0.9em|\,}^{\mathrm{d}}_{\mathrm{cl}(B)} \mathrm{cl}(BC)$. SU*-rank is calculated as is* U*-rank under the assumption of good control.*

We give two examples illustrating the situation without good control:

EXAMPLES. (1) Recall Example 3.9. Here, every completion of T_ω is superstable. None of them is ω-stable, not even small, since there are $2^{\aleph_0}$ d-generic types over $\emptyset$. More generally, if both T_1 and T_2 are trivial, then every completion of the corresponding T_ω is superstable, as types over cl-closed sets are stationary.

(2) $T_0 :=$ theory of an equivalence relation E, all classes consisting of 3 elements, except one one-element class $\{0\}$,
$T_1 := \overline{\mathbb{F}_2}$-vector spaces, expanding E via $Exy: \Leftrightarrow x = \alpha y$ for some $\alpha \in \mathbb{F}_4^* = \{1, \lambda, \lambda^2\}$,
$T_2 :=$ theory of the following action of a group G strictly containing a subgroup $G_2 = \{\text{id}, \mu, \mu^2\} \cong \mathbb{Z}/3$: G acts trivially on 0 and freely on the complement of $\{0\}$. Here, $Exy: \Leftrightarrow x = \alpha y$ for some $\alpha \in G_2$.

In this example (in any completion T of T_ω), no d-generic type is stationary over any parameter set, in particular T is unstable. To see this, take x, b d-generic and independent over $A = \text{cl}(A)$. Now, $p := \text{tp}(x/A)$ has a non-forking extension to Ab with $\mu(x + b) = \lambda(x + b)$ and one with $\mu(x + b) = \lambda^2(x + b)$, so p is not stationary.

§7. **Non-orthogonality of types in** T_ω. In order to understand the "dimensions" of T_ω (i.e., the non-orthogonality classes of regular types) we need not consider imaginary types. This can be seen as follows. Primitive types are s.m. by 6.4, so in particular regular. The generic type is clearly regular. Thus, 4.7. shows that every real type is (almost) analysable in real regular types. Since imaginary types are always analysable in real types, this means that every type (even imaginary) is (almost) analysable in real regular types. We are thus lead to a study of non-orthogonality of types giving rise to primitive extensions.

DEFINITION 7.1. Let $k \leq K^*$ and l/k be primitive of length n. $p := \text{tp}_{\mathscr{L}}(\overline{a}/k)$ is an *admissible type* if $\overline{a}$ is an $\mathscr{L}_0$-basis of some A over k, where A is as in Lemma 4.11. One has $p \in S^n(k)$, and we put $n_p := n$, the *length* of the admissible type p. We will usually not distinguish between p and its parallelism class.

We note that admissible types are stationary. It follows from Proposition 6.8 that every admissible type (once based over a finitely generated strong fusion) is isolated by a totally categorical strongly minimal formula. All dimensions of T_ω (except the generic, of course) are given by admissible types, since for every minimal type orthogonal to the generic there is an admissible type which is not orthogonal to it. So it is sufficient to study (almost) orthogonality within this class of types.

Let $p, q \in S(k)$ be admissible types $(k = \text{cl}(k))$. If $p \not\perp_k^a q$, there are $\overline{a} \models p$ and $\overline{b} \models q$ such that $\overline{a}$ and $\overline{b}$ are interalgebraic over k and so $\text{cl}(k\overline{a}) = \text{cl}(k\overline{b})$ by Lemma 5.6. By the uniqueness part of Lemma 4.11 and the definition of admissible types, this means that $\overline{a}$ and $\overline{b}$ are inter-$\mathscr{L}_0$-algebraic over k.

When p, q are stationary types based on k, by abuse of notation, $p \perp_k^a q$ always means $p|k \perp_k^a q|k$.

We will only be interested in the case that T_0 is non-trivial. First, note that totally categorical strongly minimal trivial theories are basically equivalence relations with finite classes, and therefore, not of great interest. Second, if T_0 is trivial then by

Corollary 6.6 so is every strongly minimal set in T_ω. It is then easy to check that the original construction of the collapsed fusion given in [17] goes through virtually unaltered.

Recall that if T_0 is an uncountably categorical theory containing a strongly minimal set whose geometry is non-trivial, then in some expansion of T_0 by constants there is a definable infinite 1-based group (see [22, Ch. III]). Since we do not mind adding finitely many constants to T_0 we may as well assume that the group is already definable in T_0. The assumption that T_0 has a definable strongly minimal 1-based group is enough to obtain most of the results in this section (see also Remark 8.4), but for the sake of clarity and completeness we make the following

CONVENTION 7.2. From now on, T_0 is the theory of a 1-based totally categorical (not necessarily pure) group.

For convenience, we further assume that $M_0 := \mathrm{acl}_0(\emptyset) = \mathrm{dcl}_0(\emptyset)$. By ω-categoricity, M_0 is a finite subgroup of any model M of T_0. As we still assume good control, it might help to think of T_0 being such that $M_0 = \{0\}$, which ensures good control for any expansions T_1 and T_2, as $\mathrm{dcl}_0 = \mathrm{acl}_0$ in this case (see Lemma 3.6).

Let $\pi_i\colon M^{2n} \to M^n$ be the projection maps (π_1 mapping on the first n, π_2 on the last n coordinates). Consider an $\mathscr{L}_0(\emptyset)$-definable subgroup $W \leq M^{2n}$, connected of Morley rank n (in T_0), such that $\pi_1(A) = \pi_2(A) = M^n$. Call such a W a *correspondence*. For generic $(\overline{a}, \overline{b}) \in W$ this means that $\mathrm{d}_0(\overline{a}) = \mathrm{d}_0(\overline{b}) = n$ and $\overline{a}, \overline{b}$ are inter-$\mathscr{L}_0$-algebraic. Now, as T_0 is a locally modular strongly minimal group, its pre-geometry is determined by the field of definable quasi-endomorphisms of $M \models T_0$ which has to be a finite field $\mathbb{F}_q$ by ω-categoricity (see [19, Section 4.5] for details on locally modular groups). Thus, on the quotient M/M_0, the correspondence W coincides with the graph of some $\Gamma_W \in GL_n(\mathbb{F}_q)$. Since T_0 is a modular group, whenever $p \not\perp_k^a q$ for admissible types $p, q \in S^n(k)$ there is a correspondence W and $\overline{c}, \overline{c}' \in k$ such that for some $\overline{a} \models p$, $\overline{a}' \models q$ one has $(\overline{a} - \overline{c}, \overline{a}' - \overline{c}') \in W$. As $\pi_1(W) = M^n$ and $k \models T_0$, w.l.o.g. $\overline{c} = (0, \ldots, 0)$. For $p \in S^n(B)$ admissible we put

$$\mathrm{Tstab}_W(p) := \{\overline{c} \mid \text{There are } \overline{a}, \overline{a}' \models p \mid \overline{c} \text{ with } (\overline{a}, \overline{a}' - \overline{c}) \in W\}.$$

Modulo M_0, this just means $\overline{a}' = \Gamma_W \overline{a} + \overline{c}$. Because p is a definable type, $\mathrm{Tstab}_W(p)$ is definable over B. It is a twisted (slightly coarsened) stabiliser.

Moreover, if p is admissible and $(\overline{a}_1, \overline{a}_2) \models p^{(2)}$, we put $\Delta p := \mathrm{stp}(\overline{a}_1 - \overline{a}_2/\mathrm{Cb}(p))$.

LEMMA 7.3. *Suppose p is admissible, with infinite $\mathrm{Tstab}_W(p)$ for some correspondence W as above. Then $\mathrm{stab}(p)$ is strongly minimal with generic type Δp. Moreover, Δp is admissible.*

PROOF. $\mathrm{Tstab}_W(p)$ is infinite if and only if it has U-rank 1 (since $\mathrm{U}(\mathrm{Tstab}_W(q)) \leq \mathrm{U}(q)$ holds for all types q, as in the case of non-twisted stabilisers). Let $\overline{c}_1, \ldots, \overline{c}_N$ be realisations of $\mathrm{Tstab}_W(p)$, independent and generic over $B\overline{a}$, where p is over B, $\overline{a} \models p{\restriction}B$ and N the order of Γ_W in $GL_n(\mathbb{F}_q)$. Now put $\overline{a}_1 := \overline{a}$ and choose inductively $\overline{a}_{i+1} \models p$ satisfying $(\overline{a}_i, \overline{a}_{i+1} - \overline{c}_i) \in W$. Then, the following holds:

- $\overline{a}_2 \models p \mid \overline{c}_1$ (in particular $\overline{a}_2 \underset{B}{\perp} \overline{c}_1$),

- $\overline{a}_1 \underset{B}{\perp} \overline{a}_2$ (since $U(\overline{a}_1\overline{a}_2/B) = U(\overline{a}_2\overline{c}_1/B) = U(\overline{a}_1\overline{c}_1/B) = 2$).

In the same way we get

- $\overline{a}_{i+1} \models p \mid \overline{c}_1 \dots \overline{c}_i$,
- $\overline{a}_1, \dots, \overline{a}_{N+1}$ is a Morley sequence in p.

Modulo M_0, we have

$$\overline{a}_{N+1} = \Gamma_W \overline{a}_N + \overline{c}_N = \cdots = \Gamma_W^N \overline{a}_1 + \Gamma_W^{N-1}\overline{c}_1 + \cdots + \Gamma_W c_{N-1} + c_N,$$

Since $\Gamma_W^N = \mathrm{id}$ we get that $\overline{a}_{N+1} = \overline{a}_1 + \tilde{c} \models p \mid \tilde{c}$ for some $\tilde{c}$ algebraic over $\overline{c}_1 \dots \overline{c}_N$. As $2 = U(\overline{a}_1\overline{a}_{N+1}/B) = U(\overline{a}_{N+1}\tilde{c}/B)$ and $\overline{a}_{N+1} \underset{B}{\perp} \tilde{c}$, one has $U(\tilde{c}/B) = 1$. On the other hand, since $\overline{a}_1 \underset{B}{\perp} \overline{a}_{N+1}$, we have $\tilde{c} \models \Delta p$. Now, everything follows easily. $\dashv$

We can gather all this in the following:

PROPOSITION 7.4. *Let p be an admissible type.*

(1) *p is trivial iff $U(\Delta p) = 2$ iff $\mathrm{stab}(p)$ is finite.*

(2) *p is non-trivial iff $U(\Delta p) = 1$ (and Δp is admissible) iff $\mathrm{stab}(p)$ is strongly minimal (with generic type Δp).*

(3) *In case (2) Δp is non-orthogonal to p. It is a modular representative in the non-orthogonality class of p defined over the same parameters as p (its existence is promised by a general result of Hrushovski, see e.g., [19, Prop. 5.2.1]), and p is generic for a coset of $\mathrm{stab}(p)$.*

PROOF. Assume that p is based on $B = \mathrm{cl}(B)$. If $(\overline{a}_1, \overline{a}_2) \models p^{(2)}$, then $\overline{a}_1 - \overline{a}_2$ and $\overline{a}_2$ are interalgebraic (in T_ω) over $B\overline{a}_1$. Therefore $p \not\perp \Delta p$ and $U(\Delta p) \in \{1, 2\}$.

We now show (2). If $U(\Delta p) = 1$, then (by rank considerations) $\Delta p \vdash \mathrm{stab}(p)$, whence $\mathrm{stab}(p)$ is infinite, and so strongly minimal with generic type Δp by Lemma 7.3. In the other direction, if $\mathrm{stab}(p)$ is strongly minimal, Δp is admissible of U-rank 1 by the same lemma. To prove the other equivalence, note that if $\mathrm{stab}(p)$ is strongly minimal (with generic Δp), then p is non-orthogonal to a group generic, so can not be trivial. On the other hand, if p is non-trivial, there exist $\overline{a}_0, \dots, \overline{a}_{m-1}$ pairwise independent solutions of p which are dependent. We suppose that m is minimal with these properties (in fact $m = 3$ or $m = 4$, but we don't use this). Thus, $\overline{a}_0$ and $\overline{a}_1$ are interalgebraic over $B\overline{a}_2 \dots \overline{a}_{m-1}$. It follows from Lemma 4.11 that $(\overline{a}_0, \overline{a}_1 - \overline{c}) \in W$ for some correspondence W and some $\overline{c} \in \mathrm{acl}(B\overline{a}_2 \dots \overline{a}_{m-1}) \setminus B$. Since $\overline{a}_0 \underset{B}{\perp} \overline{a}_1$ and $\overline{c} \underset{B}{\perp} \overline{a}_i$ for $i = 0, 1$ we have $\overline{c} \models \mathrm{Tstab}_W(p)$ and $U(\mathrm{Tstab}_W(p)) \geq 1$, so $\mathrm{stab}(p)$ is strongly minimal by Lemma 7.3.

This finishes the proof of (2). Using $U(\Delta p) \in \{1, 2\}$, (1) follows from (2).

Finally, observe that the locally modular geometry of Δp is in fact modular, as Δp is the generic type of a group. This finishes (3). $\dashv$

We want to stress that not all admissible types are modular, as is shown by the following example.

EXAMPLE. Let T_0 be the theory of an infinite $\mathbb{F}_2$-vector space and both T_1 and T_2 the theory of an infinite $\mathbb{F}_4$-vector space, where the prime model is named by constants. In this situation $\langle \cdot \rangle = \mathrm{cl}$ holds, since $2\mathrm{d}_i(\overline{a}/B) \geq \mathrm{d}_0(\overline{a}/B)$ for any acl_i-closed set B ($i = 1, 2$), but we do not use this. Consider the admissible type p over $k := \langle \emptyset \rangle$ given by the equations $x = \lambda_1 *_1 y$ and $x = \lambda_2 *_2 y + c$, where

$\lambda_i \in \mathbb{F}_4 \setminus \mathbb{F}_2$, $0 \neq c \in \mathrm{acl}_2(\emptyset)$ and $*_i$ denotes the multiplication in the sense of T_i. This type is non-trivial with Δp given by the equations $x = \lambda_1 *_1 y$ and $x = \lambda_2 *_2 y$. Now we suppose, for contradiction, that p is modular. Then, p and Δp are not almost orthogonal over k (see [19, Cor. 2.5.5]) and thus one can find $x_\Delta y_\Delta \models \Delta p$ and $x_0 y_0 \models p$ which are $\mathbb{F}_2$-interdefinable over k, i.e., $\begin{pmatrix} x_\Delta \\ y_\Delta \end{pmatrix} = \Gamma \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} + \begin{pmatrix} a \\ b \end{pmatrix}$ for some $\Gamma \in GL_2(\mathbb{F}_2)$ and $a, b \in k$.

Recall that $\lambda_i^3 = 1, 1 + \lambda_i = \lambda_i^2$ and $1 + \lambda_i^2 = \lambda_i$. For simplicity we will write $\lambda_i z$ instead of $\lambda_i *_i z$. This should not lead to confusions.

Case $\Gamma = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. From $x_\Delta = y_0 + a$ and $y_\Delta = x_0 + b$ we deduce

$$\lambda_1 a = \lambda_1(x_\Delta + y_0) = \lambda_1^2 y_\Delta + x_0 \lambda_1^2 y_\Delta + y_\Delta + b = \lambda_1 y_\Delta + b = x_\Delta + b.$$

Thus, $x_\Delta \in k$, a contradiction. Similar computations show that $\Gamma = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ and $\Gamma = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ are impossible.

Case $\Gamma = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$. One first shows that $a = \lambda_1 b$ and $a = \lambda_2 b + \tilde{c}$ hold, where $0 \neq \lambda_2 c =: \tilde{c} \in \mathrm{acl}_2(\emptyset)$. It is easy (albeit somewhat lengthy) to see that $\mathrm{tp}(ab/\tilde{c})$ is admissible (show that $d_0(ab/\tilde{c}) = 2$, and $d_1(A/\tilde{c}) = d_2(A/\tilde{c}) = 1$ for all $\{0, \tilde{c}\} \subsetneq A \subseteq \mathrm{acl}_0(ab\tilde{c}) = \langle ab\tilde{c} \rangle_{\mathbb{F}_2}$). On the other hand, since $\{0, \tilde{c}\} \leq k = \langle \tilde{c} \rangle$, such $a, b \in k$ cannot exist, e.g., by induction using the hierarchies $\langle \tilde{c} \rangle_i^n$.

The cases $\Gamma = \mathrm{id}$ and $\Gamma = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ are treated in the same way. $\dashv$

We now present a general way to construct almost orthogonal admissible types. Let $p(\overline{x}) \in S^n(k)$ be admissible ($k = \mathrm{cl}(k)$), and $\overline{c} \in K^*$ an n-tuple. We define a new type $p_{\overline{c}}(\overline{x})$, an admissible type over $l := \mathrm{cl}(k\overline{c})$, as follows: Choose $\varphi_i(\overline{x}, \overline{b}_i) \in \mathscr{L}_i$ such that $p{\restriction}_{\mathscr{L}_i}$ is its unique generic type over k. Now let p_1 be the non-forking extension of $p{\restriction}_{\mathscr{L}_1}$ to l (an $\mathscr{L}_1$-type) and let $p_{2,\overline{c}}$ be the unique generic $\mathscr{L}_2$-type over l contained in $\varphi_2(\overline{x} - \overline{c}, \overline{b}_2)$. It is routine to check that the $\mathscr{L}$-type $p_{\overline{c}}$ over l which we get applying Lemma 3.10 to p_1 and $p_{2,\overline{c}}$—namely $p_1 \cup p_{2,\overline{c}}$ plus "strong"—is admissible.

LEMMA 7.5. *Let $p \in S^n(k)$ be admissible and $\overline{c}$ a d-generic n-tuple over k. Then $p \perp_{k\overline{c}}^a p_{\overline{c}}$.*

PROOF. Suppose otherwise. Then, by Remark 6.9, there are $\overline{a} \models p \mid k\overline{c}$ and $\overline{a}' \models p_{\overline{c}}$ such that $\overline{a}$ and $\overline{a}'$ are inter-$\mathscr{L}_0$-algebraic over $k\overline{c}$. Put $A := \mathrm{acl}_0(k\overline{a})$ and $A' := \mathrm{acl}_0(k\overline{a}')$. One has

$(*)$ $A \underset{k}{\overset{i}{\bigcup}} \overline{c}$, for $i = 1, 2$,

$(**)$ $A' \underset{k}{\overset{1}{\bigcup}} \overline{c}$ and $A' \underset{k}{\overset{2}{\not\bigcup}} \overline{c}$.

$(*)$ is clear by construction, as is the first part of $(**)$, since $\varphi_1(\overline{x}, \overline{b}_1)$ does not $\mathscr{L}_1$-fork over k. On the other hand, $\varphi_2(\overline{x} - \overline{c}, \overline{b}_2)$ does $\mathscr{L}_2$-fork over k. This can be seen as follows. Take $\overline{a}$ in $\varphi(\overline{x}, \overline{b}_2)$, generic over l. Then, $\mathrm{MR}_2(\overline{c}/k, \overline{a}) = \mathrm{MR}_2(\overline{a} + \overline{c}/k, \overline{a})$, as $\overline{c}$ and $\overline{a} + \overline{c}$ are interdefinable over $k, \overline{a}$. Thus, $\overline{a} + \overline{c}$ is an $\mathscr{L}_2$-generic tuple over $k, \overline{a}$, and so over k as well. But $\mathrm{MR}_2(\varphi_2(\overline{x} - \overline{c}, \overline{b}_2)) =$

$MR_2(\varphi_2(\overline{x}, \overline{b}_2)) < n$, since p is an admissible type. As $\overline{a} + \overline{c} \in \varphi_2(\overline{x} - \overline{c}, \overline{b}_2)$, this formula $\mathscr{L}_2$-forks over k. The second assertion in $(**)$ follows.

Now $(*)$ together with $(**)$ forces $A' \not\subseteq A$, so $0 < d_0(A'/A \cap A')$. As $\overline{c}$ realises the d-generic n-type over A, we have $d(A'/A) = d_i(A'/A)$ for $i = 0, 1, 2$, and by modularity of T_0, these are equal to $d_0(A'/A \cap A') > 0$ (and to $d_i(A'/A \cap A')$ for $i = 1, 2$). On the other hand, using $(**)$ and putting $C := acl_0(k\overline{c})$, we get $d_1(A'/A \cap A')d_1(A'/(A \cap A')C) < d_0(A'/(A \cap A')C) = d_0(A'/A \cap A')$ by primality of $acl_0(A'C)$ over C, unless $d_0(A'/A \cap A') = 0$. Contradiction. $\dashv$

In the presence of a type-definable group G a strong type p extending "$x \in G$" will be called a *subgroup type* if it is the generic type of a type-definable connected subgroup of G. If p is the generic type of a (right) coset of a type-definable connected subgroup of G, we will call it a *coset type*. Of course, these notions depend on the group G we are considering. We first make an easy observation. Let G be type-definable in a stable theory, and let p be a strong type (over $A = acl^{eq}(A)$, say) extending "$x \in G$". Then, $stab(p)$ (the left stabiliser) is defined as

$$\{g \in G \mid a \models p \mid A, g \text{ implies } g \cdot a \models p \mid A, g\},$$

where $\cdot$ is the multiplication in G. It is type-definable over A. We put $\Delta p := tp(a \cdot a'^{-1}/A)$, where $(a, a') \models p^{(2)}$. The following is standard:

 (i) p is a subgroup type iff $p \vdash stab(p)$ (iff $p = \Delta p$).
 (ii) p is a coset type iff $\Delta p \vdash stab(p)$.

Combining (ii) with (the proof of) Proposition 7.4, we obtain the following result, where the group G in question is given by a cartesian power of the $\mathscr{L}_0$-definable group structure over which we work.

COROLLARY 7.6. *Let p be an admissible type and p_i its restriction to $\mathscr{L}_i$. Then, p is non-trivial if and only if both p_1 and p_2 are coset types.* $\dashv$

One would like to know when the primitive extension given by some admissible type is trivial, if this type is moved within a definable family of such types. The following definability result is easy (and should be known).

LEMMA 7.7. *Let T be the theory of an uncountably categorical group G and let $\varphi(x, \overline{u})$ be a formula such that whenever $\varphi(x, \overline{b})$ is consistent, then $MRD(\varphi(x, \overline{b})) = (m, 1)$, and denote $p_{\overline{b}}$ the unique generic type of $\varphi(x, \overline{b})$. The set of all $\overline{b}$ such that $p_{\overline{b}}$ is a subgroup type is then definable, as is the set of all $\overline{b}$ such that $p_{\overline{b}}$ is a coset type.*

PROOF. The proof uses the observations (i) and (ii), transformed into definable conditions in the particular situation of the lemma. Recall that in T, Morley rank is finite and definable. Moreover it is equal to U-rank, whence additive. This enables us to single out those $\overline{b}$ where $p_{\overline{b}}$ is a subgroup type (a coset type, respectively). Consider $\psi_{sub}(x, y, \overline{u}) := \varphi(x \cdot y, \overline{u})$ and put

$$\theta_{sub} := \{\overline{b} \mid MR([\varphi(x, \overline{b}) \times \varphi(y, \overline{b})] \cap \psi_{sub}(x, y, \overline{b})) = 2m\}.$$

Then, θ_{sub} is a definable set and it describes the set of all $\overline{b}$ such that $p_{\overline{b}}$ is a subgroup type. To see this, take $\alpha \models p_{\overline{b}}$ and $\beta \models p_{\overline{b}} \mid \alpha$. If $\models \theta_{sub}(\overline{b})$ we have $\models \varphi(\alpha \cdot \beta, \overline{b})$. Thus $MR(\alpha \cdot \beta/\overline{b}) \leq m$. From this and $MR(\alpha, \beta/\overline{b}) = 2m$ we deduce $\alpha \underset{\overline{b}}{\perp} \alpha \cdot \beta$. The converse is easy.

One proceeds in the same way for the set of $\overline{b}$ giving rise to coset types $p_{\overline{b}}$. First put $\psi_{co}(x, y, z, \overline{u}) := \varphi(x \cdot y^{-1} \cdot z, \overline{u})$. The set in question is then given by $\theta_{co} := \{\overline{b} \mid \mathrm{MR}([\varphi(x, \overline{b}) \times \varphi(y, \overline{b}) \times \varphi(z, \overline{b})] \cap \psi_{co}(x, y, z, \overline{b})) = 3m\}$. ⊣

PROPOSITION 7.8. *Suppose that T_1 and T_2 have DMP. Then the following holds*:

(1) *T_ω has DMP for sets of finite Morley rank.*
(2) *There is a countable set of pairs of $\mathscr{L}$-formulas $\{\varphi_i(x, y), \theta_i(y)\}$ such that*
  - *$\varphi_i(x, b)$ is strongly minimal and its generic type is admissible whenever $\models \theta_i(b)$.*
  - *For every admissible type p there exists i and $b \models \theta_i(y)$ such that p is the generic type of $\varphi_i(x, b)$.*
  - *For every i and b, if $\models \theta_i(b)$ holds and $\varphi_i(x, b)$ has a trivial (locally projective) geometry, then $\varphi_i(x, b')$ has a trivial (locally projective) geometry for every $b' \models \theta_i(y)$.*
  - *For every $i \neq j$ and every $b_i \models \theta_i(y)$, $b_j \models \theta_j(y)$ one has $\varphi_i(x, b_i) \perp \varphi_j(x, b_j)$.*

PROOF. We leave the details as an exercise to the reader. Just note that Corollary 7.6 with Lemma 7.7 allow to definably distinguish trivial admissible types from locally projective ones. ⊣

We refer to Appendix A for detailed discussion of how strongly minimal dimensions can be coded in T_ω.

LEMMA 7.9. *Let p be a locally projective admissible type based on $k = \mathrm{cl}(k)$. For $i = 1, 2$, let N_i be the $\mathscr{L}_i(k)$-definable cosets of subgroups G_i of M^n such that $p_i := p{\restriction}_{\mathscr{L}_i}$ is the unique generic of N_i (as promised by Corollary 7.6). Then $N_1 \cap N_2$ is strongly minimal with generic type p.*

PROOF. Since we can take larger k if necessary, it is sufficient to show that p is the unique non-algebraic type over k. Let $\overline{b} \in N_1 \cap N_2$ with $\overline{b} \notin k$, so $\overline{b}$ is not algebraic over k. Suppose that $N_i = G_i + \overline{c}_i$. By (the proof of) Proposition 6.8 we can choose $\mathscr{L}_i$-formulas $\varphi_i \subseteq G_i$ such that p is the unique generic of $\varphi_1 \wedge \varphi_2$. Now choose $\overline{a} \models \Delta p \mid \mathrm{cl}(k\overline{b})$. By $\mathscr{L}_i$-genericity, $\overline{a} + \overline{b} \models p_i \mid \mathrm{cl}(\overline{b})$ for $i = 1, 2$. Thus, $\overline{a} + \overline{b} \models p \mid \mathrm{cl}(k\overline{b})$. Now, since $\overline{b} \notin \mathrm{acl}_{T_\omega}(k)$ and $\mathrm{U}(\overline{a}/k\overline{b}) = 1$, we must have $\mathrm{U}(\overline{a}, \overline{a} + \overline{b}/k) = \mathrm{U}(\overline{a}, \overline{b}/k) \geq 2$, so $\overline{a} \underset{k}{\smile} \overline{a} + \overline{b}$. So, as $\overline{a} + \overline{b}$ is generic, $\overline{b} = (\overline{a} + \overline{b}) - \overline{a}$ is $\mathscr{L}_i$-generic in both N_1 and N_2, showing that $\overline{b} \models p$. ⊣

Suppose $p = \Delta p$ is an admissible type over the (finitely generated) fusion $k = \mathrm{cl}(k)$. Then, by Proposition 6.8, $G := \mathrm{stab}(p)$ is a strongly minimal totally categorical modular group. Thus, its geometry is given by its skew-field $QE(G)$ of quasi-endomorphisms. By omega-categoricity, $QE(G)$ is in fact a finite field. If $\mathbb{F}_q$ is the field of quasi-endomorphisms of $M \models T_0$, then clearly $QE(G)$ embeds into $Mat_n(\mathbb{F}_q)$. We give two examples to illustrate what kind of situations can occur.

EXAMPLES. (1) Let T_0 be the theory of vector spaces over $\mathbb{F}_2$. For $i = 1, 2$, let T_i be the theory of vector spaces over $F_i := \overline{\mathbb{F}_2}$. As in a previous example we choose primitive third roots of unity $\lambda_i \in F_i$. The admissible type p is given by $x = \lambda_1 *_1 y \wedge x = \lambda_2 *_2 y$, where one requires x, y $\mathbb{F}_2$-independent over $\emptyset$. G is given by $\{(x, y) \mid x = \lambda_i *_i y$ for $i = 1, 2\}$. Then $\mathbb{F}_4 \subseteq QE(G)$, since $(\lambda^2, \lambda): G \to G, (x, y) \mapsto (\lambda_1^2 *_1 x, \lambda_1 *_1 y)$ is a quasi-endomorphism of order 3.

(2) Let T_i be the theory of infinite vector spaces over F_i, where $F_0 := \mathbb{F}_4$, $F_1 := \overline{\mathbb{F}_2}$ and where F_2 is a skew-field containing $\mathbb{F}_4$ such that there is $\lambda_2 \in F_2$ and $\lambda \in \mathbb{F}_4$ such that $\lambda_2 \lambda \lambda_2^{-1} \notin \{\lambda, \lambda^{-1}\}$, i.e., F_2 does not normalise $\mathbb{F}_4$. Such skew-fields exist, e.g., a quotient field of $\mathbb{F}_4 *_{\mathbb{F}_2} K$ for some proper extension K of $\mathbb{F}_2$.

Now choose $\lambda_1 \in F_1$, a primitive element for $\mathbb{F}_8$. Let p be given by $x = \lambda_i *_i y$, for $i = 1, 2$. It is not too hard (but a bit tedious) to see that $QE(\mathrm{stab}(p)) = \mathbb{F}_2 \subsetneq \mathbb{F}_4$.

With the collapse already in mind, we should first have a look at the easiest fusion context, which is the case where $T_i =$ vector spaces over F_i, where F_1 and F_2 are skew-field extensions of $F_0 = \mathbb{F}_q$. More generally we can consider two strongly minimal modular expansions of T_0. Recall that we restricted our context to those T_0 which are theories of a (modular) totally categorical group.

LEMMA 7.10. T_ω is 1-based if and only if both T_1 and T_2 are 1-based.

PROOF. Note that since all theories we consider contain T_0 as a reduct, they are theories of a group. Thus an expansion is an abelian structure-exactly if it is 1-based. Now, if both T_1 and T_2 are 1-based, so is every completion of $T_1 \cup T_2$, see [11]. In particular, T_ω is 1-based. For the converse, it suffices to note that a reduct of an abelian structure stays abelian, provided the group law persists in the reduct (see [19, Prop. 4.6.4]). ⊣

DEFINITION 7.11. The situation described in the preceding lemma will be referred to as the *abelian fusion context*.

PROPOSITION 7.12. *In the abelian fusion context, we have the following*:
(1) T_ω *is non-multidimensional.*
(2) *In T_ω, no admissible type is trivial.*

PROOF. In an abelian structure, every type is non-orthogonal to a group generic which does not fork over $\emptyset$. This proves both (1) and (2).

Note though that we can deduce a more concrete argument from the analysis of primitive extensions we already presented. If $\varphi_i(\overline{x}, \overline{b})$ are stationary $\mathscr{L}_i$-formulas giving rise to a primitive extension l/k (when we apply Lemma 3.10 to the generic types p_i of φ_i over k), then since the T_i are abelian, these are just generic types of cosets of $\mathrm{acl}_i(\emptyset)$-definable subgroups H_i of M^n. If p is the admissible type given by the p_i, then Δp is the admissible type corresponding to H_1 and H_2, and it gives rise to a primitive extension of $\langle \emptyset \rangle$. ⊣

In fact, there is a general result linking the two notions.

PROPOSITION 7.13. *If in T_ω there exists an admissible type with a trivial pre-geometry, then T_ω is multidimensional.*

PROOF. Let p be an admissible type (defined over k) whose associated geometry is trivial. Let $(\overline{c}_i), i \in I$ be a long sequence of d-generic tuples, independent over k. It follows immediately from Lemma 7.5 that for $i \neq j$ $p_{\overline{c}_i} \perp p_{\overline{c}_j}$, since almost orthogonality and orthogonality are the same for trivial types. We conclude that there exists an unbounded number of non-orthogonality classes of types, which implies multidimensionality. ⊣

We have seen that the abelian fusion context implies the non-multidimensionality of T_ω which in turn implies that no admissible type is trivial. It seems to be a rather

intricate issue to decide if these implications are strict. An answer to this question will rely on the representability of certain (finite) matroids inside T_i.

§**8. The collapse in the abelian fusion context.** In this section we show that the collapse can be done in the abelian fusion context (and more generally in a 1-based setting). We give a self contained direct proof of this, and leave it to Appendix A to outline briefly how this fits into the general context of envelopes as developed in [12].

Recall that in the abelian fusion context we treat the case of two 1-based expansions T_i of the theory T_0 of a (not necessarily pure) totally categorical group. For convenience we suppose that $\mathrm{dcl}_i(\emptyset) = \mathrm{acl}_i(\emptyset)$ for $i \in \{0,1,2\}$. Of course this can be achieved by adding constants to the language. Inspecting the proof of Proposition 7.12 (and with the notation therein) we see that every admissible type is nonorthogonal to an admissible type p such that its restriction p_i to $\mathscr{L}_i$ is a subgroup type based on $\emptyset$, so p_i is the generic type of some connected $\mathscr{L}_i(\emptyset)$-definable subgroup G_i of $(V_0,+)^{n_p}$, where $(V_0,+)$ is the underlying $\mathscr{L}_0$-definable group over which we are working. Thus, $p(\overline{x})$ is the generic type of the s.m. group $G_1 \cap G_2$, and it can be isolated by the quantifier free formula $G_1(\overline{x}) \wedge G_2(\overline{x}) \wedge \ulcorner \mathrm{d}_0(\overline{x}) = n_p \urcorner$.

Let $\mathscr{D}$ be the (bounded!) set of all admissible types of this form. From the discussion in Section 7 we know that for $p, q \in \mathscr{D}$ we have $p \not\perp q$ if and only if there is a correspondence W and an $\mathscr{L}_0$-generic $(\overline{a}, \overline{a}')$ of W such that $\overline{a} \models p$ and $\overline{a}' \models q$.

Now consider any function $\mu \colon \mathscr{D} \to \mathbb{N} \cup \{\infty\}$ such that μ is invariant under non-orthogonality. Put

$$\tilde{\mathscr{C}}_0^{\mu} := \{M \in \tilde{\mathscr{C}}_0 \mid \dim_M(p) \le \mu(p) \text{ for all } p \in \mathscr{D}\}.$$

This just means that a structure in $\tilde{\mathscr{C}}_0^{\mu}$ contains at most $\mu(p)$ independent solutions of p for all $p \in \mathscr{D}$. This class is elementary. Indeed, using e.g., Lemma 4.4, it is easy to see that $\mathrm{d}_0(p^M) = n_p \cdot \dim_M(p)$ which gives $\dim_M(p)$ in a rather explicit form, showing that $\dim_M(p) \le n$ is a quantifier free definable condition for all n.

As before we denote by $\mathscr{C}_0^{\mu}$ the subclass of finitely generated structures (in the sense of $\langle \cdot \rangle$).

LEMMA 8.1 (Economic Amalgamation). $(\mathscr{C}_0^{\mu}, \le)$ *has the amalgamation property.*

PROOF. Using the decomposition of finitely generated extensions (4.7) and identifying d-bases first, it suffices to treat the following case: l is a primitive extension of k and $k \le m$ is arbitrary. We have to show that if $l, m \in \tilde{\mathscr{C}}_0^{\mu}$, there exists (in $\tilde{\mathscr{C}}_0^{\mu}$) an amalgam of l with m over k. Our goal is to take advantage of the geometrical analysis of T_ω carried out in the previous section so we may assume w.l.o.g. that k is strongly embedded in K^*, some big saturated model of T_ω. Let p be an admissible type associated to l/k.

There are now two cases:

Case 1: $p \not\perp_k^a \Delta p$. This is meaningful, by the assumption that we are working in a model of T_ω (and $\not\perp$ is used with respect to this theory).

W.l.o.g. we may assume that $p = \Delta p$. If $\dim_k(p) < \dim_m(p)$ we can find a copy of l in m over k and m is itself an amalgam (an economic one). If $\dim_k(p) = \dim_m(p)$, then we take $m' := l \otimes_k m$ as amalgam. Observe that if $K \le L$ is a primitive

extension in $\tilde{\mathscr{C}}_0^\mu$, given by a solution of p, then $\dim_L(p') = \dim_K(p') + 1$ (if these numbers are finite) for all $p' \in \mathscr{D}$ with $p' \not\perp p$, and $\dim_L(q) = \dim_K(q)$ for all $q \in \mathscr{D}$ orthogonal to p (since they are all based on $\emptyset$). Using this for m' and m we get that $\dim_{m'}(p') = \dim_l(p')$ for all $p' \not\perp p$ and $\dim_{m'}(q) = \dim_m(q)$ for all $q \perp p$. Thus, if $l, m \in \mathscr{C}_0^\mu$ we obtain that the free amalgam m' is in the class $\mathscr{C}_0^\mu$.

Case 2: $p \perp_k^a \Delta p$. Since Δp is modular (and $p \not\perp \Delta p$), we have $p \not\perp_m^a \Delta p$ if and only if m contains a solution of p which is generic over k (this follows from a result of Hrushovski, see e.g., [19, 5.2.5.]). If this is the case, then there is a copy of l/k in m and we can amalgamate "economically". Otherwise, we take the free amalgam m' and nonetheless we end up with $\dim_{m'}(\Delta p) = \dim_m(\Delta p)$ (so clearly $\dim_{m'}(q) = \dim_m(q)$ for all $q \in \mathscr{D}$), thus staying in the class. $\dashv$

Since one can amalgamate in the class $(\mathscr{C}_0^\mu, \leq)$, there exists a generic model M^μ for this class. We now consider T^μ, the $\mathscr{L}$-theory of M^μ. For every choice of a μ-function we will see that M^μ is a saturated model of T^μ. This is achieved by axiomatising T^μ via T'^μ and identifying rich structures in $\tilde{\mathscr{C}}_0^\mu$ with ω-saturated models of T'^μ, exactly as we did in the free case.

T'^μ is axiomatised as follows (in (3) we think of $p(\overline{x}) \in \mathscr{D}$ given by $G_1(\overline{x}) \wedge G_2(\overline{x}) \wedge \ulcorner d_0(\overline{x}) = n_p \urcorner$ as explained above):

(1) $\mathrm{Th}(\tilde{\mathscr{C}}_0)$,
(2) $\dim(p) = \mu(p)$ for all $p \in \mathscr{D}$,
(3) $\forall \overline{y} \exists \overline{z} \{\overline{z} \in G_1 \wedge \overline{z} \in [\overline{y} + G_2]\}$ for all $p(\overline{x}) \in \mathscr{D}$.

We note that the axioms in (3) could be replaced by an axiom scheme expressing that "every s.m. affine space definable in a model has a point". The axioms given in (3) just do not mention all definable s.m. affine spaces. Instead, a certain choice of families of affine spaces is required to have a point, which turns out to be enough.

CLAIM. The rich structures in $\tilde{\mathscr{C}}_0^\mu$ are exactly the ω-saturated models of T'^μ. In particular, T'^μ equals T^μ and is complete.

We first show that ω-saturated models of T'^μ are rich structures in $\tilde{\mathscr{C}}_0^\mu$. Approximating the generic extension by parasitic ones in the class is an easy application of compactness and left to the reader. Now, the same reduction and case distinction as in the proof of Lemma 8.1 shows that the only difficulty appears in the case of a primitive extension given by an admissible type q with $q \perp_k^a \Delta q$ (we think of a primitive extension l/k given by q). We know that there are $\mathscr{L}_i$-definable subgroups G_i and tuples $\overline{c}_1, \overline{c}_2$ such that Δq is given by $(G_1(\overline{x}), G_2(\overline{x}))$, whereas q is given by $(\overline{c}_1 + G_1(\overline{x}), \overline{c}_2 + G_2(\overline{x}))$. Putting $\overline{c} := \overline{c}_2 - \overline{c}_1$ and passing to a type not almost orthogonal to q, we can assume that q is given by $(G_1(\overline{x}), \overline{c} + G_2(\overline{x}))$. But this is exactly the situation taken care of in $T'^\mu(3)$, so there is a solution $\overline{a}$ to $(G_1, \overline{c} + G_2)$, a strongly minimal set by Lemma 7.9. Now, $\overline{a}$ is automatically generic, since every solution turns q modular. This shows that an ω-saturated model M of T'^μ is rich.

Now assume that $M \in \tilde{\mathscr{C}}_0^\mu$ is rich. We show that $M \models T'^\mu$. As in the free fusion this will suffice to establish the claim. Clearly, M satisfies (1) and (2). So suppose we have G_1, G_2 and $\overline{c} \in M$ as in (3). Put $k := \mathrm{cl}(\overline{c}) \leq M$, and let q be the admissible type given by (the generic types of) G_1 and $\overline{c} + G_2$. The proof of Lemma 8.1 shows that if $q \perp_k^a \Delta q$, then $\langle k\overline{a} \rangle \in \tilde{\mathscr{C}}_0^\mu$, where $\overline{a} \models q$, and so by richness there is an isomorphic copy of $\overline{a}$ in M. If $q \not\perp_k^a \Delta q$ there is a correspondence W and $\overline{b} \in k$

such that $(\overline{a}', \overline{a} - \overline{b}) \in W$ for some $\overline{a} \models q$, $\overline{a}' \models \Delta q$. Observe that both $G_1 \wedge G_2$ and $G_1 \wedge [\overline{c} + G_2]$ are closed under $f(\overline{x}, \overline{y}, \overline{z}) := \overline{x} + \overline{y} - \overline{z}$. Of course, W is closed under f, too. Now choose $\overline{a}'_1, \overline{a}'_2$ independent generic solutions of Δq over k, in the sense of T_ω. As $\overline{a}'_3 := \overline{a}'_1 + \overline{a}'_2 \models \Delta q \mid k$ we deduce that for the corresponding $\overline{a}_i \models q$, for $i = 1, 2, 3$ one has $\models W(\overline{a}'_i, \overline{a}_i - \overline{b})$. Applying f we see that $(0, \tilde{a} - \overline{b}) \in W$ for some $\tilde{a} \in G_1 \wedge [\overline{c} + G_2]$. So $\tilde{a} \in \mathrm{acl}_0(\overline{b}) \subseteq k$ and $\tilde{a}$ is the desired solution.

We observe that, due to non-multidimensionality, no finite-to-one condition on the μ-function is needed to obtain saturation of the generic model, as opposed to [18], [17] and [3]. It follows that a counter-example in the style of [3, Section 4] cannot be constructed in this context. On the other hand, already in the case when two vector spaces (over skew-fields F_1 and F_2, respectively) are fused over equality, such a condition is needed since T_ω is then multidimensional. We note that In this case, the resulting fused theory is not 1-based. So, quite paradoxically, the fusion over the $\mathbb{F}_q$-vector space structure for some finite common subfield $\mathbb{F}_q$ of F_1 and F_2 is less complicated than the fusion over mere equality.

Note that the aforementioned counter-example in [3] arises from admissible types based on tuples $\overline{a}$ that are not strongly embedded in the universe.

THEOREM 8.2. *Suppose $\mu(p)$ is finite for all $p \in \mathscr{D}$. Then T^μ is a complete strongly minimal theory. Moreover, it is 1-based and model complete.*

PROOF. Completeness has already been shown, and 1-basedness holds for any completion of $T_1 \cup T_2$, as we mentioned earlier. By a theorem of Lindström, an $\aleph_1$-categorical theory admitting a $\forall\exists$-axiomatisation is model complete. Thus, since the axioms we gave for T'^μ are clearly $\forall\exists$, strong minimality has model completeness as a consequence. We note that this technique to show model-completeness of a collapsed theory was first used by Holland in [16].

We now consider M^μ, the countable ω-saturated model of T^μ. Let $M^\mu \preccurlyeq N$, with $n \in N \setminus M^\mu$. As in the free case, the fact that every ω-saturated model is rich means that the type of a strong tuple is determined by its quantifier free type. Thus, there is only one type of a d-generic element. We claim that n has to realise this type, i.e., $\mathrm{d}(n/M^\mu) = 1$. Else we could find $M^\mu \leq N_1 \leq \mathrm{cl}(M^\mu n)$ with N_1/M^μ primitive, given by some admissible type $\tilde{q}$. Choose $k \leq M^\mu$ finitely generated such that $\tilde{q}$ is based on k, and set $q := \tilde{q} {\restriction} k$. Since M^μ is rich, M^μ contains a solution of q. Thus, $\tilde{q}$ is modular, so non-orthogonal to $p := \Delta q$. One has $\mu(p) = \dim_{M^\mu}(p) < \dim_{N_1}(p)$, and so $N_1 \notin \tilde{\mathscr{C}}_0^\mu$, a contradiction. Thus we have seen that there is only one non-realised 1-type over M^μ, and strong minimality follows. $\dashv$

As in the case of the fusion over the infinite structureless set we can add the following:

REMARK 8.3. (1) *The expansions $T_i \subseteq T^\mu$ preserve Morley rank and Morley degree.*

(2) *Let $M \models T^\mu$ and $V_i \subseteq M^n$ be $\mathscr{L}_i$-definable without parameters. Assume V_i has empty intersection with every $\mathscr{L}_0$-definable hyperplane in M^n, and $\mathrm{MR}(V_1) + \mathrm{MR}(V_2) < n$. Then $V_1 \cap V_2 = \emptyset$. If Z is $\emptyset$-definable in $\mathscr{L}_1$ and in $\mathscr{L}_2$, then Z is $\emptyset$-definable in $\mathscr{L}_0$.*

(3) *The expansions $T_i \subseteq T^\mu$ are essential, and thus there is no maximal abelian strongly minimal structure of bounded exponent (in a countable language).*

PROOF. Every strongly minimal expansion is rank preserving. Now suppose that $\mathrm{MD}(\varphi_1(\overline{x}, \overline{b})) = 1$. For the sake of simplicity, we give the proof for $T_0 =$ vector spaces over $\mathbb{F}_q$. Shrinking $\varphi_1(\overline{x}, \overline{b})$ if necessary and using an $\mathbb{F}_q$-definable bijection, we can assume that $\models \varphi_1(\overline{a}/\overline{b})$ implies $\mathbb{F}_q$-genericity of $\overline{a}$ over $\overline{b}$. Moreover we suppose that $\overline{b}$ is strong. For $\overline{a} \in M \models T^\mu$ generic in $\varphi_1(\overline{x}, \overline{b})$ we have $\mathrm{d}_1(\overline{a}/\overline{b}) + \mathrm{d}_2(\overline{a}/\overline{b}) - \mathrm{d}_0(\overline{a}/\overline{b}) \geq \mathrm{d}(\overline{a}/\overline{b}) = \mathrm{d}_1(\overline{a}/\overline{b})$, the last equality holding since the expansion is rank preserving. Thus, $\overline{a}$ is $\mathscr{L}_2$-generic over $\overline{b}$ and $\overline{a}\overline{b}$ is strong. By quantifier elimination for strong subsets this means there is only one $\mathscr{L}$-generic type in $\varphi_1(\overline{x}, \overline{b})$. This shows (1).

The first part of (2) is just a consequence of $\delta(\overline{a}) \geq 0$. Now let $Z \subseteq M^n$ be $\mathscr{L}_i(\emptyset)$-definable for $i = 1, 2$. Clearly, Z is $\mathscr{L}_0$-definable if for every complete $\mathscr{L}_0$-type p_0 over $\emptyset$, either $Z \supseteq p_0$ or $Z^c \supseteq p_0$. Using definable bijections, it suffices to treat the case where p_0 is the $\mathscr{L}_0$-generic type of M^n. Suppose, for contradiction, that $Z \cap p_0 \neq \emptyset \neq Z^c \cap p_0$. W.l.o.g. $\mathrm{MR}(Z \cap p_0) =: m < n$. For $\overline{a} \in Z \cap p_0$ generic we then have $m = \mathrm{d}(\overline{a}) \leq \mathrm{d}_1(\overline{a}) + \mathrm{d}_2(\overline{a}) - \mathrm{d}_0(\overline{a}) = 2m - n < m$, a contradiction.

(3) is easy, and we leave it to the reader.   ⊣

We remark that—as we already mentioned in the previous section—the same results can be obtained if we only assume that there is a definable group in T_0 (and T_1, T_2 are 1-based). The crucial point in the above proof is the non-multidimensionality of T_ω. The axiom analogous to $T'^\mu(3)$ just says that "every affine strongly minimal space has a point". For a detailed exposition of this (in a more general setting) see [12].

In the following remark, we mention another generalisation, the proof of which entirely rests on our analysis of T_ω, the theory of the free fusion. Its content is made precise and explained in detail in [15], and it provides a *formal* reduction of the collapsing problem to the case where T_0 is (essentially) the theory of an infinite vector space over a finite field and the case where T_0 is the theory of an infinite structureless set (the case treated in the original fusion paper [17]).

REMARK 8.4. *Let (T_0, T_1, T_2) be a fusion context (having good control), and such that T_0 is $\aleph_0$-categorical and modular (it might be trivial here). Suppose there is an $\mathscr{L}_0$-interpretable set D_0' such that T_i', the theory of D_0' induced by $\mathscr{L}_i$, is s.m. for $i = 0, 1, 2$. Then, if one can find a s.m. collapse of T_ω', the free fusion in the context (T_0', T_1', T_2'), one can collapse T_ω, the free fusion in the original context (T_0, T_1, T_2), onto a s.m. theory, too.*   ⊣

Either way, we can now state the following result:

THEOREM 8.5. *Let (T_0, T_1, T_2) be a fusion context having good control, and assume that T_1 and T_2 are 1-based. If T_0 is trivial, we further assume that the T_i have the DMP. Then, the theory of the free fusion T_ω can be collapsed onto a strongly minimal theory expanding both T_1 and T_2, with their expansion sharing a common reduct which models T_0.*

PROOF. One can either use Remark 8.4 to reduce to the case done in [17] if T_0 is trivial, and to Theorem 8.2 in case T_0 is modular non-trivial.

As we already mentioned above, in case T_0 is trivial, one could also simply redo the proof from [17], and, in case T_0 is modular non-trivial, use that the corresponding

T_ω is non-multidimensional, and so give a proof similar to the one in the abelian fusion context (see [12]).   ⊣

EXAMPLE 8.6. Let F_1, F_2 be skew-field extensions of $F_0 := \mathbb{F}_q$, and let T_i be the theory of infinite F_i-vector spaces. Thus, we are in the abelian fusion context. Set $R := F_1 *_{F_0} F_2$, where $*_{F_0}$ denotes the coproduct of non-commutative rings over K_0 (we refer to [10] for background on skew-fields etc.). Clearly, every $M \in \tilde{\mathscr{E}}_0$ is an R-module. It is easy to see that for a d-generic and $0 \neq r \in R$ the element $r \cdot a$ is d-generic, too. One can use an appropriate $\mathbb{F}_q$-basis of R (see [10, p. 97]) to show that if a is generic over $k = \mathrm{cl}(k)$, then $\langle ka \rangle = k \oplus R$ as R-modules, and rather explicitly $\mathrm{d}(a/k \cup \{r \cdot a\}) = 0$. Thus, for a strongly minimal fusion T^μ, R naturally embeds into the skew-field of quasi-endomorphisms. This shows in particular that R does have a field of fractions. It is known by algebraic methods that R even has a universal field of fractions (combine Thm. 4.C. and Thm. 5.3.2. in [10]). Since our construction is very canonical, we conjecture that the field of quasi-endomorphisms of T^μ (any μ) coincides with the universal field of fractions of R promised algebraically. Let us finally remark that this argument does not need the collapsed fusion, since the forking geometry of the generic type of T_ω (which is locally modular regular) is given by a skew-field, too.

There is an interesting difference between the pure algebraic content of the previous example and the slightly more general (model-theoretic) abelian fusion context, though. The construction of a (universal) skew-field extension of $K_1 *_{K_0} K_2$ works equally well over any skew-field K_0, whereas in the abelian fusion context, ω-categoricity of T_0 seems crucial for the collapse. This is illustrated by the following example mentioned in the introduction of [17].

EXAMPLE. Let T_0 be the theory of $\mathbb{Q}$-vector spaces, with two $\mathbb{Q}$-independent elements c, d named as constants. Let T_1 be the theory of $\mathbb{Q}[i]$-vector spaces, with $i \cdot c = d$, T_2 the theory of vector spaces over $\mathbb{Q}(X)$, with $X \cdot c = d$.

In any $M \models T_1 \cup T_2$, $N := \ker(i - X)$ is a proper non-trivial $\mathbb{Q}$-subspace of M, since $c \in N$ and $d \notin N$. In particular, M is not strongly minimal.

REMARK 8.7. *First order logic is not the right framework in which one should consider fusion contexts over a non-$\aleph_0$-categorical T_0. Even the class $\tilde{\mathscr{E}}_0$ in which we were working all the time is not first order in this case. But this class can be axiomatised in $\mathscr{L}_{\omega_1,\omega}$, and the construction goes through practically unaltered. In the corresponding abelian fusion context, one can axiomatise a collapsed fusion that turns out to be a quasi-minimal excellent class, so in particular categorical in all uncountable cardinalities. To this aim, one has to work with $\mathscr{L}_{\omega_1,\omega}(Q)$-axioms stating that for all finite A and all parasitic types $p \in S(A)$, $|p(M)| = \aleph_0$ in a model M of the theory. For a smooth introduction to these issues, see e.g., [2].*

Added in proof: In the period after the submission of this paper the collapse of both the fusion over vector spaces and Poizat's field with an additive subgroup were announced by Baudisch, Martin-Pizarro and Ziegler [7]. Although their exposition differs significantly from ours, their work can be easily translated into the context in which we are working. Once this translation is done a proof that in T_ω orthogonality is equivalent to almost total orthogonality (see Appendix A below) follows easily from Lemmas 5.4 and 7.3 of [7]. With this result in hand, as we projected in

the appendix, the completion of the collapse is almost a formality. We refer the interested reader to [13] for the details. It is also worth noting that the results of the Berlin Group together with the results of this paper—see Remark 8.4—completely settle the question of the fusion over sublanguages for totally categorical T_0 and T_1, T_2 strongly minimal with DMP (with the case of the non-stable free fusions taken care of in [15]).

Appendix A. Envelopes and a view towards a collapse in general. The grail of the quest we set upon in this paper is, of course, the collapse of T_ω into a strongly minimal theory in full generality. This goal, valuable for the new strongly minimal sets and geometries it will give rise to is all the more interesting due to the many geometrical similarities it shares with Poizat's red-and-white fields in positive characteristic. Indeed, the geometrical analysis of the two structures is so similar that we believe that a successful collapse of the one will be a considerable step towards collapsing the other. Our aim in this appendix is to suggest an axiomatic framework which generalizes what we did in the abelian fusion context, and into which the two contexts fit. We believe that this framework is the right one for studying (and obtaining) the collapse for both. We discuss further the relations between the fusion and bicoloured fields in Appendix B.

In Section 8 we gave a full proof of the collapse of T_ω in the abelian fusion context. In this appendix we outline the concept of envelopes which is at the origin of the collapse in the abelian fusion context, and we give an overview of what we consider to be the main obstacle on the way to obtaining the full result.

Our approach to the problem of the collapse is closely related to the theory of smoothly approximable structures, as developed in [8] and more recently in [9], whose terminology we will be using throughout this appendix. To be more precise, our view is that the collapse consists in developing a theory of envelopes for T_ω. Indeed, this point of view justifies the use of the term 'collapse', as it consists of finding and axiomatising substructures (of a monster model) of T_ω which are strongly minimal and smoothly approximate it. We give some details of the suggested strategy.

Let T be an ω-stable theory in a countable language such that

(I) T has a unique (up to domination equivalence) generic 1-type p_ω of rank ω, and the remaining regular types are all strongly minimal.

(II) Every strongly minimal set of T is locally finite.

For most purposes, instead of (I) it is sufficient to require

(I)′ T has a unique regular type p_ω of rank ω (up to non-orthogonality), and the remaining regular types are all strongly minimal.

A set of (pairs of) $\mathscr{L}$-formulas $\{\varphi_i(x, y), \theta_i(y)\}_{i \in \mathbb{N}}$ is called a *complete system of codes* (for T) if the following holds:

(0) Whenever $\models \theta_i(b)$, the set $\varphi_i(x, b)$ is strongly minimal.

(1) For every strongly minimal type p in T there exists $i \in \mathbb{N}$ and b such that $\models \theta(b)$ and $p \not\perp \varphi_i(x, b)$.

(2) For every i and b, if $\models \theta_i(b)$ holds and $\varphi_i(x, b)$ has a trivial (locally projective) geometry, then $\varphi_i(x, b')$ has a trivial (locally projective) geometry for every $b' \models \theta_i(y)$.

(3) For every $i \neq j$ and every $b_i \models \theta_i(y), b_j \models \theta_j(y)$ one has $\varphi_i(x, b_i) \perp \varphi_j(x, b_j)$.

The complete system of codes is called *almost normalised* if for every non-orthogonality class $\mathbb{P}$ of strongly minimal types there are only finitely many b_l such that $\models \theta_i(b_l)$ and the generic type of $\varphi_i(x, b_l)$ is in $\mathbb{P}$. We then call any such type an *almost canonical representative* (for $\mathbb{P}$). If there is a unique such type for any non-orthogonality class $\mathbb{P}$, the system is called *normalised*, and the corresponding types are *canonical representatives*. The following lemma is just Proposition 7.8, and it will help to clarify the definition of envelopes in the present context:

LEMMA A.1. *Assume that T_1 and T_2 are both strongly minimal satisfying the assumptions given in 3.4. Assume in addition that both have DMP. Then there exists complete system of codes for T_ω, all instances of which are admissible types.* ⊣

A little more delicate is the next lemma:

LEMMA A.2. *Working in $\mathscr{L}(T_\omega)^{eq}$, there is a complete normalised system of codes for T_ω.* ⊣

Note that in order to obtain the stronger orthogonality assumption on code families $\{\varphi_i(x, y), \theta_i(y)\}$ we must pass to imaginary sorts, the problem being that—in general—we cannot hope to find, within the class of admissible types, a canonical representative of the non-orthogonality classes for those strongly minimal sets in T_ω with trivial geometry. Of course, as follows from Section 7, the code families concerning locally-projective sets can be almost normalised already on admissible types.

Now choose any function $\mu \colon \mathbb{N} \to \mathbb{N} \cup \{\infty\}$, and define:

DEFINITION A.3. Suppose that T is ω-stable and satisfies (I)′,(II), and suppose that $\{\varphi_i(x, y), \theta_i(y)\}$ is a complete normalised system of codes for T. Let $M \models T$ be a saturated model. A set $E \subseteq M$ is a *μ-envelope* if it is maximal satisfying the following requirements:

(1) $E = \mathrm{acl}(E)$.
(2) For all $i \in \mathbb{N}$ and every $b \in \theta_i(E)$, $\dim_E(p_i(x, b)) \leq \mu(i)$ (for $p_i(x, b)$ the generic type of $\varphi_i(x, b)$).

Roughly, our approach for tackling the problem of the collapse consists of the following stages:

(1) Show that for a suitable choice of the μ-function, μ-envelopes are homogeneous (in T_ω).
(2) For every μ as above, give an axiomatisation of the theory of μ-envelopes and thus show that being a μ-envelope is first order.

More precisely, since (unlike in the abelian fusion context treated in Section 8) it will most probably not be possible to axiomatise envelopes as such, the idea will be to use pseudo Morley sequences—a definable "good enough" approximation of Morley sequences (of length $\mu(i)$)—more or less in the spirit of the axiomatisation of [17]. Thus, our strategy would be better restated in terms of pseudo-μ-envelopes, namely for:

DEFINITION A.4. Let $M \models T$ be a saturated model, $\{\varphi_i(x, y), \theta_i(y)\}$ be a complete normalised system of codes. A set $E \subseteq M$ is a *pseudo-μ-envelope* if it is maximal satisfying the following requirements:

(1) $E = \mathrm{acl}(E)$.

(2) For all $i \in \mathbb{N}$ and every $b \in \theta_i(E)$,

$$\neg(\exists \bar{x}_0, \dots, \bar{x}_{\mu(i)})\psi_{\mu(i)+1}(\bar{x}_0, \dots, \bar{x}_{\mu(i)}, b)$$

(for some "good enough" approximation $\psi_{\mu(i)+1} \in p_b^{\otimes \mu(i)+1}(x, b)$ chosen in advance).

Quite clearly, working with μ-envelopes is easier than working with pseudo-envelopes. However, already in this setting we were not able to prove that if the set $S_\mu := \{i : \mu(i) < \infty\}$ is infinite, envelopes in T_ω are homogeneous (at this stage we were only able to prove homogeneity of envelopes under the assumption that T_0 is the theory of infinite vector spaces over $\mathbb{F}_p$—p a prime number—and that S_μ is finite). As we see it, this is the main obstacle towards collapsing T_ω. This last claim is partly justified by a thorough analysis of the theory of envelopes and smooth approximations carried out in [12]. Using the terminology thereof, we believe that the most important step towards collapsing T_ω (or indeed for red-and-white fields of positive characteristic, see the following appendix) would be proving that T_ω satisfies the assumption on the extension of types for each one of the formulas $\varphi_i(x, y)$ above. Explicitly, what needs to be shown is:

PROBLEM A.5. Let $\{\varphi_i(x, y), \theta_i(y)\}$ be a complete normalised system of codes, $M \models T_\omega$ a saturated model.
- For any $i \in \mathbb{N}$,(small) $B = \mathrm{acl}(B) \subseteq M$, $b \models \theta_i(y)$ and generic $a \models p_i(x, b)\!\restriction_B$, for any $j \in \mathbb{N}$ and $b' \in B_a$ such that $\models \theta_j(b')$ define $\Delta(b, b', i, j, B) := \dim_{B_a/B}(p_j(x, b'))$, where $B_a = \mathrm{acl}(Ba)$ and

$$\dim_{C/D} p : \max\{k : \exists(b_1 \dots, b_k \models p^{\otimes k})(b_i \subseteq C, \ b_i \cap D \neq b_i)\}.$$

- Define for all $i \in \mathbb{N}$ a function $O_i : \mathbb{N} \to \mathbb{N} \cup \{\infty\}$ by

$$O_i(j) := \max\{\Delta(b, b', i, j, B) : \ \models \theta_i(b), \models \theta_j(b'), B\}.$$

Is there some $i \in \mathbb{N}$ such that $O_i(j)$ is unbounded?

The above problem merits a few words of explanation. Of course, if $i = j$ and $b = b'$ then $\Delta(b, b', i, j, B) = 1$ and if $b' \in B$ and we are not in the previous case then $p_i(x, b) \perp p_j(x, b')$ and therefore $\Delta(b, b', i, j, B) = 0$. The interesting case is therefore when $b' \in B_a \setminus B$, in which case $\Delta(b, b', i, j, B)$ reduces to $\dim_{B_a}(p_j(x, b))$. Under the assumption that $b' \in \mathrm{acl}(ba)$ it is proved in [12] that $\dim_{B_a}(p_j(x, b')) = \dim_{\mathrm{acl}(ba)}(p_j(x, b'))$, and we are therefore in safe territory. Defining:

DEFINITION A.6. Let $b \models \theta_i$. Say that $p_i(x, b)$ *admits an obstruction*, if there exists some $B = \mathrm{acl}(B) \subseteq M$ such that for every $a \models p\!\restriction_B$ there exists some $j \in \mathbb{N}$ and $d \in B_a \setminus (B \cup \mathrm{acl}(ba))$ satisfying θ_j such that $\dim_{B_a}(p_j(x, d)) > 0$, and call $\dim_{B_a}(p_j(x, d))$ the size of the obstruction.

So Problem A.5 reduces to:

PROBLEM A.7. Is there a definable family of strongly minimal sets in T_ω witnessing obstructions of unbounded size?

It is worth noting that, e.g., in the fusion over T_0 with trivial geometry the corresponding theory T_ω admits only bounded obstructions. Generalising the notion of a type admitting no obstruction we suggest the following:

DEFINITION A.8. Let $p, q \in S(\mathfrak{C})$ be strongly minimal types. p is *almost totally orthogonal* to q if there exists $n(q) \in \mathbb{N}$ such that for every $B \supseteq Cb(p)$, every $m \in \mathbb{N}$ and every $a \models p^{\otimes m} \restriction_B$, setting $B_a = \text{acl}(Ba)$, $\dim_{B_a/B} q \leq n(q)$. If $n(q) = 0$ we say that p is *totally orthogonal* to q.

We note that total orthogonality agrees with orthogonality whenever the latter is defined. I.e., if $p \perp q$ and $Cb(p), Cb(q) \subseteq B$ then clearly for all $a \models p \restriction_B$, $\dim_{B_a} q = \dim_B q$, so $\dim_{B_a/B} q = 0$, with the other direction being obvious. The same is clearly true of almost total orthogonality. It follows that total orthogonality is of interest only when $Cb(p) \subseteq B$ and $q \perp B$ (using the transitivity of non-orthogonality for strongly minimal sets). Thus, if p and q are strongly minimal sets based on $\emptyset$ and $p \perp q$, they are totally orthogonal to each other. In fact, in that case we have an even stronger property, if $p \perp q$ then for every set B and every $a \models p \restriction_B$, $\dim_{B_a} q = \dim_B q$. In particular, the same is true in the abelian fusion context (see the proof of the Economic Amalgamation Lemma 8.1). As we have already mentioned in the case that T_0 is trivial, in the corresponding theory T_ω orthogonality (of strongly minimal types) is equivalent to almost total orthogonality, and if T_0 is the theory of infinite vector spaces over $\mathbb{F}_p$, then orthogonality is equivalent to almost total orthogonality. So it is clear that this notion is meaningful even in theories with a fairly complicated structure. It would be interesting to characterise, say, those ω-stable theories T in which (almost) total orthogonality is equivalent to orthogonality, and possibly even more relevant to the present problem, those theories in which almost total orthogonality is uniform, i.e., $n(q)$ in the above definition can be chosen independently of q.

Having said all of the above, we now show how the collapse in the abelian fusion context can be seen as a special case of a general result about envelopes. Axiomatically, the situation is as follows:

Suppose the complete ω-stable theory T satisfies (I),(II) and that it is non-multidimensional. This means that the following holds, too:

(III)   There exists a set $\mathscr{D}$ of strongly minimal sets such that:
- For all $J \in \mathscr{D}$, J is definable over $\text{acl}(\emptyset)$.
- For all $J_1 \neq J_2 \in \mathscr{D}$, $J_1 \perp J_2$.
- For every strongly minimal set D definable in T there exists $J \in \mathscr{D}$ such that $J \not\perp D$.

We may choose the elements of $\mathscr{D}$ strictly minimal and modular, and we may assume that T has QE. Call this the *bounded context*.

Now consider any μ-function $\mu: \mathscr{D} \to \mathbb{N} \cup \{\infty\}$ (there is no finite-to-one restriction). Then, μ-envelopes exist in T, and they are homogeneous [12]. If E is a μ-envelope, $\dim_E(J) = \mu(J)$ for all $J \in \mathscr{D}$. Note that this property usually does not hold for (pseudo-)envelopes in an unbounded context. This is the reason why it is more difficult to establish homogeneity for (pseudo)-envelopes in general.

For any $\mathscr{L}$-structure $A \models T^\forall$, say that *A is a μ-envelope* if A is equal to its μ-envelope when considered as a substructure of the prime model of T over A. The following result is shown in [12]:

THEOREM A.9. *In the bounded context, the class of μ-envelopes is axiomatisable. Its theory T^μ has quantifier elimination in $\mathscr{L}$. If $\mu(J)$ is finite for all $J \in \mathscr{D}$, then T^μ is strongly minimal.*

Appendix B. Relations to bicoloured fields. In this appendix we gather some results and observations which show the structural similarities between the fusion and the bicoloured field context. In fact, we show that bicoloured fields (black-and-white in any characteristic, red-and-white in positive characteristic) are covered by the axiomatic framework we presented in the previous appendix, as far as the geometrical analysis of the corresponding theory T_ω is concerned, thereby also suggesting a common strategy for a future collapse. We do not treat Poizat's *green fields*—algebraically closed fields (of characteristic 0) with a distinguished subgroup of the multiplicative group of the field—in what follows. We note that the original fusion and black-and-white fields were (partially) treated in a unified way in [3].

First, we give a quick overview of black-and-white fields. A *black-and-white field* is an algebraically closed field K together with a subset $N = N(K) \subseteq K$ (its *black points*; the others are white), such that any $A \subseteq_\omega K$ has positive predimension, i.e., $\delta(A) := 2 \cdot \mathrm{td}(A) - |N \cap A| \geq 0$. One works with fields of any fixed characteristic. The corresponding notion of *strong embedding* is defined using δ as in the case of fusions. The class $(\mathscr{C}_0, \leq)$ of all 'finitely generated' black-and-white fields has the joint embedding property, the amalgamation property and is countable.

As in the fusion context one gets a dimension function d and a closure operator $A \mapsto \overline{A}$. Field theoretic algebraic closure acl_f corresponds to $\langle \cdot \rangle$, and *selfsufficient closure* is defined as $\mathrm{cl}(A) := \mathrm{acl}_f(\overline{A})$. Every finitely generated strong extension of black-and-white fields L/K can be decomposed into a finite tower of extensions K_i/K_{i-1} which are either *primitive*, or *black generic* (i.e., $K_i = \mathrm{acl}_f(K_{i-1}b)$ for some black point b with $\mathrm{d}(b/K_{i-1}) = 1$), or *white generic* (i.e., $K_i = \mathrm{acl}_f(K_{i-1}a)$ for some white point a with $\mathrm{d}(a/K_{i-1}) = 2$). Let T_ω be the theory of the generic model. The following theorem can be quite easily obtained using [20] and [3].

THEOREM B.1. (1) T_ω *is ω-stable of rank $\omega \cdot 2$.*
(2) *Let p be the black generic (a regular type). Then, dimension corresponds to p-weight w_p.*
(3) *Parasitic types (more generally types q orthogonal to p, i.e., such that $w_p(q) = 0$) are 1-based. Primitive types are trivial strongly minimal and can be isolated by totally categorical formulas.*
(4) *Non-orthogonality between admissible types (chosen in a similar way to the fusion context) is definable. Only permutations of the variables can give rise to non-orthogonality. There is an almost normalised system of codes such that every instance is admissible (exactly as in [17]).* ⊣

Using this knowledge about T_ω, the corresponding collapsed black-and-white field (of Morley Rank 2) can be obtained exactly as the collapsed original fusion. This was recently worked out by Baudisch, Martin-Pizarro and Ziegler in [6], whereas Baldwin-Holland [3] proceeded in a slightly different way. In any case what one can show is that—for a suitable choice of the μ-function and corresponding pseudo-Morley sequences—pseudo-μ-envelopes are homogeneous and axiomatisable.

Combining [20], [3] and subsequent work by Baldwin-Holland [5], one gets the following theorem, which is entirely reproved in [6].

THEOREM B.2. *There is a black-and-white field (K^μ, N) with the following properties*:

(1) $\mathrm{MR}(K) = 2$ *and N is strongly minimal.*
(2) *The theory of (K, N) is model-complete.*
(3) *Morley rank and Lascar rank are the same and equal to* d.   ⊣

In [12] it is shown that B.2 can formally be deduced from the results in B.1.

As was noted in [3], the theory of algebraically closed fields in the above can be replaced by any other strongly minimal theory T which has DMP. Elimination of imaginaries in T makes things technically easier, but this property is not essential. As far as the non-collapsed theory T_ω is concerned, no DMP assumption is needed. This was already mentioned in [20].

In fact, in analogy with the *fusion over sublanguages*, we can also work in a relative context and just consider a strongly minimal expansion $T_0 \subseteq T_1$ and study bicoloured structures with respect to this setting, i.e., naming a predicate R for an $\mathscr{L}_0$-substructure and considering the predimension $\delta(A) := 2 \cdot \mathrm{d}_1(A) - \mathrm{d}_0(R \cap A)$. For the same reasons as in the relative fusion context, it is essential to work with a modular theory T_0. In what follows we also assume that T_0 is $\aleph_0$-categorical, although in the context which seems to be most interesting for applications (Poizat's *green fields*), this is not the case.

In order to make the exposition clearer we will now explain the case when T_0 is the theory of infinite vector spaces over some finite field F, the *red-and-white context*. If T_1 is the theory of algebraically closed fields of characteristic $p > 0$ and $F = \mathbb{F}_p$, we are dealing with Poizat's *red-and-white fields*. This is the context we will work in, but the same results hold in the general setting. We denote by T_ω the theory of the Fraïssé limit of the class $(\mathscr{C}_0, \leq)$ of red-and-white fields in characteristic $p > 0$ as defined in [21].

Performing the natural adaptations (e.g., for $A \subseteq K \in \tilde{\mathscr{C}}_0$ we take $\overline{A}$ as the smallest $\mathbb{F}_p$-subspace of K containing A which is strong in K), we get results analogous to Theorem B.1. The first two parts of the following theorem can be extracted from [21], and the remaining part is shown in a similar way to the corresponding results in the fusion over sublanguages.

THEOREM B.3. (1) *T_ω is ω-stable of rank $\omega \cdot 2$.*
(2) *Let p be the red generic. Then, d-dimension corresponds to p-weight, and the white generic has p-weight* 2.
(3) *Parasitic types (more generally types q orthogonal to p, i.e., such that $w_p(q) = 0$) are 1-based. Primitive types are locally modular strongly minimal and can be isolated by totally categorical formulas.*
(4) *Non-orthogonality between admissible types (chosen in a similar way to the fusion context) is definable. It is given by the action of the group of affine linear transformations (w.r.t. the $\mathbb{F}_p$-vector space structure) on admissible types.*
(5) *There is a complete system of codes for admissible types. Moreover, one can definably distinguish trivial types from locally projective ones, and the transition to the modular representative in the same non-orthogonality class is uniformly definable.*   ⊣

Among other things, it follows from B.3 that conditions (I)′ and (II) hold for T_ω in the red-and-white context, and that the strategy of its collapse by the construction of envelopes, as presented in the previous appendix, may prove a fruitful one.

As we already mentioned, there is still no collapse available for the *red-and-white* fields, the difficulties being completely analogous to those in the fusion over vector spaces.

As in the case of the fusion over vector spaces (see Lemma A.2), one cannot avoid the introduction of imaginary elements in order to obtain (almost) canonical representatives (for the types with trivial geometry). For details on this see [15].

REFERENCES

[1] J. T. BALDWIN, **Fundamentals of stability theory**, Springer Verlag, 1988.

[2] ———, *Notes on quasiminimality and excellence*, preprint, 2004.

[3] J. T. BALDWIN and K. HOLLAND, *Constructing ω-stable structures*: *Rank 2 fields*, this JOURNAL, vol. 65 (2000), pp. 371–391.

[4] ———, *Constructing ω-stable structures*: *Computing rank*, **Fundamenta Mathematicae**, vol. 170 (2001), pp. 1–20.

[5] ———, *Constructing ω-stable structures*: *Model completeness*, **Annals of Pure and Applied Logic**, vol. 125 (2004), pp. 159–172.

[6] A. BAUDISCH, A. MARTIN-PIZARRO, and M. ZIEGLER, *Fields and colors*, preprint, 2005.

[7] ———, *Fusion over a vector space*, preprint, 2005.

[8] G. CHERLIN, L. HARRINGTON, and A. LACHLAN, $\aleph_0$-*Categorical*, $\aleph_0$-*stable structures*, **Annals of Pure and Applied Logic**, vol. 28 (1985), pp. 103–135.

[9] G. CHERLIN and E. HRUSHOVSKI, **Finite structures with few types**, Annals of Mathematics Studies, vol. 152, Princeton University Press, 2003.

[10] P. M. COHN, **Skew field constructions**, London Mathematical Society Lecture Note Series, vol. 27, Cambridge University Press, 1977.

[11] H. P. GUTE and K. K. REUTER, *The last word on elimination of quantifiers in modules*, this JOURNAL, vol. 55 (1990), pp. 670–673.

[12] A. HASSON, *Collapsing structure and a theory of envelopes*, preprint, 2004.

[13] ———, *Fusion over sublanguges revisited*, preprint, 2006, (Available on author's webpage).

[14] M. HILS, *La fusion au-dessus d'un sous-langage*: *le cas simple*, preprint, 2004.

[15] ———, Ph.D. thesis, in preperation.

[16] K. HOLLAND, *Model completeness of the new strongly minimal sets*, this JOURNAL, vol. 64 (1999), pp. 946–962.

[17] E. HRUSHOVSKI, *Strongly minimal expansions of algebraically closed fields*, **Israel Journal of Mathematics**, vol. 79 (1992), pp. 129–151.

[18] ———, *A new strongly minimal set*, **Annals of Pure and Applied Logic**, vol. 62 (1993), pp. 147–166.

[19] A. PILLAY, **Geometric stability theory**, Clarendon Press, Oxford, 1996.

[20] B. POIZAT, *Le carré de l'égalité*, this JOURNAL, vol. 64 (1999), pp. 1339–1355.

[21] ———, *L'égalité au cube*, this JOURNAL, vol. 66 (2001), pp. 1647–1676.

[22] B. ZIL'BER, **Uncountably categorical theories**, Translations of Mathematical Monographs, vol. 117, American Mathematical Society, Providence, R.I., 1993.

EINSTEIN INSTITUTE OF MATHEMATICS, EDMOND J. SAFRA CAMPUS, GIVAT RAM
   THE HEBREW UNIVERSITY OF JERUSALEM, JERUSALEM, 91904, ISRAEL
Current address: The Mathematical Institute, 24–29 St. Giles, Oxford, OX1 3LB, UK
E-mail: hasson@maths.ox.ac.uk

INSTITUT GIRARD DESARGUES (LYON 1), UNIVERSITÉ CLAUDE BERNARD
   21, AVENUE CLAUDE BERNARD, 69622 VILLEURBANNE-CEDEX, FRANCE
Current address: Institut für Mathematik, Humboldt-Universität zu Berlin, D-10099 Berlin, Germany
E-mail: hils@mathematik.hu-berlin.de