J. Aust. Math. Soc. **108** (2020), 262–277
doi:10.1017/S1446788719000399

THE QUANDARY OF QUANDLES: A BOREL COMPLETE KNOT INVARIANT

ANDREW D. BROOKE-TAYLOR and SHEILA K. MILLER

(Received 21 February 2017; accepted 29 July 2019; first published online 30 October 2019)

Communicated by G. Willis

Abstract

We show that the isomorphism problems for left distributive algebras, racks, quandles and kei are as complex as possible in the sense of Borel reducibility. These algebraic structures are important for their connections with the theory of knots, links and braids. In particular, Joyce showed that a quandle can be associated with any knot, and this serves as a complete invariant for tame knots. However, such a classification of tame knots heuristically seemed to be unsatisfactory, due to the apparent difficulty of the quandle isomorphism problem. Our result confirms this view, showing that, from a set-theoretic perspective, classifying tame knots by quandles replaces one problem with (a special case of) a much harder problem.

2010 *Mathematics subject classification*: primary 20N99; secondary 03E15, 57M27.

Keywords and phrases: quandle, left distributive, Borel complete, knot invariant.

1. Introduction

Left distributivity arises in the study of many well-known mathematical objects such as groups, knots and braids, and also in the study of large cardinal embeddings in set theory. Specifically, left distributive algebras are structures with one binary operation $*$ satisfying the left self-distributivity law $a * (b * c) = (a * b) * (a * c)$. Familiar examples include the conjugation operation on any group and the implication operation on any Boolean algebra; symmetric spaces in differential geometry provide further examples [1]. The first nontrivial example of a free left distributive algebra on one generator is due to Laver [17], who showed that the algebra generated by closing

Much of this work was carried out while both authors were Visiting Fellows at the Isaac Newton Institute for Mathematical Sciences in the programme 'Mathematical, Foundational and Computational Aspects of the Higher Infinite' (HIF) funded by EPSRC grant EP/K032208/1. The first author was supported by the UK Engineering and Physical Sciences Research Council Early Career Fellowship EP/K035703/1 and EP/K035703/2, *Bringing set theory and algebraic topology together,* and undertook some of the work whilst visiting the Centre de Recerca Matemàtica for the programme *Large cardinals and strong logics.* The second author was supported by grants from PSC-CUNY and the City Tech PDAC.

a *rank-to-rank embedding* under the application operation is such an algebra. (The existence of these embeddings is one of the strongest known set-theoretic axioms. For more on them, see [17–19], [5] or [6].) Other interesting classes of structures are obtained by adding further algebraic axioms to the left distributive law. Racks are left distributive algebras such that, for every a and c in the algebra, there is a unique b such that $a * b = c$. Quandles are racks satisfying $a * a = a$ for every element a.

Quandles were rediscovered and named by Joyce in his thesis, published in [13]. (The objects that Joyce called *involutory quandles*, namely, those satisfying $x * (x * y) = y$ for all x, y in the underlying set, were first considered by Takasaki [23] in 1943 under the name *kei*, with the kei of reflections in the plane under conjugation being a central example.) Joyce established many foundational relationships in his thesis, including those between quandles and group conjugation and quandles and knots. Indeed, he showed that the equational theory of quandles is precisely the equational theory of the conjugation operation: any identity true in every group with its conjugation operation is also true in every quandle, and hence provable from the quandle axioms.

The three quandle axioms may also be viewed as algebraic versions of the familiar Reidemeister moves for passing between different regular projections of equivalent tame knots. One may consequently associate to any tame knot K a quandle $Q(K)$ generated by the arcs of the knot and with identities dictated by the crossings. (Tame knots essentially correspond to one's intuitive notion of finite knots in three-dimensional space and, in particular, are not assumed to be endowed with an orientation.) Further, for any (possibly wild) knot K embedded in a space X, Joyce defined the *fundamental quandle* $Q(K, X)$ analogously to the fundamental group. For tame knots Joyce showed that $Q(K)$ can be derived from $Q(K, S^3)$ and, moreover, that the quandle $Q(K)$ constitutes a complete invariant for K: two tame knots K and K' are equivalent if and only if their associated quandles $Q(K)$ and $Q(K')$ are isomorphic.

However, knot theorists express some dissatisfaction with quandles as knot invariants, in part, because of the apparent difficulty in determining whether two quandles are isomorphic. Indeed (as the anonymous referee kindly pointed out to us), substantial work has been done in obtaining secondary invariants that can be derived from the knot quandle that are more practical (see, for example, [14] for a survey). Our result makes rigorous this impression that the quandle itself is difficult to work with: we show that the isomorphism problem for arbitrary countable quandles is as complex as possible for algebraic structures, in the sense of *Borel reducibility* (discussed below). By contrast, the problem of distinguishing tame knots up to equivalence is trivial in this context as there are only countably many possibilities. Thus, if the goal is to simplify the problem of distinguishing different objects, moving to the algebraic framework of quandles may be viewed as a step in the wrong direction; any reasonable algorithm for dealing with knot quandles will necessarily use information about them beyond the mere fact of being countable quandles.

The complexity of classification problems and the study of complete invariants for structures have emerged as major themes in set theory. Broadly, a classification can be

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

thought to assign mathematical objects of one type—considered up to isomorphism or some other such equivalence relation—to mathematical objects of another type (again, up to an equivalence relation), where the former act as invariants. Frequently, the objects in question, both those to be classified and the invariants, can be encoded by real numbers. For example, countable structures with underlying set $\mathbb{N}$, such as groups, rings and, indeed, left distributive algebras and quandles can be encoded in a natural way by sets of finite tuples of natural numbers, and hence by reals. Classification then amounts to finding a reasonably definable map from the reals encoding the structures to the reals encoding the invariants that respects the relevant equivalence relations. The words 'reasonably definable' are important here—a nonconstructive proof of the existence of such a map using, for example, the Axiom of Choice should not be considered a classification. A natural way to exclude such uninformative maps would be to require the map to be continuous, but this interpretation is too restrictive to be practical. For example, it is reasonable to encode reals by elements of $2^{\mathbb{N}}$, but no such encoding map can be continuous, by connectedness considerations. The more liberal constraint that the map be Borel, however, permits almost all constructions that arise in practice whilst being restrictive enough to obtain meaningful theorems about the framework.

Classifying structures using Borel maps between sets of encoding reals gives rise to the notion of Borel reducibility. Given two equivalence relations E and F on real numbers, we say that E is Borel reducible to F, written $E \leq_B F$, if there is a Borel function f from $\mathbb{R}$ to $\mathbb{R}$ such that, for all x and y in $\mathbb{R}$, $x\ E\ y$ holds if and only if $f(x)\ F\ f(y)$ holds. Establishing that one equivalence relation is *not* Borel reducible to another has been used in a number of cases to show that a classification problem is impossible to resolve. For example, Farah, Toms and Törnquist [8] used this analysis to show that unital simple separable nuclear C^*-algebras are not classifiable by countable structures (note that each adjective makes the theorem stronger), and Foreman, Rudolph and Weiss [9] showed that ergodic measure-preserving transformations of the unit interval are not classifiable by countable structures (and, indeed, much more). For more on this area see, for example, the books of Hjorth [12] and Gao [11].

Against this background it is natural to ask: what is the Borel reducibility complexity of the isomorphism relation on the class of countable left distributive algebras? This question was indeed posed to the second author by Matt Foreman. In this article, we show that it has the maximum possible complexity for an isomorphism relation on a first-order class of countable structures: in the standard terminology introduced in the seminal paper of Friedman and Stanley [10], isomorphism of left distributive algebras is *Borel complete*. Moreover, the same is true for the subclasses of racks, quandles and kei (see Section 2 for definitions). We show directly that isomorphism of kei (Definition 2.1.4) is Borel complete; the result for the other, more general classes follows. We also observe that the related class of expanded left distributive algebras satisfying the set of axioms that Laver [18] denoted by Σ (Definition 2.2) is Borel complete, as this follows from the fact, due to Mekler [20], that the class of groups is Borel complete.

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

Our results add to the list of isomorphism relations for countable structures that are known to be Borel complete. Other examples include those for graphs, for linear orders, for trees [10] and for groups [20]. It remains an important open question whether the isomorphism relation for the class of countable *abelian* groups is Borel complete.

Whilst Camerlo and Gao [3] have shown that the isomorphism relation for countable Boolean algebras is Borel complete, our result for left distributive algebras does not obviously follow, because Boolean algebras have algebraic structure in addition to the left distributive implication operation. Furthermore, as a Boolean algebra with the implication operation is not a rack, our other results are orthogonal to those of Camerlo and Gao.

We present the technical preliminaries in Section 2, and we give the main result and corollaries in Section 3. We prove that the class of countable kei, a subclass of the countable left distributive algebras, is Borel complete. The proof proceeds by reducing the question to the folklore result that the class of countable irreflexive directed graphs is Borel complete. We provide an explicit construction of a countable kei $け_G$ associated with every countable irreflexive directed graph G, making use of Kamada's notion of a *dynamical quandle* [15]. (け is the Japanese hiragana letter 'ke'. In the paper [23] introducing them, Takasaki used the kanji character 圭 for kei.) Our construction of け clearly constitutes a Borel map; in fact, it is continuous between the relevant topological spaces (described in Section 2). The main technical difficulty lies in showing that the map け is injective on isomorphism classes: if two kei associated to graphs are isomorphic, then the graphs from which they were obtained are isomorphic, that is, if $け_G \cong け_{G'}$, then $G \cong G'$. Whilst not every isomorphism φ between such kei $け_G$ and $け_{G'}$ arises from an isomorphism between the graphs G and G', by considering the combinatorics of the kei operations we show that the existence of such a φ guarantees the existence of *some* isomorphism between G and G'.

We thus show that countable quandles are well above tame knots in the Borel complexity hierarchy. Meanwhile, Kulikov [16] has recently shown that equivalence of arbitrary knots—including wild knots with infinitely many crossings—is strictly more complex than the isomorphism relation on any first-order class of countable structures, and hence, in particular, the isomorphism relation of countable quandles. This raises natural questions about the relationship between quandles and knots, which we discuss along with other related questions in Section 4.

2. Preliminaries

As we will be discussing the related classes of left distributive algebras, racks, quandles and kei, we begin by giving some intuition for them. These classes of structures can usefully be understood in terms of the behaviour of the action of left multiplication by an element of the algebra. For structures with underlying set A and binary operation $*$, and for each a in A, denote by m_a the map from A to A that acts by multiplication on the left by a, that is, $m_a(b) = a * b$. Then *left distributive algebras*

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

are those for which m_a is a homomorphism from A to itself for each a in A. A *rack* is a left distributive algebra in which each m_a is an automorphism (indeed, Brieskorn [2] referred to racks as *automorphic sets*). In a *quandle*, m_a is an automorphism and a is a fixed point of m_a for each a in A. Finally, a *kei* (plural kei; also called an *involutory quandle*) is a quandle such that each m_a is its own inverse. For the history of the nomenclature in this area, see the Preface of [7].

Formally, these structures can be defined using the following axioms.

(i) For every a, b and c in A, $a * (b * c) = (a * b) * (a * c)$.
(ii) For all a and c in A, there is a unique b in A such that $a * b = c$.
(iii) For every a in A, $a * a = a$.
(iv) For all a and b in A, $a * (a * b) = b$.

Definition 2.1. For a set A with one binary operation $*$ (an *algebra*), we define the following.

(1) A *left distributive algebra* is an algebra satisfying axiom (i).
(2) A *rack* is an algebra satisfying axioms (i) and (ii).
(3) A *quandle* is an algebra satisfying axioms (i), (ii) and (iii).
(4) A *kei* is an algebra satisfying axioms (i), (ii), (iii) and (iv).

There are a number of choices to be made in presenting the above definitions. Instead of using axiom (ii), one can formulate racks using a second operation $\bar{*}$ such that the function $m_a : b \mapsto a * b$ is inverse to the function $b \mapsto a \bar{*} b$: formally, one requires that, for all a and b, $a \bar{*} (a * b) = a * (a \bar{*} b) = b$ holds. This has the advantage of eliminating the existential quantifier. Whether to consider self distributive structures as left distributive, like we do here, or right distributive (with axioms (ii) and (iv) reformulated for right multiplication) is an arbitrary choice. Many relevant references on racks, quandles and kei use right distributivity; we chose left distributivity in order to easily view these classes of structures as subclasses of the left distributive algebras.

There is another well-studied left distributive structure, this one with two operations: the left distributive operation $*$ and another operation $\circ$ that behaves like composition. These algebras were first studied by Laver [17] as algebras of large cardinal embeddings in which the operation $\circ$ is, in fact, composition.

Definition 2.2. We denote by Σ the following collection of four identities

$$\begin{aligned} a \circ (b \circ c) &= (a \circ b) \circ c \\ (a \circ b) * c &= a * (b * c) \\ a * (b \circ c) &= (a * b) \circ (a * c) \\ (a * b) \circ a &= a \circ b. \end{aligned}$$

Note that left distributivity follows from the second and fourth identities via the equalities $a * (b * c) = (a \circ b) * c = ((a * b) \circ a) * c = (a * b) * (a * c)$. Dehornoy refers to algebras satisfying Σ as *LD-monoids* (for example, in [5] and [6]); we use Laver's original phrase 'algebras satisfying Σ' to avoid any potential confusion with

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

other uses of 'monoid'. If $\circ$ is a group operation on A, then the fourth equational condition of Σ determines that the other operation $*$ must be the conjugation operation $a * b = a \circ b \circ a^{-1}$. Taking $*$ to denote conjugation in the group in question, it is straightforward to check that the other identities of Σ are also satisfied, so any group with its multiplication and conjugation operations is an algebra satisfying Σ.

Laver showed, among other things, that Σ is a conservative extension of the left distributive law [17]. Thus any free left distributive algebra may be expanded to a free algebra on the same generators satisfying Σ: any identity on elements of the free left distributive algebra will hold in the algebra satisfying Σ if and only if it is a consequence of the left distributive law. For more on this, the linearity of several orderings on the free left distributive algebra (from the large cardinal hypothesis) and a normal form for terms in the free left distributive algebra, see [17] and [18]. For a simpler proof and fuller account of the theory of left distributive algebras, see [19]. Using braid groups, Dehornoy showed within the standard axioms of set theory that the above-mentioned orderings on the free left distributive algebra are linear [4]; Dehornoy has also contributed substantially to the literature on algebras satisfying Σ (see, for example, [5]).

We now move on to preliminaries regarding Borel reducibility. Recall that a subset of a topological space is *Borel* if it lies in the least σ-algebra containing the open sets, and that a function between two topological spaces is Borel if the inverse image of any Borel set (or, equivalently, of any open set) is Borel. Thus, to discuss Borel reducibility between classes of countable structures, we first define a topology on each of these classes. We briefly sketch this definition here, and refer the reader to Section 2.3 of Hjorth's book [12] for further details.

We exclusively consider countable structures, and so we may assume that each structure has underlying set $\mathbb{N}$. Furthermore, all of the classes of structures we consider are first-order, namely, the structures have finitely many relations and operations, and the class is defined by formulas involving these relations and operations. The relations and operations of a structure in one of these classes can thus be represented by a set of tuples from $\mathbb{N}$. Indeed, we follow the common practice of identifying a directed graph $(\mathbb{N}, E)$ (with vertex set $\mathbb{N}$) with the set $\{(m, n) \mid m\ E\ n\} \subseteq \mathbb{N}^2$, and we may identify an algebra $(\mathbb{N}, *)$ with the set $\{(\ell, m, n) \mid \ell * m = n\} \subset \mathbb{N}^3$. The space of countable structures for a given signature with finitely many operation and relation symbols can thus be identified with a subset of Cantor space via the usual identification of a power set $\mathcal{P}(X)$ with the space of characteristic functions 2^X; the set X here is a product of sets of the form $\mathbb{N}^k$, one for each relation and operation, and is, in particular, countable. The topology considered on these classes is the standard topology on the Cantor space. Note that a clopen subbase for this topology is given by the sets defined by determining a single 'bit' from 2^X—for example, on the space of countable algebras with underlying set $\mathbb{N}$, the subbase is the collection as ℓ, m and n vary over $\mathbb{N}$ of all sets either of the form $\{(\mathbb{N}, *) \mid \ell * m = n\}$ or of the form $\{(\mathbb{N}, *) \mid \ell * m \neq n\}$.

We deviate from this conventional framework in one detail: for expositional clarity, the kei that we construct will have underlying set $\mathbb{N} \times \{0, 1\}$ rather than $\mathbb{N}$. However,

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

this discrepancy can be easily overcome using the canonical identification of $\mathbb{N} \times \{0, 1\}$ with $\mathbb{N}$ via the map $(n, i) \mapsto 2n + i$.

Note that the Cantor space 2^X with X countable is a separable topological space (that is, it has a countable dense set) and may be endowed with a complete metric: identifying X with $\mathbb{N}$, let $d(x, y) = 2^{-n}$, where n is least such that $x(n) \neq y(n)$. Separable, completely metrisable spaces such as 2^X and $\mathbb{R}$ are known as *Polish spaces*. As outlined in the Introduction, we have the following standard definitions.

Definition 2.3. Let X and Y be Polish spaces, let E be an equivalence relation on X, and let F be an equivalence relation on Y. We say that E is *Borel reducible to* F, written $E \leq_B F$, if there is a Borel function f from X to Y such that, for all x and x' in X, $x\ E\ x'$ holds (that is, x is E-equivalent to x') if and only if $f(x)\ F\ f(x')$ holds.

We say that E is *continuously reducible* to F, written $E \leq_c F$, if there is a continuous function f from X to Y such that, for all x and x' in X, $x\ E\ x'$ if and only if $f(x)\ F\ f(x')$.

If F is the isomorphism relation for a first-order class of countable structures for a finite signature, each with underlying set $\mathbb{N}$, we say that F is *Borel complete* if every other such class has isomorphism relation Borel reducible to F.

Continuous maps are, of course, Borel, and all maps we construct in what follows will be continuous.

3. The class of kei is Borel complete

It is folklore that the class of countable irreflexive directed graphs is Borel complete—see Section 13.1 of Gao's book [11] for a proof of the stronger statement that the subclass of countable irreflexive symmetric graphs is Borel complete. The general strategy of this section is to construct a kei from an arbitrary irreflexive directed graph, and then to show that the resulting kei are isomorphic if and only if the original graphs are isomorphic. Since the map taking each irreflexive directed graph to the corresponding kei will be Borel (indeed, continuous), this will establish that the class of countable kei is also Borel complete. To this end, we shall describe how to build what Kamada [15] calls a *dynamical quandle*; the specific dynamical quandles we construct will, in fact, be kei.

In all of the following we exclusively discuss graphs that are irreflexive and directed, but, for the sake of the casual reader, we will repeat these hypotheses each time they are used.

Let A be a set and let τ be a bijection from A to itself. Let φ be a map from A to the power set $\mathcal{P}(A)$ such that, for every $a \in A$, $\varphi(a)$ contains a, $\varphi(a)$ is closed under τ and τ^{-1}, and $\varphi(a) = \varphi(\tau a)$. We will refer to such maps φ as *τ-replete*. Kamada observes [15, Theorem 4] that with the operation $*$ defined by

$$a * b = \begin{cases} b & \text{if } a \in \varphi(b), \\ \tau b & \text{if } a \notin \varphi(b), \end{cases}$$

the structure $(A, *)$ is a quandle. Kamada uses an equivalent definition with a function θ defined on τ-orbits rather than our orbit-invariant function φ on elements of A. Axioms

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

(ii) and (iii) of Definition 2.1 are immediate from the assumptions on φ, and (i) follows by checking cases

$$a * (b * c) = (a * b) * (a * c) = \begin{cases} c & \text{if } a \in \varphi(c) \text{ and } b \in \varphi(c), \\ \tau c & \text{if } a \in \varphi(c) \text{ and } b \notin \varphi(c), \\ \tau c & \text{if } a \notin \varphi(c) \text{ and } b \in \varphi(c), \\ \tau^2 c & \text{if } a \notin \varphi(c) \text{ and } b \notin \varphi(c). \end{cases}$$

Moreover, if τ is an involution, then, clearly, axiom (iv) also holds and so the quandle is a kei. Following Kamada, but using our φ rather than Kamada's θ, we call this $(A, *)$ the *quandle derived from* (A, τ) *relative to* φ. Kamada named the objects so constructed *dynamical quandles*, in line with a view of the pair (A, τ) as a dynamical system, and we shall call those dynamical quandles that are kei *dynamical kei*.

To encode an irreflexive directed graph $G = (V, E)$ into a kei け_G, we use the dynamical quandle construction with underlying set being a *pair* of copies of the vertex set V of G. Our involution τ simply switches between the two copies of the vertex set, and the function φ corresponds to choosing the set of neighbours (in one direction) for each vertex of G, irrespective of which copy of V contains the vertices.

Definition 3.1. Suppose $G = (V, E)$ is an irreflexive directed graph. Let τ be the involution on $V \times \{0, 1\}$ taking $(v, 0)$ to $(v, 1)$ and $(v, 1)$ to $(v, 0)$ for every v in V. Let $\bar{\varphi}_G$ be the function from V to $\mathcal{P}(V)$ defined by $u \in \bar{\varphi}_G(v)$ if and only if $u\,E\,v$ or $u = v$. Let φ_G from $V \times \{0, 1\}$ to $\mathcal{P}(V \times \{0, 1\})$ be the function obtained from $\bar{\varphi}_G$ by ignoring second coordinates: $(u, i) \in \varphi_G(v, j)$ if and only if $u \in \bar{\varphi}_G(v)$, that is, if and only if $u\,E\,v$ or $u = v$. Note that φ_G is τ-replete. The *kei* け_G *associated to* G is the quandle derived from $(V \times \{0, 1\}, \tau)$ relative to φ_G, and we denote the operation on け_G by $*_G$.

Thus, け_G is a kei on underlying set $V \times \{0, 1\}$ with operation $*$ such that $(u, i) * (v, j)$ equals (v, j) if there is an edge from u to v in G or if $u = v$, and $(u, i) * (v, j)$ is $(v, 1 - j)$ otherwise.

We now begin working towards Theorem 3.5, which says that the dynamical kei け_G and $\text{け}_{G'}$ constructed from graphs G and G' are isomorphic if and only if the graphs G and G' are isomorphic. First, we prove the existence of a particular, useful involution of the kei け_G.

Lemma 3.2. *For every irreflexive directed graph G with underlying set V and for every $W \subseteq V$, the function $I_W : \text{け}_G \to \text{け}_G$, defined by*

$$I_W(v, j) = \begin{cases} (v, j) & \text{if } v \in W, \\ (v, 1 - j) & \text{if } v \notin W, \end{cases}$$

is an involution of け_G.

Proof. By inspection, I_W is a bijection and, moreover, $(I_W)^2$ is the identity map. To see that I_W respects the quandle operation $*$ of け_G, we must verify that $I_W((u, i) * (v, j)) =$

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

$I_W(u,i) * I_W(v,j)$. Note that, for each $(v,j) \in \text{け}_G$, either both of $(u,0)$ and $(u,1)$ are in $\varphi_G(v,j)$ or neither is, so

$$(u,i) * (v,j) = (I_W(u,i)) * (v,j) = \begin{cases} (v,j) & \text{if } (u,i) \in \varphi(v,j), \\ (v,1-j) & \text{if } (u,i) \notin \varphi(v,j). \end{cases}$$

So

$$I_W((u,i) * (v,j)) = \begin{cases} I_W(v,j) & \text{if } (u,i) \in \varphi(v,j), \\ I_W(v,1-j) & \text{if } (u,i) \notin \varphi(v,j), \end{cases}$$

and

$$I_W((u,i)) * I_W(v,j) = \begin{cases} I_W(v,j) & \text{if } (u,i) \in \varphi(I_W(v,j)) = \varphi((v,j)), \\ (v,1-j) = I_W(v,1-j) & \text{if } v \in W \text{ and } (u,i) \notin \varphi((v,j)), \\ (v,j) = I_W(v,1-j) & \text{if } v \notin W \text{ and } (u,i) \notin \varphi((v,j)). \end{cases}$$

Thus it is established that I_W is a homomorphism, and is indeed an involution of け_G. □

A slicker if less direct proof of Lemma 3.2 is to consider the graph G' on $V \dot\cup \{v_0\}$ (where $\dot\cup$ denotes disjoint union) with the restriction $G' \upharpoonright V$ of G' to the vertices V being G, and taking $v_0 \mathrel{E} v$ if and only if v is in W for each v in V. Then the restriction $\text{け}_{G'} \upharpoonright V \times \{0,1\}$ of $\text{け}_{G'}$ to $V \times \{0,1\}$ is simply け_G, and $m_{v_0} \upharpoonright V \times \{0,1\} = I_W$.

The kei constructed in Definition 3.1 are, in fact, quite general dynamical kei. Indeed, the only extra constraint we need on dynamical kei to get a kei け_G associated to a graph G is that the involution τ has no fixed points.

Definition 3.3. A kei $(A,*)$ is called a *folded kei*[1] if there is an involution τ of A with no fixed points and a τ-replete function φ such that $(A,*)$ is the quandle derived from (A,τ) relative to φ.

By definition, the kei け_G associated to any graph G is a folded kei. As alluded to above, we also have a converse to this.

Proposition 3.4. *Every folded kei is isomorphic to a kei of the form け_G for some irreflexive directed graph G.*

Proof. Let $(A,*)$ be a folded kei and, in particular, suppose $(A,*)$ is the quandle derived from (A,τ) relative to φ for τ an involution of A without fixed points and φ is a τ-replete function from A to $\mathcal{P}(A)$. Choose a subset V of A such that, for each pair $\{a,\tau a\}$ of elements of A, exactly one of a and τa is in V, and express A as the disjoint union $A = V \cup \{\tau v \mid v \in V\}$. For each v in V, let $\bar\varphi(v)$ denote the set $\varphi(v) \cap V$; since $(A,*)$ is the quandle derived from (A,τ) relative to φ, we have that $\bar\varphi(v)$ is the set of u in V such that $u * v = v$ (this $\bar\varphi$ will be $\bar\varphi_G$, as in Definition 3.1, for the graph G we now construct). Take the directed graph G on vertex set V with edge relation defined by $u \mathrel{E} v$ if and only if $u \in \bar\varphi(v)$ holds and $u \neq v$. Then it is straightforward to check that the map from け_G to A taking $(v,0)$ to v and $(v,1)$ to τv is an isomorphism of kei. □

[1]In baking, one *folds* ingredients to achieve complete mixing with minimal disruption.

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

We will now state the main result.

THEOREM 3.5. *For irreflexive directed graphs G and G' and the associated kei け_G and $\text{け}_{G'}$, $G \cong G'$ if and only if $\text{け}_G \cong \text{け}_{G'}$.*

PROOF. One direction is a fairly straightforward observation.

REMARK 3.6. Isomorphic irreflexive directed graphs have isomorphic associated kei.

PROOF OF REMARK. Recall that a graph isomorphism is a bijection between vertices that preserves both the edge relation and the failure of the edge relation. Given graphs $G = (V, E)$ and $G' = (V', E')$ with an isomorphism $h : G \to G'$ between them, $u\ E\ v$ in G if and only if $h(u)\ E'\ h(v)$ in G', so u is in $\bar{\varphi}_G(v)$ if and only if $h(u)$ is in $\bar{\varphi}_{G'}(h(v))$. Therefore, by construction of the quandles け_G and $\text{け}'_G$, h induces an isomorphism $h_{\text{け}}$ from け_G to $\text{け}_{G'}$ taking (u, i) to $(h(u), i)$. Indeed, for vertices u and v in G, we have that $(u, i) \in \varphi_G(v, j)$ holds if and only if $(h(u), i) \in \varphi_{G'}(h(v), j)$ holds. The verification that $x *_G y = z$ if and only if $h_{\text{け}}(x) *_{G'} h_{\text{け}}(y) = h_{\text{け}}(z)$ follows immediately. □

For the converse, we will show that any two isomorphic kei of the form け_G and $\text{け}_{G'}$ admit an isomorphism induced by an isomorphism of the underlying graphs G and G'. Not all kei isomorphisms between け_G and $\text{け}_{G'}$ arise from graph isomorphisms; indeed, Lemma 3.2 gives continuum many others. Also, if the graph K is the complete irreflexive directed graph on V, then け_K is the trivial kei on $V \times \{0, 1\}$, with $(u, i) * (v, j) = (v, j)$ for all (u, i) and (v, j). Of course, there are many automorphisms of the trivial kei that are not of the form given by Lemma 3.2 or induced by a graph isomorphism: any permutation of the underlying set $V \times \{0, 1\}$ is an automorphism of this kei. We will see in the claim that follows that any kei isomorphism ρ between folded kei splits into two parts, one being of the type described by Lemma 3.2 and one given by an automorphisms of a trivial kei. Each of these can be converted into a partial isomorphism of the desired form, and the pieces can be recombined to yield the graph isomorphism required for the theorem. To aid with intuition, for any graph $G = (V, E)$ with associated kei $\text{け}_G = (V \times \{0, 1\}, *_G)$, we refer to $V \times \{0\} \subset \text{け}_G$ as the *bottom* of け_G and to $V \times \{1\} \subset \text{け}_G$ as the *top* of け_G. Also, for any v in V, we refer to each of $(v, 0)$ and $(v, 1)$ as the *twin* of the other.

PROPOSITION 3.7. *Suppose $G = (V_G, E_G)$ and $G' = (V_{G'}, E_{G'})$ are irreflexive directed graphs such that there is a kei isomorphism ρ from け_G to $\text{け}_{G'}$. Then there is bijection f from V_G to $V_{G'}$ such that, viewed as a map from G to G', f is a graph isomorphism.*

PROOF OF PROPOSITION 3.7. For any graph $H = (V, E)$, we split the underlying set V into two components, which we call the 'fixed points' and the 'moving points' based on their behaviour in the quandle け_H. The purely graph-theoretic definitions of the fixed points and moving points is simpler, so we give them first: the fixed points are those which are complete for inward edges, and the moving points are those that are not. That is,

$$F_H = \{v \in V \mid \forall u \in V(u = v \text{ or } u\ E\ v)\}.$$

From the quandle point of view, the fixed points may, equivalently, be defined as those v for which left multiplication by any element of け_H does not swap $(v,0)$ with $(v,1)$, that is,

$$F_H = \{v \in V \mid \forall (u,i) \in \text{け}_H[(u,i) *_H (v,0) = (v,0)]\}.$$

The moving points are then those not in F_H, that is, $M_H = V \setminus F_H$.

When we come to define the function $f : V_G \to V_{G'}$ it will be piecewise, giving separately the restrictions of f to the fixed points F_G and the moving points M_G. In fact, these restrictions will themselves be bijections from F_G to $F_{G'}$ and from M_G to $M_{G'}$, as is clearly necessary for f to be a graph isomorphism.

We are given an isomorphism $\rho : \text{け}_G \to \text{け}_{G'}$. Let us denote by $\rho_V(v,i)$ and $\rho_I(v,i)$, respectively, the first and second components of $\rho(v,i)$: that is, $\rho(v,i) = (\rho_V(v,i), \rho_I(v,i))$.

First, we define f on the moving points. If v is in M_G, then there is some (u,i) in け_G that moves $(v,0)$. That is, the value of $(u,i) *_G (v,0)$ is not $(v,0)$, and hence, by the definition of $*_G$, it must be that $(u,i) * (v,0)$ is $(v,1)$ and, furthermore, that $(u,i) * (v,1)$ is $(v,0)$. Applying the kei isomorphism ρ, we have that $\rho(u,i) * \rho(v,0) = \rho(v,1)$ holds, and by injectivity $\rho(v,1) \neq \rho(v,0)$. By the definition of $*_{G'}$, the first components of $\rho(v,0)$ and $\rho(v,1)$ must be equal. We take $f(v)$ to be this value: $f(v) = \rho_V(v,0) = \rho_V(v,1)$.

Clearly, $f \restriction M_G$ (f restricted to M_G) so defined is injective since ρ is a bijection. Moreover, $f \restriction M_G$ surjects onto $M_{G'}$. Indeed, for w in $M_{G'}$ and (t,i) in $\text{け}_{G'}$ such that $(t,i) *_{G'} (w,0) \neq (w,0)$, we have $\rho^{-1}(t,i) *_G \rho^{-1}(w,0) \neq \rho^{-1}(w,0)$, and so the first component of $\rho^{-1}(w,0)$ lies in M_G and has image w under f.

To complete the definition of f, it remains to give the value of $f(v)$ for those v in F_G. Let v_0 be an element of F_G. Unlike for elements of M_G, it need not be the case that $\rho_V(v_0,0)$ is the same as $\rho_V(v_0,1)$. However, since ρ is surjective, we may find v_1 in F_G and i_{v_1} in $\{0,1\}$ such that $\rho_V(v_1, i_{v_1}) = \rho_V(v_0,1)$ and $\rho_I(v_1, i_{v_1}) = 1 - \rho_I(v_0,1)$: that is, if $\rho(v_0,1)$ is on the bottom of the kei, then (v_1, i_{v_1}) is chosen such that $\rho(v_1, i_{v_1})$ is its twin on the top, and, conversely, if $\rho(v_0,1)$ is on the top of the kei, then (v_1, i_{v_1}) is chosen such that $\rho(v_1, i_{v_1})$ is its twin on the bottom. Likewise, we may find v_{-1} in F_G and $i_{v_{-1}}$ in $\{0,1\}$ such that $\rho_V(v_{-1}, 1 - i_{v_{-1}}) = \rho_V(v_0,0)$ and $\rho_I(v_0,0) = 1 - \rho_I(v_{-1}, 1 - i_{v_{-1}})$. We may inductively extend our definitions, obtaining for all k in $\mathbb{Z}$ a vertex v_k in V_G and i_{v_k} in $\{0,1\}$ (with $i_{v_0} = 0$) such that $\rho_V(v_k, 1 - i_{v_k}) = \rho_V(v_{k+1}, i_{v_{k+1}})$ and $\rho_I(v_k, 1 - i_{v_k}) \neq \rho_I(v_{k+1}, i_{v_{k+1}})$. Note that if there is some k such that $v_k = v_0$, then i_{v_k} defined in this way will be equal to i_{v_0}, so our notation i_{v_j} gives a well-defined function from vertices v_j in F_G to members of $\{0,1\}$. Indeed (construing for now i_{v_j} as a function of j rather than v_j), consider the first repetition in the sequence $(v_0, i_{v_0}), (v_0, 1 - i_{v_0}), (v_1, i_{v_1}), \ldots$. Clearly, if (v_k, i_{v_k}) is distinct from all of its predecessors in the sequence, then so too is $(v_k, 1 - i_{v_k})$. Thus, the first repetition in the sequence must be of the form (v_k, i_{v_k}). If $(v_k, i_{v_k}) = (v_j, 1 - i_{v_j})$ for some $j < k$, then, of course, $\rho(v_k, i_{v_k}) = \rho(v_j, 1 - i_{v_j})$, so swapping between the top and bottom of the kei, we have from the inductive construction that $\rho(v_{k-1}, 1 - i_{v_{k-1}}) = \rho(v_{j+1}, i_{v_{j+1}})$. But then, by the minimality of k as

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

giving a repetition, we must have $j = k - 1$, so $(v_k, i_{v_k}) = (v_{k-1}, 1 - i_{v_{k-1}})$, which violates the fact from the construction that $\rho(v_k, i_{v_k}) \neq \rho(v_{k-1}, 1 - i_{v_{k-1}})$.

The set $\{v_j \mid j \in \mathbb{Z}\}$ may be finite or infinite, but the corresponding subset $\{\rho_V(v_j, i_{v_j}) \mid j \in \mathbb{Z}\}$ has the same cardinality: $(v_j, i_{v_j}) = (v_k, i_{v_k})$ if and only if $\rho(v_j, i_{v_j}) = \rho(v_k, i_{v_k})$. Note also that, for each k, the left multiplication maps $m_{\rho(v_k, 1-i_{v_k})}$ and $m_{\rho(v_{k+1}, i_{v_{k+1}})}$ on $\text{け}_{G'}$ are the same since $\rho(v_k, 1 - i_{v_k})$ and $\rho(v_{k+1}, i_{v_{k+1}})$ have the same first component. Therefore $m_{(v_k, 1-i_{v_k})}$ and $m_{(v_{k+1}, i_{v_{k+1}})}$ are the same on け_G. It follows that v_k and v_{k+1} have outward edges to the same other vertices in G, as well as to each other, and by induction the same is true of all members of the set $\{v_k \mid k \in \mathbb{Z}\}$; likewise, all members of the set $\{\rho_V(v_k, i_{v_k}) \mid k \in \mathbb{Z}\}$ have edges to one another and to the same other vertices.

The set F_G may be expressed as the disjoint union of such 'cycles' of vertices $\{v_k \mid k \in \mathbb{Z}\}$ by choosing a starting vertex v_0 in each cycle. With such choices made, we, in particular, have an assignment of i_v in $\{0, 1\}$ to each v in F_G, and we may define f on F_G by $f(v) = \rho_V(v, i_v)$. Clearly, with this definition, $f \upharpoonright F_G$ is a bijection from F_G to its image. Moreover its image is all of $F_{G'}$: if $(t, i) *_{G'} (w, 0) = (w, 0)$ for all (t, i) in $\text{け}_{G'}$, then $\rho^{-1}(t, i) *_G \rho^{-1}(w, 0) = \rho^{-1}(w, 0)$ for all (t, i) in $\text{け}_{G'}$, that is, $(u, j) *_G \rho^{-1}(w, 0) = \rho^{-1}(w, 0)$ for all (u, j) in け_G.

We have thus constructed a bijection $f : V_G \to V_{G'}$, and it remains to show that f is, in fact, a graph isomorphism from G to G'. So let u and v be vertices of G. If v is in F_G, then $f(v)$ is in $F_{G'}$, so both $u \, E_G \, v$ and $f(u) \, E_{G'} \, f(v)$ hold. Suppose v is in M_G. If u is in F_G, we have i_u in $\{0, 1\}$, as defined above; otherwise take $i_u = 0$. Then

$$(u, i_u) *_G (v, 0) = \begin{cases} (v, 0) & \text{if } u \, E_G \, v \text{ or } u = v, \\ (v, 1) & \text{otherwise,} \end{cases}$$

so

$$\rho(u, i_u) *_G \rho(v, 0) = \begin{cases} \rho(v, 0) & \text{if } u \, E_G \, v \text{ or } u = v, \\ \rho(v, 1) & \text{otherwise.} \end{cases}$$

Since the first component of $\rho(u, i_u)$ is $f(u)$ and the first component of $\rho(v, 0)$ is $f(v)$, we have that $f(u) \, E'_G \, f(v)$ if and only if $u \, E_G \, v$, which completes the proof that f is a graph isomorphism from G to G'. □

With Proposition 3.7, we have shown that, whilst not every isomorphism of kei け_G and $\text{け}_{G'}$ need arise from a graph isomorphism, such an isomorphism can be used to define a graph isomorphism of G and G', which, by the remark, gives rise to a (potentially different) isomorphism of け_G and $\text{け}'_G$. This completes the proof of Theorem 3.5. □

THEOREM 3.8. *The classes of kei, quandles, racks, left distributive algebras and algebras satisfying Σ are each Borel complete.*

PROOF. Implicit in the statement that these classes of structures are Borel complete is that we are considering the classes of countable such structures with underlying set $\mathbb{N}$, with each class topologised as described in Section 2.

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

The map $\text{け} : G \mapsto \text{け}_G$ from the class of graphs to the class of kei is not only Borel but is, in fact, continuous. Recall from Section 2 that the subbasic open sets in the space of graphs are of the form either $\{G \mid m\ E\ n\}$ or $\{G \mid m\ \not E\ n\}$. Similarly, for quandles with underlying set $\mathbb{N}$, the subbasic open sets are of the form $\{(\mathbb{N}, *) \mid u * v = w\}$ or $\{(\mathbb{N}, *) \mid u * v \neq w\}$. Then, by the construction of our dynamical kei, it is clear that the inverse image of any open set is open (as we defined $*$ in terms of the edge relation of E). Hence the map け taking G to け_G is continuous and so certainly Borel, and so, since the class of graphs is Borel complete, it follows that the class of kei is Borel complete. Moreover, since the kei form a subclass of the classes of quandles, of racks and of left distributive algebras, the map け likewise shows that the classes of quandles, of racks and of left distributive algebras are Borel complete.

Because the language of Σ is different from that of left distributive algebras, a different argument is needed to show that the class of algebras satisfying Σ is Borel complete. For this, we use the result of Mekler [20] that the class of groups is Borel complete (see [10, Section 2.3] for a sketch of the argument). As discussed after Definition 2.2, every group endowed with its conjugation operation and its group operation satisfies Σ. The inclusion map $(G, \circ) \mapsto (G, \circ, *)$, where $\circ$ denotes the group operation and $*$ denotes conjugation, is easily seen to be continuous and so is certainly Borel. Of course, since the group operation is one of the two operations in the language of Σ, and the other is conjugation which is determined by the group operation, two groups are isomorphic if and only if their corresponding structures satisfying Σ are isomorphic. We thus have that group isomorphism Borel reduces to isomorphism as algebras satisfying Σ, and therefore that the latter is Borel complete. □

4. Concluding remarks

We have shown that, in the sense of Borel reducibility, the complexity of the isomorphism problem for countable quandles is well above that of tame knots. As mentioned in the introduction, Kulikov [16] has shown, in contrast, that the equivalence relation of arbitrary knots is strictly greater in complexity than any first-order isomorphism relation. This intermediate status of quandle isomorphism relative to knot equivalence raises a number of questions.

Question 4.1. (a) Is there a natural characterisation of those quandles arising from tame knots? (b) If so, can the characterisation be used to simplify the computation of whether two such quandles are isomorphic?

While it is clear that the quandles associated with tame knots are finitely generated, it seems unlikely that this completely characterises them.

Even though wild knots are more complex in Borel reducibility terms than quandles, and the fundamental quandle map associates a quandle to every knot, this does not imply that all countable quandles can be obtained from a knot.

Question 4.2. Which quandles can be obtained as the quandle associated with a (possibly wild) knot?

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

Question 4.2 suggests the following one, which was first posed to us by Marcin Sabok.

QUESTION 4.3. Is there a natural class of knots broader than the tame knots that is completely classified by their associated quandles?

A standard way to think of a knot is as the image of an embedding of the circle S^1 into three-dimensional Euclidean space $\mathbb{R}^3$. One possible approach to Question 4.3 is to generalise by increasing the number of components and thus by considering links rather than knots, where a *link* is the image of an embedding of multiple copies of the circle into $\mathbb{R}^3$. In particular, Joyce's definition [13] of the fundamental quandle is also applicable to links. Call a link in $\mathbb{R}^3$ *locally tame* if its intersection with every compact subset of $\mathbb{R}^3$ is tame (that is, a finite union of pieces of tame knots). The fundamental quandle of a link with infinitely many components will not be finitely generated, so this will be a proper extension of the tame knot case if the following question has a positive answer.

QUESTION 4.4. Is the fundamental quandle a complete invariant for locally tame links in $\mathbb{R}^3$?

A different formalisation of the question of complexity is in a category-theoretic setting. Just as the class of graphs is maximal in the sense of Borel completeness (and, indeed, our proof made use of this fact), the *category* of graphs is universal in the sense that every algebraic category fully embeds into it [22, Theorem 5.3]. There are many such universality results for other categories—see, for example, [22]—raising the following natural question.

QUESTION 4.5. Does the category of graphs fully embed into the category of left distributive algebras?

Of course, the same question may also be asked of the category of racks, the category of quandles and the category of kei, in each case taking homomorphisms as the morphisms of the category. We note that the construction of $け_G$ from G in Theorem 3.5 is not even functorial in a natural way, since graph homomorphisms need not preserve nonedges. A potentially more problematic obstacle, however, is the fullness requirement—we have seen that dynamical kei admit many more homomorphisms than simply those arising from graph homomorphisms, at least in our construction. On the other hand, even if it turns out that the category of graphs cannot be fully embedded into the category of kei because kei always admit many homomorphisms, there may be interesting minimal-nonfullness, maximal-complexity results to be obtained in this direction. As an analogy, there can be no full embedding of the category of graphs into the category of abelian groups, as any two abelian groups A and B admit at least one homomorphism between them (the zero map) and the set of homomorphisms between them $\mathrm{Hom}(A, B)$ naturally forms an abelian group. Nevertheless, Przeździecki [21] has shown that there is an embedding $\mathcal{A}$ from the category of graphs into the category of abelian groups such that $\mathrm{Hom}(\mathcal{A}G, \mathcal{A}G')$ is

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press

the free abelian group generated by $\text{Hom}(G, G')$—the best possible result given these constraints.

Camerlo and Gao's result that isomorphism of countable Boolean algebras is Borel complete shows that Ketonen's classification of countable Boolean algebras uses objects for complete invariants that 'cannot be improved in an essential way' [3]. In contrast, our result illustrates the need for finer analysis of quandles as knot invariants.

References

[1] W. Bertram, *The Geometry of Jordan and Lie Structures*, Lecture Notes in Mathematics, 1754 (Springer, Berlin, 2000).

[2] E. Brieskorn, 'Automorphic sets and braids and singularities', in: *Braids (Santa Cruz, CA, 1986)*, Contemporary Mathematics, 78 (American Mathematical Society, Providence, RI, 1988), 45–115.

[3] R. Camerlo and S. Gao, 'The completeness of the isomorphism relation for countable Boolean algebras', *Trans. Amer. Math. Soc.* **353**(2) (2001), 491–518.

[4] P. Dehornoy, 'Braid groups and self-distributive operations', *Trans. Amer. Math. Soc.* **345**(1) (1994), 115–151.

[5] P. Dehornoy, *Braids and Self Distributivity*, Progress in Mathematics, 192 (Birkhäuser, Basel, 2000).

[6] P. Dehornoy, 'Elementary embeddings and algebra', in: *The Handbook of Set Theory* (eds. M. Foreman and A. Kanamori) (Springer, The Netherlands, 2010), 737–774.

[7] M. Elhamdadi and S. Nelson, *Quandles: An Introduction to the Algebra of Knots*, Student Mathematical Library, 74 (AMS, Providence, RI, 2015).

[8] I. Farah, A. Toms and A. Törnquist, 'Turbulence, orbit equivalence, and the classification of nuclear C^*-algebras', *J. reine angew. Math. (Crelle's Journal)* **688** (2014), 101–146.

[9] M. Foreman, D. Rudolph and B. Weiss, 'The conjugacy problem in ergodic theory', *Ann. Math.* **173** (2011), 1529–1586.

[10] H. Friedman and L. Stanley, 'A Borel reducibility theory for classes of countable structures', *J. Symbolic Logic* **54**(3) (1989), 894–914.

[11] S. Gao, *Invariant Descriptive Set Theory*, Pure and Applied Mathematics (Chapman & Hall/CRC Press, an imprint of Taylor & Francis Group, Boca Raton, 2009).

[12] G. Hjorth, *Classification and Orbit Equivalence Relations*, Mathematical Surveys and Monographs, 75 (American Mathematical Society, Providence, RI, 2002).

[13] D. Joyce, 'A classifying invariant of knots, the knot quandle', *J. Pure Appl. Algebra* **23** (1982), 37–65.

[14] S. Kamada, 'Knot invariants derived from quandles and racks', in: *Invariants of Knots and 3-Manifolds (Kyoto, 2001)*, Geometry and Topology Monographs, 4 (eds. T. Ohtsuki, T. Kohno, T. Le, J. Murakami, J. Roberts and V. Turaev) (Geometry and Topology Publications, Coventry, 2004).

[15] S. Kamada, 'Quandles derived from dynamical systems and subsets which are closed under quandle operations', *Topol. Appl.* **157** (2010), 298–301.

[16] V. Kulikov, 'A non-classification result for wild knots', *Trans. Amer. Math. Soc.* **369** (2017), 5829–5853.

[17] R. Laver, 'The left distributive law and the freeness of an algebra of elementary embeddings', *Adv. Math.* **91** (1992), 209–231.

[18] R. Laver, 'A division algorithm for the free left distributive algebra', in: *Logic Colloquium '90 (Helsinki, 1990)*, Lecture Notes in Logic, 2 (Springer, Berlin, 1993), 155–162.

[19] R. Laver and S. K. Miller, 'The free one-generated left distributive algebra: basics and a simplified proof of the division algorithm', *Cent. Eur. J. Math.* **11**(12) (2013), 2150–2175.

[20] A. H. Mekler, 'Stability of nilpotent groups of class 2 and prime exponent', *J. Symbolic Logic* **46**(4) (1981), 781–788.

[21] A. J. Przeździecki, 'An almost full embedding of the category of graphs into the category of abelian groups', *Adv. Math.* **257** (2014), 527–545.
[22] A. Pultr and V. Trnková, *Combinatorial, Algebraic and Topological Representations of Groups, Semigroups and Categories* (Academia, Prague, 1980).
[23] M. Takasaki, 'Abstraction of symmetric transformations', *Tohoku Math. J.* **49** (1943), 145–207.

ANDREW D. BROOKE-TAYLOR, School of Mathematics,
University of Leeds, Leeds LS2 9JT, UK
e-mail: andrewbt@gmail.com

SHEILA K. MILLER, 28 Archuleta Road,
Ranchos de Taos, NM 87557, USA
e-mail: sheila.miller@colorado.edu

https://doi.org/10.1017/S1446788719000399 Published online by Cambridge University Press