Logical Methods in Computer Science
Vol. 13(1:1)2017, pp. 1–11
https://lmcs.episciences.org/

Submitted Jan. 21, 2016
Published Jan. 23, 2017

LOGICAL COMPACTNESS AND CONSTRAINT SATISFACTION PROBLEMS

DANNY RORABAUGH [a], CLAUDE TARDIF [b], AND DAVID WEHLAU [c]

[a] Department of Mathematics and Statistics, Queen's University, Kingston, Ontario, Canada

[b,c] Department of Mathematics and Computer Science, Royal Military College of Canada, Kingston, Ontario, Canada

Abstract. We investigate a correspondence between the complexity hierarchy of constraint satisfaction problems and a hierarchy of logical compactness hypotheses for finite relational structures. It seems that the harder a constraint satisfaction problem is, the stronger the corresponding compactness hypothesis is. At the top level, the NP-complete constraint satisfaction problems correspond to compactness hypotheses that are equivalent to the ultrafilter axiom in all the cases we have investigated. At the bottom level, the simplest constraint satisfaction problems correspond to compactness hypotheses that are readily provable from the axioms of Zermelo and Fraenkel.

1. Introduction

A relational structure $\mathbb{A}$ is said to be *compact* if for any structure $\mathbb{B}$ of the same type, $\mathbb{B}$ admits a homomorphism to $\mathbb{A}$ whenever every finite substructure $\mathbb{B}'$ of $\mathbb{B}$ admits a homomorphism to $\mathbb{A}$. The compactness theorem of logic implies that every finite structure is compact. However, the compactness theorem is equivalent to the ultrafilter axiom, a consequence of the axiom of choice that is not provable from the axioms of Zermelo and Fraenkel. In this paper, we restrict our attention to the case where $\mathbb{A}$ is finite, but we do not assume the ultrafilter axiom. Instead, for each structure $\mathbb{A}$, we consider the statement "$\mathbb{A}$ is compact" as an hypothesis that is consistent with the axioms of Zermelo and Fraenkel, but not necessarily provable from these.

It turns out that the strength of such compactness hypotheses varies widely. For some structures, compactness can be proved from the axioms of Zermelo and Fraenkel. At the other extreme, for some other structures, compactness implies the ultrafilter axiom, hence the compactness of all finite structures. Perhaps it makes sense to call the latter structures "compactness-complete". This designation is borrowed from that of complexity classes, but our results also parallel complexity results: The structures $\mathbb{A}$ which we prove to be compactness-complete have their corresponding constraint satisfaction problems NP-complete. In contrast, the structures $\mathbb{A}$ for which we show that the statement "$\mathbb{A}$ is compact" is provable from the axioms of Zermelo and Fraenkel are the structures of "width one". The corresponding constraint satisfaction problems are arguably the simplest polynomial cases.

Key words and phrases: Compactness theorem, Relational structures, Constraint satisfaction problems.

DOI:10.23638/LMCS-13(1:1)2017

Also, the compactness of a structure implies that of any structure which can be primitively positively defined from it. Therefore it is possible that the algebraic approach to the complexity classification of constraint satisfaction problems would be relevant to the study of the hierarchy of compactness hypotheses as well.

The "dichotomy conjecture" of Feder and Vardi [1] states that every constraint satisfaction problem is either polynomial or NP-complete. There is no such dichotomy between the compactness hypotheses that are equivalent to the ultrafilter axiom and those that are provable from the axioms of Zermelo and Fraenkel. Indeed we will exhibit structures with polynomial constraint satisfaction problems, for which compactness hypotheses are equivalent to some intermediate "cardinal-specific" versions of the axiom of choice.

The axiom of choice and its variants grew out of the need to distinguish between existential and constructive aspects of mathematical proofs. In time, tools such as forcing theory became available to establish independence results. The possibility of connecting such independence results with constraint satisfaction problems is the main motivation of our investigation.

The paper is structured as follows. The next section will present the basics on homomorphisms of relational structures and an alternative characterisation of compactness. In the following two sections, we will present our results on compactness hypotheses derivable from the axioms of Zermelo and Fraenkel, and on compactness hypotheses equivalent to the ultrafilter axiom. Up to then, the relevant set-theoretic concepts are pretty standard, but afterwards we will present lesser known axioms weaker than the ultrafilter axiom, before moving on to intermediate compactness hypotheses.

2. Relational structures and compactness

A *type* is a finite set $\sigma = \{R_1, \dots, R_m\}$ of *relation symbols*, each with an *arity* r_i assigned to it. A σ-structure is a relational structure $\mathbb{A} = (A; R_1(\mathbb{A}), \dots, R_m(\mathbb{A}))$ where A is a nonempty set called the *universe* of $\mathbb{A}$, and $R_i(\mathbb{A})$ is an r_i-ary relation on A for each i. Let $\mathbb{A}$ and $\mathbb{B}$ be σ-structures, with universes A and B respectively. A *homomorphism* of $\mathbb{B}$ to $\mathbb{A}$ is a map $f : B \to A$ such that $f(R_i(\mathbb{B})) \subseteq R_i(\mathbb{A})$ for all $i = 1, \dots, m$. If B is a subset of A and the identity is a homomorphism from $\mathbb{B}$ to $\mathbb{A}$, then $\mathbb{B}$ is called a *substructure* of $\mathbb{A}$. The constraint satisfaction problem associated to a structure $\mathbb{A}$ is the problem of determining whether an input structure admits a homomorphism to $\mathbb{A}$.

As stated in the introduction, a relational structure $\mathbb{A}$ is called *compact* if for any structure $\mathbb{B}$ of the same type, $\mathbb{B}$ admits a homomorphism to $\mathbb{A}$ whenever every finite substructure $\mathbb{B}'$ of $\mathbb{B}$ admits a homomorphism to $\mathbb{A}$. The main result of this section is an alternative characterisation of compact structures in terms of "filter-tolerant" powers. Recall that a filter $\mathcal{F}$ on a set I is a family of nonempty subsets of I that is closed under intersection and contains every superset of each of its members. A filter which is maximal with respect to inclusion is called an ultrafilter. Equivalently, a filter is an ultrafilter if and only if it contains precisely one of each pair of complementary subset of I. The ultrafilter axiom states that every filter extends to an ultrafilter.

Let $\mathcal{F}$ be a filter on a set I, and $\mathbb{A}$ a relational structure of some type σ. The $\mathcal{F}$*-tolerant power* $\mathbb{A}^I_{\mathcal{F}}$ is the structure defined as follows. The universe of $\mathbb{A}^I_{\mathcal{F}}$ is the set A^I of all functions of I to the universe A of $\mathbb{A}$, and for each $R \in \sigma$ of arity k, $R(\mathbb{A}^I_{\mathcal{F}}) \subseteq (A^I)^k$ is the set of all k-tuples $(f_1, \dots, f_k)$ such that the set $\{i \in I | (f_1(i), \dots, f_k(i)) \in R(\mathbb{A})\}$ belongs to $\mathcal{F}$.

Proposition 2.1. *Let $\mathbb{A}$ be a finite relational structure. Then $\mathbb{A}$ is compact if and only if for every set I and every filter $\mathcal{F}$ on I, $\mathbb{A}^I_{\mathcal{F}}$ admits a homomorphism to $\mathbb{A}$.*

Proof. First, suppose that $\mathbb{A}$ is compact. Let $\mathbb{B}$ be a finite substructure of $\mathbb{A}^I_{\mathcal{F}}$. For every R in σ and every $(f_1,\dots,f_k) \in R(\mathbb{B})$, the set

$$S_{R,(f_1,\dots,f_k)} = \{i \in I | (f_1(i),\dots,f_k(i)) \in R(\mathbb{A})\}$$

is an element of $\mathcal{F}$. Thus, the finite intersection

$$S_{\mathbb{B}} = \bigcap \{S_{R,(f_1,\dots,f_k)} | R \in \sigma, (f_1,\dots,f_k) \in R(\mathbb{B})\}$$

is an element of $\mathcal{F}$ hence it is not empty. For every $i \in S_{\mathbb{B}}$, the map $\phi : \mathbb{B} \to \mathbb{A}$ defined by $\phi(f) = f(i)$ is a homomorphism, hence $\mathbb{B}$ admits a homomorphism to $\mathbb{A}$. If $\mathbb{A}$ is compact, this implies that $\mathbb{A}^I_{\mathcal{F}}$ admits a homomorphism to $\mathbb{A}$.

Now suppose that $\mathbb{A}^I_{\mathcal{F}}$ admits a homomorphism to $\mathbb{A}$ for every filter $\mathcal{F}$ on a set I. Let $\mathbb{B}$ be a structure such that every finite substructure of $\mathbb{B}$ admits a homomorphism to $\mathbb{A}$. Let I be the set of all maps from the universe of $\mathbb{B}$ to that of $\mathbb{A}$. For $R \in \sigma$ and $(x_1,\dots,x_k) \in R(\mathbb{B})$, let

$$(R,(x_1,\dots,x_k))^+ = \{i \in I | (i(x_1),\dots,i(x_k)) \in R(\mathbb{A})\}.$$

For a finite collection $\{(R_j,(x_{1,j},\dots,x_{k_j,j}))^+ | j = 1,\dots,\ell\}$ of these sets, let $\mathbb{B}'$ be the substructure of $\mathbb{B}$ spanned by $B' = \bigcup_{j=1}^{\ell} \{x_{1,j},\dots,x_{k_j,j}\}$. Then $\mathbb{B}'$ is a finite substructure of $\mathbb{B}$. By hypothesis, there exists a homomorphism $\phi : \mathbb{B}' \to \mathbb{A}$. Any extension of such a homomorphism ϕ to the universe of $\mathbb{B}$ belongs to $\bigcap\{(R_j,(x_{1,j},\dots,x_{k_j,j}))^+ | j = 1,\dots,\ell\}$. Thus, the family $\{(R,(x_1,\dots,x_k))^+ | R \in \sigma, (x_1,\dots,x_k) \in R(\mathbb{B})\}$ generates a filter $\mathcal{F}$ on I. By construction, the natural map $\psi : \mathbb{B} \to \mathbb{A}^I_{\mathcal{F}}$ defined by $\psi(x) = f$, where $f(i) = i(x)$, is a homomorphism. If there exists a homomorphism $\phi : \mathbb{A}^I_{\mathcal{F}} \to \mathbb{A}$, then $\phi \circ \psi : \mathbb{B} \to \mathbb{A}$ is a homomorphism. □

In view of this result, the standard derivation of the compactness of $\mathbb{A}$ from the ultrafilter axiom is a direct consequence of the following result.

Proposition 2.2. *If $\mathcal{F}$ is contained in an ultrafilter, then for every finite structure $\mathbb{A}$, $\mathbb{A}^I_{\mathcal{F}}$ admits a homomorphism to $\mathbb{A}$.*

Proof. Let $\mathcal{U}$ be an ultrafilter containing $\mathcal{F}$. Then for every f in A^I, there is a unique x in A such that $f^{-1}(x) \in \mathcal{U}$. We put $\phi(f) = x$. We show that ϕ is a homomorphism. Let $R \in \sigma$ be a relation of arity k, and $(f_1,\dots,f_k) \in R(\mathbb{A}^I_{\mathcal{F}})$. Then

$$S = \{i \in I | (f_1(i),\dots,f_k(i)) \in R(\mathbb{A})\} \in \mathcal{F} \subseteq \mathcal{U}.$$

Therefore the set $S \cap \left(\bigcap_{j=1}^k f_j^{-1}(\phi(f_j))\right)$ is in $\mathcal{U}$, hence it is not empty. For $i \in S \cap \left(\bigcap_{j=1}^k f_j^{-1}(\phi(f_j))\right)$, we have

$$(\phi(f_1),\dots,\phi(f_k)) = (f_1(i),\dots,f_k(i)) \in R(\mathbb{A}).$$ □

Our results on compactness will use the characterisation of Proposition 2.1, or its variant in terms of the natural quotient discussed next. Let $\sim$ be the equivalence defined on the universe of $\mathbb{A}^I_{\mathcal{F}}$ by $f \sim g$ if $\{i \in I | f(i) = g(i)\} \in \mathcal{F}$. Note that when $\mathcal{F}$ is an ultrafilter, $\mathbb{A}^I_{\mathcal{F}}/\!\sim$ is the standard ultrapower construction.

In general, for $(f_1,\dots,f_k) \in R$ and $g_j \sim f_j, j = 1,\dots,k$, we have $(g_1,\dots,g_k) \in R$. Thus, $\mathbb{A}^I_{\mathcal{F}}$ is the "lexicographic sum" of the equivalence classes of $\sim$. Any homomorphism from $\mathbb{A}^I_{\mathcal{F}}/\sim$ to $\mathbb{A}$ can be composed with the quotient map of $\mathbb{A}^I_{\mathcal{F}}$ to $\mathbb{A}^I_{\mathcal{F}}/\sim$. Conversely, when $\mathbb{A}$ is finite, any homomorphism $\phi : \mathbb{A}^I_{\mathcal{F}} \to \mathbb{A}$ allows to define a homomorphism $\psi : \mathbb{A}^I_{\mathcal{F}}/\sim \to \mathbb{A}$ as follows: given an ordering of the universe of $\mathbb{A}$, define $\psi(x/\sim)$ to be the smallest a such that there exists $y \in x/\sim$ with $\phi(y) = a$. Therefore we have the following

Proposition 2.3. *Let $\mathbb{A}$ be a finite relational structure. Then $\mathbb{A}$ is compact if and only if for every set I and every filter $\mathcal{F}$ on I, $\mathbb{A}^I_{\mathcal{F}}/\sim$ admits a homomorphism to $\mathbb{A}$.*

3. Structures of width one

Feder and Vardi [1] described various heuristics for finding homomorphisms to given relational structures. For such a heuristic, it is desirable to characterise the class of structures for which it is decisive. The simplest among these heuristics is the consistency-check algorithm of width one. This is the intuitive heuristic that it is unconsciously adopted by Sudoku puzzle solvers with no scientfic background. The structures for which width one consistency-check is decisive are called structures of width one, or structures with tree duality. We refer the reader to [1] for details. We will use the non-algorithmic structural characterisation given below.

For a structure $\mathbb{A}$, let $\mathcal{P}(\mathbb{A})$ be the structure defined as follows. The universe of $\mathcal{P}(\mathbb{A})$ is the set of nonempty subsets of the universe A of $\mathbb{A}$. For $R \in \sigma$ of arity k, $R(\mathcal{P}(\mathbb{A}))$ consists of the k-tuples $(S_1, \ldots, S_k)$ such that $\mathrm{pr}_i(R(\mathbb{A}) \cap (\Pi_{i=1}^k S_i)) = S_i$ for $i = 1, \ldots, k$. (Where pr_i is the i-th projection.) A structure $\mathbb{A}$ is said to have *width one* if there exists a homomorphism from $\mathcal{P}(\mathbb{A})$ to $\mathbb{A}$.

Proposition 3.1. *Every structure of width one is compact.*

Proof. Let $\mathbb{A}$ be a structure of width one. We will show that for every filter $\mathcal{F}$ on a set I, there exists a homomorphism from $\mathbb{A}^I_{\mathcal{F}}$ to $\mathbb{A}$.

For f in the universe of $\mathbb{A}^I_{\mathcal{F}}$, put

$$\mathcal{S}(f) = \{S \subseteq A : f^{-1}(S) \in \mathcal{F}\}$$

(where A is the universe of $\mathbb{A}$). In particular, $A \in \mathcal{S}(f)$, $\emptyset \notin \mathcal{S}(f)$, and $\mathcal{S}(f)$ is closed under intersections. Since A is finite, $\mathcal{S}(f)$ contains a smallest member $\phi(f)$. We show that $\phi : \mathbb{A}^I_{\mathcal{F}} \to \mathcal{P}(\mathbb{A})$ is a homomorphism.

Let $R \in \sigma$ be a relation of arity k, and $(f_1, \ldots, f_k)$ an element of $R(\mathbb{A}^I_{\mathcal{F}})$. Put

$$S = \{i \in I : (f_1(i), \ldots, f_k(i)) \in R(\mathbb{A})\}.$$

Then S is in $\mathcal{F}$, hence $T = S \cap \left(\bigcap_{i=1}^k f_i^{-1}(\phi(f_i))\right) \in \mathcal{F}$. For $i = 1, \ldots, k$, we have $f_i(T) \subseteq \phi(f_i)$, but by minimality of $\phi(f_i)$, the inclusion cannot be strict since $f_i^{-1}(f_i(T)) \in \mathcal{F}$. Thus every $x \in \phi(f_i)$ is the i-th coordinate of some $(f_1(t), \ldots, f_k(t))$ in $R(\mathbb{A})$, with $t \in T$ and $f_j(t) \in \phi(f_j)$ for $j = 1, \ldots, k$. In other words, $\mathrm{pr}_i(R(\mathbb{A}) \cap (\Pi_{i=1}^k \phi(f_i))) = \phi(f_i)$ for $i = 1, \ldots, k$, whence $(\phi(f_1), \ldots, \phi(f_k)) \in R(\mathcal{P}(\mathbb{A}))$. This shows that ϕ is a homomorphism.

Since $\mathbb{A}$ has width one, there exists a homomorphism ψ of $\mathcal{P}(\mathbb{A}) \to \mathbb{A}$. The composition $\psi \circ \phi : \mathbb{A}^I_{\mathcal{F}} \to \mathbb{A}$ is then a homomorphism. Thus, for every filter $\mathcal{F}$ on a set I, there exists a homomorphism from $\mathbb{A}^I_{\mathcal{F}}$ to $\mathbb{A}$. By Proposition 2.1, this implies that $\mathbb{A}$ is compact. $\square$

It is not clear whether other structures can be proved to be compact within the axioms of Zermelo-Fraenkel.

Problem 3.2. Let $\mathbb{A}$ be a structure for which compactness follows from the axioms of Zermelo-Fraenkel. Does $\mathbb{A}$ necessarily have width one?

4. Compactness results equivalent to the ultrafilter axiom

In graph theory, the De Bruijn-Erdős theorem is the statement that a graph is k-colourable if and only if all of its finite subgraphs are k-colourable. In our terminology, this is the statement that the complete graphs are compact. (The complete graph $\mathbb{K}_n$ on n vertices is the structure with universe $\{0, \ldots, n-1\}$ and the binary adjacency relation $\neq$.) Various proofs were known in the early fifties. Then in 1971, Läuchli [5] proved that the ultrafilter axiom is a consequence of the compactness of $\mathbb{K}_n$ for any $n \geq 3$.

In this section, we present our proof of Läuchli's result, and extend it to many relational structures using primitive positive definability. Incidentally, the complete graphs with at least three elements correspond to the first constraint satisfaction problems that were shown to be NP-complete.

Lemma 4.1. *Let $n \geq 3$ be an integer and $\mathcal{F}$ a filter on set I. If $(\mathbb{K}_n)^I_{\mathcal{F}}$ admits a homomorphism to $\mathbb{K}_n$, then $\mathcal{F}$ is contained in an ultrafilter.*

Proof. Let $\phi : (\mathbb{K}_n)^I_{\mathcal{F}} \to \mathbb{K}_n$ be a homomorphism. We write id_k for the constant function with constant value $k \in \{0, \ldots, n-1\}$. Since the restriction of ϕ to the constant functions is bijective, we can assume without loss of generality that $\phi(\mathrm{id}_k) = k$ for all $k \in \{0, \ldots, n-1\}$. If $n > 3$, every element of $(\mathbb{K}_3)^I_{\mathcal{F}} \subseteq (\mathbb{K}_n)^I_{\mathcal{F}}$ is adjacent to id_k for all $k \geq 3$. Therefore the restriction of ϕ to $(\mathbb{K}_3)^I_{\mathcal{F}}$ is an idempotent homomorphism to $\mathbb{K}_3$. Therefore we can assume that $n = 3$.

For $X \subseteq I$, we write $\mathbb{1}_X$ for the characteristic map of X, that is, $\mathbb{1}_X(i) = 1$ if $i \in X$ and $\mathbb{1}_X(i) = 0$ otherwise. Then $\mathbb{1}_X$ is adjacent to id_2, therefore $\phi(\mathbb{1}_X) \in \{0, 1\}$ for all $X \subseteq I$. Put

$$\mathcal{U} = \{X | \phi(\mathbb{1}_X) = 1\}.$$

We will show that $\mathcal{U}$ is an ultrafilter containing $\mathcal{F}$.

For $F \in \mathcal{F}$, $\mathbb{1}_F$ is adjacent to id_0. Since $\phi(\mathrm{id}_0) = 0$, we have $\phi(\mathbb{1}_F) = 1$. Thus, $\mathcal{F} \subseteq \mathcal{U}$. Also, for any $X \subseteq I$, $\mathbb{1}_X$, $\mathbb{1}_{\overline{X}}$ and id_2 are mutually adjacent (where $\overline{X}$ denotes the complement of X). Since $\phi(\mathrm{id}_2) = 2$, we must have $\{\phi(\mathbb{1}_X), \phi(\mathbb{1}_{\overline{X}})\} = \{0, 1\}$, that is, $\mathcal{U}$ contains precisely one of X and $\overline{X}$.

For $X \in \mathcal{U}$, define $f_X, g_X : I \to \{1, 2\}$ by

$$(f_X(i), g_X(i)) = \begin{cases} (2, 1) & \text{if } i \in X, \\ (1, 2) & \text{otherwise.} \end{cases}$$

Then f_X is adjacent to id_0 and $\mathbb{1}_X$, thus $\phi(f_X) = 2$. Since g_X is adjacent to id_0 and f_X, we then have $\phi(g_X) = 1$. Now for any $Y \subseteq I$ containing X, $\mathbb{1}_{\overline{Y}}$ is adjacent to g_X, thus $\phi(\mathbb{1}_{\overline{Y}}) = 0$ and $\phi(\mathbb{1}_Y) = 1$. This shows that if Y contains $X \in \mathcal{U}$, then $Y \in \mathcal{U}$.

For $X, Y \in \mathcal{U}$, define $f_{X \cap Y}, f_{X \setminus Y}, f_{\overline{X}} : I \to \{0, 1, 2\}$ by

$$(f_{X\cap Y}(i), f_{X\setminus Y}(i), f_{\overline{X}}(i)) = \begin{cases} (0, 1, 2) & \text{if } i \in X \cap Y, \\ (2, 0, 1) & \text{if } i \in X \setminus Y, \\ (1, 2, 0) & \text{if } i \in \overline{X}. \end{cases}$$

Then $f_{X\cap Y}, f_{X\setminus Y}$ and $f_{\overline{X}}$ are mutually adjacent. Therefore

$$\{\phi(f_{X\cap Y}), \phi(f_{X\setminus Y}), \phi(f_{\overline{X}})\} = \{0, 1, 2\}.$$

Since $f_{\overline{X}}$ is adjacent to $\mathbb{1}_{\overline{X}}$ and $\phi(\mathbb{1}_{\overline{X}}) = 0$, we have $\phi(f_{\overline{X}}) \neq 0$. Similarly, $X\setminus Y \subseteq \overline{Y}$ whence $\phi(\mathbb{1}_{X\setminus Y}) = 0$, and $\mathbb{1}_{X\setminus Y}$ is adjacent to $f_{X\setminus Y}$, so that $\phi(f_{X\setminus Y}) \neq 0$. Therefore $\phi(f_{X\cap Y}) = 0$. Since $\mathbb{1}_{X\cap Y}$ is adjacent to $f_{X\cap Y}$, we then have $\phi(\mathbb{1}_{X\cap Y}) = 1$, that is, $X \cap Y \in \mathcal{U}$. This shows that $\mathcal{U}$ is an ultrafilter. $\square$

The above proof is essentially the correspondence between the n-colourings of powers of $\mathbb{K}_n$, $n \geq 3$ and the 0-1 measures on a set established by Greenwell and Lovász [3].

Corollary 4.2 (Läuchli [5])**.** *For every $n \geq 3$, the ultrafilter axiom is equivalent to the statement that $\mathbb{K}_n$ is compact.*

We can then expand the class of structures for which compactness is equivalent to the ultrafilter axiom by using primitive positive definitions. The same method also allows reductions amongst constraint satisfaction problems. Let σ be a type and $\mathbb{A}$ a σ-structure with universe $\{0, \dots, n-1\}$, that is, the same universe as that of $\mathbb{K}_n$. For a σ-structure $\mathbb{B}$ with distinguished elements x, y, the binary relation $R_{(\mathbb{B},x,y)} \subseteq \{0, \dots, n-1\}^2$ is defined by

$$R_{(\mathbb{B},x,y)} = \{(\phi(x), \phi(y)) \mid \phi : \mathbb{B} \to \mathbb{A} \text{ is a homomorphism } \}.$$

Then $\mathbb{K}_n$ is said to be *primitively positively definable* from $\mathbb{A}$ if for some $(\mathbb{B}, x, y)$, the adjacency relation $\neq$ of $\mathbb{K}_n$ coincides with $R_{(\mathbb{B},x,y)}$. For instance, if $\mathbb{B}$ is an undirected path with three edges and x, y are its endpoints, then $R_{(\mathbb{B},x,y)}$ on the undirected cycle $\mathbb{A}$ with five vertices is the adjacency relation of $\mathbb{K}_5$.

Proposition 4.3. *Let $\mathbb{A}$ be a structure with universe $\{0, \dots, n-1\}$, where $n \geq 3$. If $\mathbb{K}_n$ is primitively positively definable from $\mathbb{A}$, then the ultrafilter axiom is equivalent to the statement that $\mathbb{A}$ is compact.*

Proof. Suppose that $\phi : \mathbb{A}^I_{\mathcal{F}} \to \mathbb{A}$ is a homomorphism. Then ϕ is a map from the universe of $(\mathbb{K}_n)^I_{\mathcal{F}}$ to that of $\mathbb{K}_n$. We will show that ϕ is a homomorphism of $(\mathbb{K}_n)^I_{\mathcal{F}}$ to $\mathbb{K}_n$.

By hypothesis, the adjacency relation $\neq$ of $\mathbb{K}_n$ coincides with some $R_{(\mathbb{B},x,y)}$. For each $i, j \in \{0, \dots, n-1\}$ with $i \neq j$, we can fix $\psi_{(i,j)} : \mathbb{B} \to \mathbb{A}$ such that $\psi_{(i,j)}(x) = i$ and $\psi_{(i,j)}(y) = j$.

Now let f, g be adjacent elements of the universe of $(\mathbb{K}_n)^I_{\mathcal{F}}$, and put

$$J_{(f,g)} = \{i \in I \mid f(i) \neq g(i)\}.$$

Consider the map $\psi_{(f,g)} : \mathbb{B} \to \mathbb{A}^I_{\mathcal{F}}$ defined by $\psi_{(f,g)}(z) = h$, where

$$h(i) = \begin{cases} \psi_{(f(i),g(i))}(z) \text{ if } i \in J_{(f,g)}, \\ f(i) \text{ otherwise.} \end{cases}$$

Then $\psi_{(f,g)}$ is a homomorphism from $\mathbb{B}$ to $\mathbb{A}^I_{\mathcal{F}}$ since the set of coordinates on which all relations are preserved contains $J_{(f,g)}$, which is a member of $\mathcal{F}$. Therefore $\phi \circ \psi_{(f,g)} : \mathbb{B} \to \mathbb{A}$ is a homomorphism. Since $R_{(\mathbb{B},x,y)}$ is the adjacency relation of $\mathbb{K}_n$, we then have

$$\phi(f) = \phi \circ \psi_{(f,g)}(x) \neq \phi \circ \psi_{(f,g)}(y) = \phi(g).$$

This shows that ϕ is a homomorphism of $(\mathbb{K}_n)^I_{\mathcal{F}}$ to $\mathbb{K}_n$.

Therefore, if $\mathbb{A}$ is compact, then $\mathbb{K}_n$ is compact, and the ultrafilter axiom holds. $\square$

Most structures are "projective" (see [9]), hence they primitively positively define the complete graph on their universe. Thus the method used to prove that most constraint satisfaction problems are NP-complete also shows that for most structures, compactness implies the ultrafilter axiom. It would be interesting to see whether the correspondence can be pushed further.

Problem 4.4. Let $\mathbb{A}$ be a structure for which compactness implies the ultrafilter axiom. Is the corresponding constraint satisfaction problem necessarily NP-complete?

The algebraic approach to the dichotomy conjecture proposes a criterion for determining precisely which constraint satisfaction problems are polynomial in terms of polymorphisms of relational structures (see [10]). To answer the above problem affirmatively, it would be sufficient to show that when the criterion is not satisfied on a structure $\mathbb{A}$, then the compactness of $\mathbb{A}$ implies the ultrafilter axiom. However, the converse cannot be proved without proving that NP is different from P, since some compactness hypotheses are provably not equivalent to the ultrafilter axiom.

5. Axioms of set theory

In this section we present a few axioms that are weaker than the ultrafilter axiom, but not provable from the axioms of Zermelo and Fraenkel. The results of this section can be found in "The Axiom of Choice" by Thomas Jech [4] and "Zermelo's Axiom of Choice - Its Origins, Development and Influence" by Gregory H. Moore [6]. Our purpose in considering such axioms is to obtain independence results without going into forcing theory. We will use the following axioms.

Order extension principle. Every partial ordering of a set X can be extended to a linear ordering of X.

Ordering principle. Every set can be linearly ordered.

Axiom of choice for finite sets. For every set X of nonempty finite sets, there is a function $f : X \to \cup X$ such that $f(x) \in x$ for all $x \in X$.

The ultrafilter axiom implies the order extension principle, which implies the ordering principle, which implies the axiom of choice for finite sets. None of the implications is reversible. For $n \geq 2$, the axiom of choice for n-sets is the following statement.

Choice(n): For every set X of sets of cardinality n, there is a function $f : X \to \cup X$ such that for each $x \in X$, $f(x)$ is an element of x.

The axiom of choice for finite sets implies the conjunction of Choice(n) for all n, but the implication is not reversible. There are various dependencies amongst the axioms of choice for finite sets. For instance, Choice(kn) implies Choice(n) for all k. Also, Choice(2) is equivalent to Choice(4), but independent from Choice(3). Research about such dependencies seems to have been an active area up to the seventies, as indicated in the exercises to chapter 7 of [4]. It culminated in the following result.

Proposition 5.1 (Gauntt [2])**.** *Let $m \geq 2$ be an integer and S a set of integers. Then* Choice(m) *follows from the conjunction of* Choice(n), $n \in S$ *if and only if any fixed-point free subgroup G of the symmetric group S_n contains a fixed-point free subgroup H and a finite sequence $H_1, \ldots, H_k$ of proper subgroups of H such that the sum of indices $[H : H_1] + \cdots + [H : H_k]$ belongs to S.*

The condition in Proposition 5.1 had been formulated by Mostowski [7], who proved its sufficiency. We note that [2] is only an announcement of the proof of necessity. Apparently the proof has never been published, but the result seems to be accepted by the community. An earlier result states that Choice(m) follows from the conjunction of Choice(k), $2 \leq k \leq n$ if and only if whenever m is expressed as a sum of primes, one of the primes is at most n. In particular, this implies that none of the statements Choice(k) can be proved from the axioms of Zermelo and Fraenkel.

One further axiom we need to consider is the following.

Kinna-Wagner Principle. For every set X of sets of cardinality at least 2, there is a function $f : X \to \mathcal{P}(\cup X)$ such that for all $x \in X$, $f(x)$ is a nonempty proper subset of x.

The Kinna-Wagner Principle implies the ordering principle, but is independent from the ultrafilter axiom. We could not find references to cardinality-specific versions of the Kinna-Wagner principle, but these will be useful to us.

KW(n): For every set X of sets of cardinality n, there is a function $f : X \to \mathcal{P}(\cup X)$ such that for all $x \in X$, $f(x)$ is a nonempty proper subset of x.

Clearly, there is no loss of generality in assuming that $|f(x)| \leq |x|/2$ for all $x \in X$. Thus, KW(2) is equivalent to Choice(2) and KW(3) is equivalent to Choice(3). In general, KW(n) can be weaker than Choice(n), but since Choice(n) follows from KW(n) and the conjunction of all Choice(k), $k \leq n/2$, KW(n) is still independent from the axioms of Zermelo and Fraenkel.

6. Intermediate compactness results

For $n \geq 2$, the directed n-cycle $\mathbb{C}_n$ is the structure with universe $\mathbb{Z}_n = \{0, 1, \ldots, n-1\}$ and one binary relation $R(\mathbb{C}_n) = \{(i, i+1) \mid i \in \mathbb{Z}_n\}$. It has long been known that the constraint satisfaction problem associated with $\mathbb{C}_n$ is polynomial for each n. In this section, we show that the hypothesis that $\mathbb{C}_n$ is compact is weaker than the ultrafilter axiom, but not provable from the axioms of Zermelo and Fraenkel.

Proposition 6.1. *For $n \geq 2$,* Choice(n) *implies that $\mathbb{C}_n$ is compact.*

Proof. Let $\mathcal{F}$ be a filter on a set I. First note that for the equivalence relation $\sim$ of Proposition 2.3, $(\mathbb{C}_n)^I_{\mathcal{F}}/\!\sim$ is a disjoint union of copies of $\mathbb{C}_n$. Indeed, we have $f \to g$ if and only if $\{i \in I | g(i) = f(i) + 1\} \in \mathcal{F}$ (where addition is modulo n).

Let $\mathcal{C}$ be the set of connected components of $(\mathbb{C}_n)^I_{\mathcal{F}}/\sim$. For each $C \in \mathcal{C}$, the set of homomorphisms from C to $\mathbb{C}_n$ has cardinality exactly n. If we assume Choice(n), we can select a homomorphism $\phi_C : C \to \mathbb{C}_n$ for each $C \in \mathcal{C}$. Then, $\phi = \bigcup_{C \in \mathcal{C}} \phi_C$ is a homomorphism of $(\mathbb{C}_n)^I_{\mathcal{F}}/\sim$ to $\mathbb{C}_n$. This implies that $(\mathbb{C}_n)^I_{\mathcal{F}}$ admits a homomorphism to $\mathbb{C}_n$, hence $\mathbb{C}_n$ is compact by Proposition 2.3. □

Thus, the hypothesis that $\mathbb{C}_n$ is compact is weaker than the ultrafilter axiom. The next two results show that it cannot be proved from the axioms of Zermelo and Fraenkel.

Proposition 6.2. *Let p be a prime integer. If $\mathbb{C}_p$ is compact, then* KW(p) *holds.*

Proof. Let X be a set of sets of cardinality p. Let I be the set of partial choice functions on X, that is,

$$I = \{c : Y \to \cup X \mid Y \subseteq X \text{ and } c(x) \in x \text{ for all } x \in Y\}.$$

For $x \in X$, put

$$x^+ = \{c \in I \mid x \text{ is in the domain of } c\}.$$

Then the family $\{x^+ \mid x \in X\}$ is closed under finite intersection, hence it generates a filter $\mathcal{F}$ on I. By hypothesis, there exists a homomorphism ϕ of $(\mathbb{C}_p)^I_{\mathcal{F}}$ to $\mathbb{C}_p$. We will use ϕ to construct a function $f : X \to \mathcal{P}(\cup(X))$ such that for all $x \in X$, $f(x)$ is a nonempty proper subset of x.

For $x \in X$ and $j \in x$, put

$$x^+_j = \{c \in I \mid x \text{ is in the domain of } c \text{ and } c(x) = j\}.$$

Then I can be partitioned into the sets x^+_j, $j \in x$, and x^-, where x^- is the set of elements of I that do not have x in their domain. Since x is a set of cardinality p, there exist $p!$ bijections of x to $\{0, \dots, p-1\}$. For a bijection $\psi : x \to \{0, \dots, p-1\}$, we let f_ψ be the element of the universe of $(\mathbb{C}_p)^I_{\mathcal{F}}$ defined by $f_\psi(i) = \psi(j)$ if $i \in x^+_j$, and $f_\psi(i) = 0$ if $i \in x^-$. Then $\{f_\psi, f_{\psi+1}, \dots, f_{\psi+p-1}\}$ is a copy of $\mathbb{C}_p$ in $(\mathbb{C}_p)^I_{\mathcal{F}}$, where $\psi + k : x \to \{0, \dots, p-1\}$ is defined by $\psi + k(j) = \psi(j) + k$. Therefore ϕ maps $\{f_\psi, f_{\psi+1}, \dots, f_{\psi+p-1}\}$ bijectively to the universe of $\mathbb{C}_p$. Thus there exists exactly one $k \in \{0, \dots, p-1\}$ such that $\phi(f_{\psi+k}) = 0$. We call $(\psi + k)^{-1}(0) \in x$ the *element of x distinguished by* $\{\psi, \psi + 1, \dots, \psi + p - 1\}$.

The $p!$ bijections of x to $\{0, \dots, p-1\}$ are partitionned into $(p-1)!$ classes of the form $\{\psi, \psi + 1, \dots, \psi + p - 1\}$. The number of times an element of x appears as distinguished element is not constant, since p does not divide $(p-1)!$. Therefore the set y_x of elements of x that appear the most times as distinguished element is a nonempty proper subset of x. Therefore the function $f : X \to \mathcal{P}(\cup X)$ defined by $f(x) = y_x$ satisfies the required properties. □

The fact that Choice(2) = KW(2) is equivalent to the compactness of $\mathbb{C}_2$ was proved by Mycielski [8]. (Note that $\mathbb{C}_2 = \mathbb{K}_2$.) The above results show that Choice(3) = KW(3) is equivalent to the compactness of $\mathbb{C}_3$. It is not hard to show that for prime p, KW(p) is equivalent to the compactness of $\mathbb{C}_p$, and that for some composite numbers n, Choice(n), KW(n) and the compactness of $\mathbb{C}_n$ are inequivalent hypotheses. However, for our purposes, it is sufficient to show that the compactness of $\mathbb{C}_n$ can never be proved from the axioms of Zermelo and Fraenkel. Our next result completes the proof.

Proposition 6.3. *$\mathbb{C}_n$ is compact if and only if $\mathbb{C}_p$ is compact for every prime factor p of n.*

Proof. We first show that if $\mathbb{C}_{kp}$ is compact, then $\mathbb{C}_p$ is compact. Suppose that $\phi : (\mathbb{C}_{kp})^I_{\mathcal{F}} \to \mathbb{C}_{kp}$ is a homomorphism. For f in the universe of $(\mathbb{C}_p)^I_{\mathcal{F}}$, we let kf be the element of the universe of $(\mathbb{C}_{kp})^I_{\mathcal{F}}$ defined by $kf(i) = k \cdot f(i)$. We define $\psi : (\mathbb{C}_p)^I_{\mathcal{F}} \to \mathbb{C}_p$ by letting $\psi(f)$ be the unique $j \in \mathbb{Z}_p$ such that $\phi(kf) \in \{kj, kj+1, \dots, kj+k-1\}$. If $(f, g) \in R((\mathbb{C}_p)^I_{\mathcal{F}})$, then there is a directed path of length k from kf to kg in $(\mathbb{C}_{kp})^I_{\mathcal{F}}$. Therefore, $\psi(g) = \psi(f) + 1$. This shows that ψ is a homomorphism. Therefore if $\mathbb{C}_{kp}$ is compact, then $\mathbb{C}_p$ is compact.

We next show that if p and q are relatively prime, and $\mathbb{C}_p$, $\mathbb{C}_q$ are compact, then $\mathbb{C}_{pq}$ is compact. This easily follows from the fact that $\mathbb{C}_{pq}$ is naturally isomorphic to the product $\mathbb{C}_p \times \mathbb{C}_q$ with universe $\mathbb{Z}_p \times \mathbb{Z}_q$ and relation $R(\mathbb{C}_p \times \mathbb{C}_q) = \{((i,j), (i+1, j+1)) \mid (i,j) \in \mathbb{Z}_p \times \mathbb{Z}_q\}$. Thus, a homomorphism ϕ to $\mathbb{C}_{pq}$ naturally corresponds to a pair ϕ_p, ϕ_q of homomorphisms to $\mathbb{C}_p$ and $\mathbb{C}_q$ respectively. The definition of compactness then directly implies that if $\mathbb{C}_p$ and $\mathbb{C}_q$ are compact, then $\mathbb{C}_{pq}$ is compact.

It only remains to show that if p is prime and $\mathbb{C}_p$ is compact, then $\mathbb{C}_{p^k}$ is compact for all $k \geq 1$. We proceed by induction on k. Suppose that $\mathbb{C}_p$ and $\mathbb{C}_{p^k}$ are compact. There is a homomorphism ϕ of $\mathbb{C}_{p^{k+1}}$ to $\mathbb{C}_{p^k}$ corresponding to the natural quotient of $\mathbb{Z}_{p^{k+1}}$ to $\mathbb{Z}_{p^k}$. Applying ϕ coordinatewise yields a homomorphism $\phi' : \left(\mathbb{C}_{p^{k+1}}\right)^I_{\mathcal{F}} \to \left(\mathbb{C}_{p^k}\right)^I_{\mathcal{F}}$. If $\mathbb{C}_{p^k}$ is compact, there exists a homomorphism $\phi'' : \left(\mathbb{C}_{p^k}\right)^I_{\mathcal{F}} \to \mathbb{C}_{p^k}$. Thus $\phi'' \circ \phi' : \left(\mathbb{C}_{p^{k+1}}\right)^I_{\mathcal{F}} \to \mathbb{C}_{p^k}$ is a homomorphism. The natural quotient by the equivalence $\sim$ of Proposition 2.3 yields a homomorphism $\psi : \left(\mathbb{C}_{p^{k+1}}\right)^I_{\mathcal{F}}/\!\sim\, \to \mathbb{C}_{p^k}$. The connected components of $\left(\mathbb{C}_{p^{k+1}}\right)^I_{\mathcal{F}}/\!\sim$ are copies of $\mathbb{C}_{p^{k+1}}$. Each of these connected components admits exactly p homomorphisms to $\mathbb{C}_{p^{k+1}}$ with the property that their composition with ϕ coincides with ψ. We need to choose one of these homomorphisms for each connected component.

We cannot prove that the compactness of $\mathbb{C}_p$ implies Choice(p) for p larger than 3, but in the present context, the full force of Choice(p) is not required. Let $\mathbb{A}$ be the structure with universe $\psi^{-1}(0)$, and one binary relation consisting of the pairs (f, g) such that there is a directed path of length p^k from f to g in $\left(\mathbb{C}_{p^{k+1}}\right)^I_{\mathcal{F}}/\!\sim$. The connected components of $\mathbb{A}$ are copies of $\mathbb{C}_p$, one in each connected component of $\left(\mathbb{C}_{p^{k+1}}\right)^I_{\mathcal{F}}/\!\sim$. Similarly, let $\mathbb{B}$ be the structure with universe $\phi^{-1}(0)$, and one binary relation consisting of the pairs (i, j) such that there is a directed path of length p^k from i to j in $\mathbb{C}_{p^{k+1}}$. Then $\mathbb{B}$ is a copy of $\mathbb{C}_p$. Since $\mathbb{C}_p$ is compact, there exists a homomorphism $\chi^* : \mathbb{A} \to \mathbb{B}$. We can then define $\chi : \left(\mathbb{C}_{p^{k+1}}\right)^I_{\mathcal{F}}/\!\sim\, \to \mathbb{C}_{p^{k+1}}$ componentwise as the unique homomorphism with the property that $\chi(f) = \chi^*(f)$ for all $f \in \psi^{-1}(0)$. □

The results above imply that the conjunction of Choice(n) for all n is equivalent to the compactness of $\mathbb{C}_n$ for all n. There are other structures that can be proved to be compact from the conjunction of Choice(n) for all n. However it seems unlikely that this is the case for all structures that are not compactness complete. We conclude this section with one possible witness to this assertion.

Proposition 6.4. *Let $\mathbb{A}$ be the structure on the universe $\{0, 1\}$ with the two binary relations $\neq$ and $\leq$. The order extension principle implies that $\mathbb{A}$ is compact.*

Proof. Let $\mathcal{F}$ be a filter on a set I. We will show that the order extension principle implies that $\mathbb{A}^I_{\mathcal{F}}/\!\sim$ admits a homomorphism to $\mathbb{A}$, where $\sim$ is the equivalence of Proposition 2.3. The elements of $\mathbb{A}^I_{\mathcal{F}}$ will be represented by their characteristic function. In $\mathbb{A}^I_{\mathcal{F}}/\!\sim$, the only element adjacent (under $\neq$) to $\mathbb{1}_X/\!\sim$ is $\mathbb{1}_{\overline{X}}/\!\sim$. The relation $\leq$ of $\mathbb{A}^I_{\mathcal{F}}$ is characterised by

$$\mathbb{1}_X/\!\sim\, \leq \mathbb{1}_Y/\!\sim\, \Leftrightarrow \{i \in I \mid \mathbb{1}_X(i) \leq \mathbb{1}_Y(i)\} \in \mathcal{F}.$$

The order extension principle implies that it extends to a linear order $\preceq$ on $\mathbb{A}^I_{\mathcal{F}}/\!\sim$. We define $\phi : \mathbb{A}^I_{\mathcal{F}}/\!\sim\, \to \mathbb{A}$ by

$$\phi(\mathbb{1}_X/\!\sim) = \begin{cases} 0 \text{ if } \mathbb{1}_X/\!\sim\, \preceq \mathbb{1}_{\overline{X}}/\!\sim, \\ 1 \text{ otherwise.} \end{cases}$$

For all $\mathbb{1}_X/\!\sim$, we have $\{\phi(\mathbb{1}_X/\!\sim), \phi(\mathbb{1}_{\overline{X}}/\!\sim)\} = \{0, 1\}$. Therefore ϕ preserves the adjacency relation $\neq$. For $\mathbb{1}_X/\!\sim\, \leq \mathbb{1}_Y/\!\sim$, if $\phi(\mathbb{1}_Y/\!\sim) = 0$, then $\mathbb{1}_X/\!\sim\, \leq \mathbb{1}_Y/\!\sim\, \preceq \mathbb{1}_{\overline{Y}} \leq \mathbb{1}_{\overline{X}}$, hence $\phi(\mathbb{1}_X/\!\sim) = 0$. Therefore ϕ preserves $\leq$. This shows that ϕ is a homomorphism. □

We have not been able to prove or disprove that the compactness of the structure $\mathbb{A}$ in Proposition 6.4 follows from the conjunction of Choice(n) for all n. On the other hand, the hypothesis that $\mathbb{A}$ is compact seems to be fairly close to the order extension principle: It implies that every order extends to a relation R that is trichotomic and has the property that if (x, y), (y, z), (z, x) belong to R, then $\{x, y, z\}$ is an antichain in the order. Again we cannot prove or disprove that the latter property is weaker than the order extension principle. Perhaps the trick of comparing to previously investigated axioms has its limits, and forcing theory would be needed at some point.

7. Conclusion

We have established a partial correspondence between the hierarchy of strength of compactness hypotheses and the complexity hierarchy of constraint satisfaction problems. A positive answer to Problem 4.4 would strenghten this correspondence. The conjectured algebraic criterion for characterising NP-complete constraint satisfaction problems may be the key to settling this question. Polynomial constraint satisfaction problems do not all correspond to equivalent compactness hypotheses. This is a bit unsettling. Is there an algorithmic significance to the varying strength of compactness hypotheses? More significantly, is there a possible interplay between descriptive complexity and axiomatic set theory that goes beyond a simple correspondence?

Acknowledgements. The authors wish to thank Jaroslav Nešetřil for pointing out the result in reference [3]. The authors gratefully acknowledge the partial support of NSERC for this work.

References

[1] T. Feder, M. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. SIAM J. Comput. 28 (1998), 57–104.
[2] R. J. Gauntt, Axiom of choice for finite sets - A solution to a problem of Mostowski, Notices Amer. Math. Soc. 17 (1970), 454.
[3] D. Greenwell, L. Lovász, Applications of product colouring. Acta Math. Acad. Sci. Hungar. 25 (1974), 335–340.
[4] T. Jech, The Axiom of Choice, Studies in Logic and the Foundations of Mathematics, Vol. 75, North-Holland Publishing Co., Amsterdam-London, American Elsevier Publishing Co., New York, 1973, xi+202 pp.
[5] H. Läuchli, Coloring infinite graphs and the Boolean prime ideal theorem, Israel J. Math. 9 (1971), 422–429.
[6] G. H. Moore, Zermelo's Axiom of Choice. Its Origins, Development, and Influence, Studies in the History of Mathematical and Physical Sciences 8, Springer-Verlag, New York, 1982, xiv+410 pp.
[7] A. Mostowski, Axiom of choice for finite sets, Fund. Math. 33 (1945) 137–168.
[8] J. Mycielski, Some remarks and problems on the colouring of infinite graphs and the theorem of Kuratowski, Acta Math. Acad. Sci. Hungar. 12 (1961), 125–129.
[9] T. Łuczak, J. Nešetřil, A probabilistic approach to the dichotomy problem, SIAM J. Comput. 36(3) (2006), 835–843.
[10] M. Siggers, A strong Malcev condition for locally finite varieties omitting the unary type, Algebra Universalis 64 (2010), 15–20.