

Clones on regular cardinals

by

Martin Goldstern (Wien) and Saharon Shelah (Jerusalem)

Abstract. We investigate the structure of the lattice of clones on an infinite set X . We first observe that ultrafilters naturally induce clones; this yields a simple proof of Rosenberg's theorem: there are 2^{2^λ} maximal (= "precomplete") clones on a set of size λ . The clones we construct do not contain all unary functions. We then investigate clones that do contain all unary functions. Using a strong negative partition theorem from pcf theory we show that for cardinals λ (in particular, for all successors of regulars) there are 2^{2^λ} such clones on a set of size λ . Finally, we show that on a weakly compact cardinal there are exactly 2 precomplete clones which contain all unary functions.

1. Introduction

1.1. DEFINITION. Let X be a nonempty set. The *full clone* on X , denoted by \mathcal{O} or $\mathcal{O}(X)$, is the set of all finitary functions on X : $\mathcal{O} = \bigcup_{n=1}^{\infty} \mathcal{O}^{(n)}$, where $\mathcal{O}^{(n)}$ is the set of all functions from X^n into X .

A *clone* (on X) is a set $\mathcal{C} \subseteq \mathcal{O}$ which contains all projections and is closed under composition. That is,

(1) For all $1 \leq k \leq n$, the function $\pi_k^n : X^n \rightarrow X$, $\pi_k^n(x_1, \dots, x_n) = x_k$, is in \mathcal{C} .

(2) Whenever $f_1, \dots, f_k \in \mathcal{C} \cap \mathcal{O}^{(n)}$, $g \in \mathcal{C} \cap \mathcal{O}^{(k)}$, then the function

$$(x_1, \dots, x_n) \mapsto g(f_1(x_1, \dots, x_n), \dots, f_k(x_1, \dots, x_n))$$

(which we sometimes call $g(f_1, \dots, f_k)$) is also in \mathcal{C} .

2000 *Mathematics Subject Classification*: Primary 08A40; Secondary 03B50, 03E05.

Key words and phrases: precomplete clones; maximal clones; weakly compact cardinal; negative square bracket partition relation.

The first author is grateful to the Hebrew University of Jerusalem for the hospitality during his visit, and to the Austrian Science foundation for supporting the joint research under FWF grant P13325-MAT.

The second author is supported by the Israel Science Foundation founded by the Israel Academy of Sciences and Humanities.

Alternatively, \mathcal{C} is a clone if \mathcal{C} is the set of term functions of some universal algebra over X .

The set of clones over X forms a complete algebraic lattice with largest element \mathcal{O} . The coatoms of this lattice are called *precomplete clones* or *maximal clones*.

Many results for clones on finite sets, and in particular a classification of all precomplete clones on finite sets can be found in [Sz86].

Rosenberg proved in [Ro76] that if X is an infinite set of cardinality λ then there are 2^{2^λ} precomplete clones on X . In Section 2 we will give a short new proof of this theorem, using ultrafilters.

Let $\mathcal{O}^{(1)}$, the *full unary clone*, be the clone generated by $\mathcal{O}^{(1)}$, i.e., the set of functions which depend only on one argument:

$$\mathcal{O}^{(1)} := \{f \circ \pi_k^n : f \in \mathcal{O}^{(1)}, 1 \leq k \leq n\}.$$

The clones that we construct in Section 2, as well as the clones in the family constructed by Rosenberg, all have the property that they induce a proper submonoid of the monoid $\mathcal{O}^{(1)}$ of all unary functions. This raises the following question: What is the structure of those clones that contain the full monoid of all unary functions, i.e., the interval $[\mathcal{O}^{(1)}, \mathcal{O}]$? In particular, what can we say about the precomplete elements in this interval?

If X is a *finite* set with k elements, then it is known that this interval is actually a finite chain (with $k + 1$ elements). In particular, there is a unique precomplete clone above the full unary clone, namely, the set of all functions which are either essentially unary or not onto.

We now turn to *infinite* sets. Again we will be mainly interested in the maximal or “precomplete” clones above $\mathcal{O}^{(1)}$. Since \mathcal{O} is finitely generated over \mathcal{O}_1 , it is clear that the interval $[\mathcal{O}^{(1)}, \mathcal{O}]$ is dually atomic, that is, every $\mathcal{C} \in [\mathcal{O}^{(1)}, \mathcal{O}]$ is contained in some precomplete $\mathcal{C}' \in [\mathcal{O}^{(1)}, \mathcal{O}]$. (See Fact 1.3.)

For the case of countable X , Gavrilov proved in [Ga65] that there are exactly 2 precomplete clones in this interval, and Davies and Rosenberg (see [DR85]) gave an explicit example of one precomplete clone in this interval for every infinite X .

It turns out that (for any infinite set X of regular cardinality), the clones on X above $\mathcal{O}^{(1)}$ can be naturally divided into 2 classes, depending on whether the binary functions of the clone are all “almost unary” or if there is a “heavily binary” function among them (see Definitions 3.1 and 5.1).

In Section 3 we show that among the clones whose binary part is almost unary, there is a unique precomplete clone (namely, the clone from [DR85], which we call \mathcal{U} , see 3.3).

Finally, we discuss the case which was hitherto unknown, and which turns out to be the most interesting from the set-theoretical point of view: clones with heavily binary functions. The structure of the set of these clones depends on partition properties of the cardinality of the underlying set:

(1) If the cardinality of the underlying set is a weakly compact cardinal (or \aleph_0), then there is a unique precomplete clone in $[\mathcal{O}^{(1)}, \mathcal{O}]$ which is heavily binary (so altogether there are exactly two precomplete clones above $\mathcal{O}^{(1)}$: $\hat{\mathcal{U}}$ from 3.3 and \mathcal{T} from 5.11). This result, which generalizes Gavrilov's theorem for \aleph_0 , is proved in Section 5.

(2) If the cardinality λ of the underlying set satisfies a certain negative partition property $\text{Pr}(\lambda)$ (see also 6.4—in particular, we know $\text{Pr}(\kappa^+)$ for all regular uncountable κ), then there are 2^{2^λ} precomplete clones above $\mathcal{O}^{(1)}$ which are heavily binary. This result is proved in Section 4.

In an appendix we briefly discuss partition relations and the combinatorial principle $\text{Pr}(\lambda)$.

All sections of the paper can be read independently, but they all rely on notation, facts and concepts established in this introduction.

We plan to investigate clones on singular cardinals in a separate paper.

1.2. NOTATION. We fix an infinite set X . For $n \in \{1, 2, \dots\}$ we write $\mathcal{O}^{(n)}$ for the set of all functions from X^n to X , and $\mathcal{O} = \bigcup_{n=1}^{\infty} \mathcal{O}^{(n)}$.

For any set of functions $\mathcal{F} \subseteq \mathcal{O}$ we let $\text{cl}(\mathcal{F})$ be the smallest clone containing \mathcal{F} as well as all unary functions.

We will write $\lambda = |X|$ for the cardinality of X . It will often be convenient to have a well-order of X available; we will then identify X with the ordinal λ .

We call a function $f : X \times X \rightarrow X$ a *pairing function* if $f \upharpoonright \{(x, y) : x \neq y\}$ is 1-1. For the rest of the paper we fix a pairing function pr . We will assume that the cardinality of the complement of the range of pr is equal to the cardinality of X : $|X \setminus \text{ran}(\text{pr})| = \lambda$.

We fix a value $0 \in X$, and we will assume that 0 is not in the range of pr .

When we consider terms in which several functions are nested, we may write fx or gxy for $f(x)$ or $g(x, y)$ to avoid too many parentheses.

We identify X^n with the set of functions from $\{1, \dots, n\}$ to X . If $s \cap t = \emptyset$, $s \cup t = \{1, \dots, n\}$, $a : s \rightarrow X$, $b : t \rightarrow X$, then $a \cup b$ is in X^n .

If $\mathcal{C} \subseteq \mathcal{O}$ is a clone, we let $\mathcal{C}^{(n)} = \mathcal{C} \cap \mathcal{O}^{(n)}$.

1.3. FACT. (1) *If $f : X \times X \rightarrow X$ is a pairing function, then there are unary functions g, g_1, g_2 such that the function $(x, y) \mapsto g \circ f(g_1x, g_2y)$ is a bijection from $X \times X$ to X .*

(2) *If $\mathcal{C} \subseteq \mathcal{O}$, $\{\text{pr}\} \cup \mathcal{O}^{(1)} \subseteq \mathcal{C}$, where pr is any pairing function, then $\mathcal{C} = \mathcal{O}$. [Use (1).]*

- (3) If $\mathcal{O}^{(1)} \subseteq \mathcal{C} \subseteq \mathcal{O}$, $\text{pr} \notin \mathcal{C}$, then the clones which are maximal in $\{\mathcal{D} : \mathcal{C} \subseteq \mathcal{D} \subseteq \mathcal{O}, \text{pr} \notin \mathcal{D}\}$

are exactly the precomplete clones extending \mathcal{C} . (By Zorn's lemma this easily implies that the interval $[\mathcal{O}^{(1)}, \mathcal{O}]$ is dually atomic: Every clone above $\mathcal{O}^{(1)}$, except for \mathcal{O} itself, is contained in a precomplete one.)

1.4. REMARK. As we shall see in Section 3, we cannot relax the assumption "pr is 1-1 on $\{(x, y) : x \neq y\}$ " in 1.3(2) to "pr is 1-1 on $\{(x, y) : x < y\}$."

1.5. DEFINITION. Let J be any index set, and $R \subseteq X^J$. Let $f \in \mathcal{O}^{(n)}$. We say that f respects R if whenever $\bar{\rho}^1 = \langle \rho_i^1 : i \in J \rangle, \dots, \bar{\rho}^n = \langle \rho_i^n : i \in J \rangle$ are all in R , then also $\langle f(\rho_i^1, \dots, \rho_i^n) : i \in J \rangle \in R$.

We let $\text{Pol } R$ be the set of all functions respecting R .

We will usually be interested in the case where R is a set of n -ary functions on X , i.e., $R \subseteq X^{X^n}$.

The following observations follow easily from the definitions and from the facts above:

- 1.6. FACT. (1) For any relation R , $\text{Pol } R \subseteq \mathcal{O}$ is a clone.
 (2) If \mathcal{C} is a clone, then $\mathcal{C} \subseteq \text{Pol } \mathcal{C}^{(n)}$.
 (3) If \mathcal{C} is a clone and $\mathcal{C}^{(n)} \neq \mathcal{O}^{(n)}$, then $\text{Pol } \mathcal{C}^{(n)} \subsetneq \mathcal{O}$. In fact, $(\text{Pol } \mathcal{C}^{(n)})^{(n)} = \mathcal{C}^{(n)}$.
 (4) If \mathcal{C} is a precomplete clone and $\mathcal{C}^{(1)} \neq \mathcal{O}^{(1)}$, then $\mathcal{C} = \text{Pol } \mathcal{C}^{(1)}$.
 (5) If \mathcal{C} is a precomplete clone and $\mathcal{C}^{(1)} = \mathcal{O}^{(1)}$, then $\mathcal{C}^{(2)} \neq \mathcal{O}^{(2)}$, and $\mathcal{C} = \text{Pol } \mathcal{C}^{(2)}$.

1.7. MORE NOTATION. Let \mathcal{C} be a clone on the set X . We let $\tilde{\mathcal{C}}$ be the set of all functions $\bar{f} : X^n \rightarrow X^k$ ($n, k > 0$) such that each function $\pi_i^k \circ \bar{f}$ is in \mathcal{C} .

The "closure under composition" of the clone \mathcal{C} just means that $\tilde{\mathcal{C}}$ is closed under the usual notion of composition, i.e., whenever $\bar{f} : X^n \rightarrow X^m$ and $\bar{g} : X^m \rightarrow X^k$ are in $\tilde{\mathcal{C}}$ then also $\bar{g} \circ \bar{f} \in \tilde{\mathcal{C}}$.

1.8. Acknowledgments. We are grateful to Lutz Heindorf for his thoughtful remarks on an earlier version of the paper, and for alerting us to Gavrilov's results.

2. A new proof of Rosenberg's theorem. Let X be an infinite set. Rosenberg [Ro76] has shown that there are $2^{2^{|X|}}$ precomplete clones on X . Using transfinite induction he first constructs $2^{2^{|X|}}$ clones with certain orthogonality properties and then shows that they can be extended to pairwise different precomplete clones.

Here we give an alternative proof of Rosenberg's theorem, utilizing the well-known fact (see e.g. [CN2]) that on every infinite set X there are $2^{2^{|X|}}$

ultrafilters. We will find an explicit 1-1 map from the ultrafilters to precomplete clones.

2.1. DEFINITION. Let $I \subseteq \mathcal{P}(X)$ be a maximal ideal. We define

$$\mathcal{C}_I := \bigcup_{n=1}^{\infty} \{f \in \mathcal{O}^{(n)} : \forall A \in I \ f[A^n] \in I\}.$$

2.2. FACT. (1) $\mathcal{C}_I \subsetneq \mathcal{O}$.

(2) \mathcal{C}_I is a clone.

(3) If $f : X^k \rightarrow X$ and $\text{ran}(f) \in I$, then $f \in \mathcal{C}_I$. More generally, if $A \in I$, $f : X^k \rightarrow X^n$ and the range of f is contained in A^n , then $f \in \mathcal{C}_I$.

(4) I can be reconstructed from \mathcal{C}_I as

$$I = \{A \subseteq X : \text{For all } f : X \rightarrow X : \text{If } \text{ran}(f) \subseteq A, \text{ then } f \in \mathcal{C}_I\},$$

so in particular the map $I \mapsto \mathcal{C}_I$ is 1-1.

(5) \mathcal{C}_I is a precomplete clone, i.e., for all $f \in \mathcal{O} \setminus \mathcal{C}_I$ the clone generated by $\mathcal{C}_I \cup \{f\}$ contains all of \mathcal{O} .

Proof. Parts (1)–(3) are clear. We only check (4) and (5).

For (4), let

$$I' = \{A \subseteq X : \text{For all } f : X \rightarrow X : \text{If } \text{ran}(f) \subseteq A, \text{ then } f \in \mathcal{C}_I\}.$$

By (3) above, $I \subseteq I'$, so we check $I' \subseteq I$. Let $A \notin I$. We want to prove $A \notin I'$, i.e., we are looking for a function $f : X \rightarrow X$ such that

$$\text{ran}(f) \subseteq A, \quad \text{but} \quad f \notin \mathcal{C}_I.$$


If $|A| \leq |X \setminus A|$, then let $A_0 := A$, otherwise we must have $|A| = |X|$, so we can write A as a disjoint union $A = A_0 \cup A_1$ with $|A_0| = |X| = |A_1|$, $A_0 \notin I$. In either case we have found a set $A_0 \subseteq A$ with $|A_0| \leq |X \setminus A_0|$ and $A_0 \notin I$. So there is a function $f : X \rightarrow X$ with $f[X] = f[X \setminus A_0] = A_0$. So $f \notin \mathcal{C}_I$ while $\text{ran}(f) \subseteq A_0 \subseteq A$. Thus, the function f witnesses that $A \notin I'$.

We now turn to the proof of (5). We may assume that I is nonprincipal, i.e., all finite sets are in I . [Otherwise $I = \{A \subseteq X : a_0 \notin A\}$ for some $a_0 \in X$, and in this case \mathcal{C}_I is just $\text{Pol } R$ (see 1.5 with a singleton index set J), where $R = X \setminus \{a_0\}$. It is well known (and easy to see) that the clone $\text{Pol } R$ is precomplete, for all unary relations $R \neq \emptyset, X$.]

Call a function f *conservative* if it satisfies $f(a_1, \dots, a_n) \in \{a_1, \dots, a_n\}$ for all $a_1, \dots, a_n \in X$. Clearly all conservative functions are in \mathcal{C}_I .

Let $f : X^k \rightarrow X$, $f \notin \mathcal{C}_I$. So there is some set $A \in I$ with $f[A^k] \notin I$. Let $B_0 = f[A^k]$ and $B_1 = X \setminus B_0$. So $B_0 \notin I$ and $B_1 \in I$.

Now let $g : X^n \rightarrow X$ be arbitrary. We have to show that g is in the clone generated by \mathcal{C}_I and f . Pick two distinct elements $0, 1$, where $\{0, 1\} \in I$.



The function

$$H(x, y, z) = \begin{cases} y & \text{if } x = 0, \\ z & \text{if } x \neq 0, \end{cases}$$

is conservative, hence in \mathcal{C}_I .

Let $C_0 = g^{-1}[B_0]$, $C_1 = g^{-1}[B_1]$, and define two “approximations” g_0, g_1 to g as follows:

$$g_0(\bar{x}) = \begin{cases} g(\bar{x}) & \text{if } \bar{x} \in C_0, \\ 0 & \text{if } \bar{x} \in C_1, \end{cases} \quad g_1(\bar{x}) = \begin{cases} 0 & \text{if } \bar{x} \in C_0, \\ g(\bar{x}) & \text{if } \bar{x} \in C_1, \end{cases}$$

Note that the range of g_1 is contained in $B_1 \cup \{0\}$, hence $g_1 \in I$.

Let $\chi(\bar{x}) = 0$ if $\bar{x} \in C_0$, and $\chi(\bar{x}) = 1$ if $\bar{x} \in C_1$. By definition of H , $g(\bar{x}) = H(\chi\bar{x}, g_0\bar{x}, g_1\bar{x})$, so all we have to show is that H, χ, g_0, g_1 are all in the clone generated by \mathcal{C}_I and f . We already know that

- $H \in \mathcal{C}_I$ (because H is conservative),
- $g_1 \in \mathcal{C}_I$ (because the range of g_1 is in I),
- $\chi \in \mathcal{C}_I$ (because the range of χ is in I).

It remains to show g_0 is in the clone generated by \mathcal{C}_I and f .

Let $f^* : B_0 \rightarrow A^k$ be an “inverse” of f , i.e.,

$$\forall b \in B_0 : f(f^*(b)) = b.$$

($f^* \upharpoonright B_1$ can be an arbitrary function with range $\subseteq A^k$.) Define $g'_0 : X^n \rightarrow X^k$ by $g'_0(c) = f^*(g_0(c))$. Note that the range of g'_0 is $\subseteq A^k$, $A \in I$, so $g'_0 \in \tilde{\mathcal{C}}_I$.

Now we have, for all $\bar{c} \in X^n$, $g_0(\bar{c}) = f(f^*(g_0(\bar{c}))) = f(g'_0(\bar{c}))$, so g_0 is in the clone generated by \mathcal{C}_I and f . ■

2.3. CONCLUSION. *On any infinite set X there are exactly $2^{2^{|X|}}$ precomplete clones.*

Proof. The upper bound follows from $|\mathcal{C}| = 2^{|X|}$. For the lower bound: it is known that there are $2^{2^{|X|}}$ maximal ideals, and we have just shown that the function $I \mapsto \mathcal{C}_I$ maps them injectively to precomplete clones. ■

3. Almost unary clones. In this section we will consider clones on an infinite set X of regular cardinality. We will call a set *small* if its cardinality is smaller than the cardinality of X , and we will say that there are *few* objects with some property if the set of those objects is small.

For example, if X is countable, then “small” will mean “finite”. If X has cardinality \aleph_1 , then “small” will mean “finite or countably infinite”.

[With this terminology, the property “ X has regular cardinality” can be rephrased as “ X cannot be written as a union of few small sets”.]

3.1. DEFINITION. Let $g : X^n \rightarrow X$. We say that g is *almost unary* if there is a function G which is defined on X , each $G(x)$ a small subset of X , such that for some k ,

$$\forall (x_1, \dots, x_n) \in X^n : g(x_1, \dots, x_n) \in G(x_k).$$

If X itself is a cardinal, then we can equivalently say: g is almost unary if for some k and $G : X \rightarrow X$, for all $x_1, \dots, x_n \in X$: $g(x_1, \dots, x_n) \leq G(x_k)$.

3.2. DEFINITION. Let $\mathcal{U} \subseteq \mathcal{O}$ be the set of all almost unary functions.

In Definition 5.1 we will call functions in $\mathcal{O}^{(2)} \setminus \mathcal{U}$ *heavily binary*.

The set $\mathcal{U}^{(2)}$ is called T_1 in [Ga65].

3.3. DEFINITION. Let $\hat{\mathcal{U}} := \text{Pol } \mathcal{U}^{(2)}$ (see 1.5).

That is, a function $f \in \mathcal{O}^{(n)}$ is in $\hat{\mathcal{U}}$ iff

$$\forall g_1, \dots, g_n \in \mathcal{U}^{(2)} : f(g_1, \dots, g_n) \in \mathcal{U}^{(2)}$$

where $f(g_1, \dots, g_n)$ is the function $(x, y) \mapsto f(g_1(x, y), \dots, g_n(x, y))$.

Note that $\hat{\mathcal{U}} \cap \mathcal{O}^{(2)} = \mathcal{U} \cap \mathcal{O}^{(2)}$, and $\mathcal{U} \subseteq \hat{\mathcal{U}}$.

3.4. EXAMPLE. Let $X = \lambda$ be a cardinal, so the small subsets of X are exactly the bounded subsets of λ .

(1) The function \min is almost unary: $\min \in \mathcal{U}^{(2)}$.

(2) The function \max is not almost unary.

(3) The median function med , defined by

$$\text{med}(x, y, z) = \max(\min(x, y), \min(y, z), \min(x, z)),$$

is not almost unary, but it is easy to check that med respects all binary almost unary functions, so $\text{med} \in \hat{\mathcal{U}} \setminus \mathcal{U}$.

(4) Let pr_Δ be defined by

$$\text{pr}_\Delta(x, y) = \begin{cases} \text{pr}(x, y) & \text{if } x > y, \\ 0 & \text{otherwise,} \end{cases}$$

where pr is a pairing function. Then $\text{pr}_\Delta \in \mathcal{U}$.

The following was already observed by Davies and Rosenberg [DR85]. In the countable case, Gavrilov [Ga65, Theorem 4] showed that $\text{Pol } \mathcal{U}^{(2)}$ is precomplete.

3.5. CONCLUSION. Assume $\mathcal{C} \in [\mathcal{O}^{(1)}, \mathcal{O}]$. If $\text{pr}_\Delta \in \mathcal{C}$ (see 3.4), and if \mathcal{C} contains a binary function not in $\mathcal{U}^{(2)}$, then $\mathcal{C} = \mathcal{O}$. Hence, $\text{Pol } \mathcal{U}^{(2)}$ is an example of a precomplete clone containing all unary functions.

Proof. Let $p_1, p_2 : X \rightarrow X$ be two 1-1 functions such that the ranges of p_1, p_2, pr are disjoint. Since \mathcal{C} contains a function which is not almost unary, there is some $H \in \mathcal{C}^{(2)}$ with $H(x, p_2 0) = x = H(p_1 0, x)$ for all x in the range of pr . Then the function

$$(x, y) \mapsto H(p_1(\text{pr}_\Delta(x, y)), p_2(\text{pr}_\Delta(y, x)))$$

is a pairing function. ■

(We will meet a similar argument again in the proof of 5.10.)

We now show a kind of converse to this theorem: $\text{Pol } \mathcal{U}^{(2)}$ is the *unique* precomplete clone which contains all unary functions and only “almost unary” binary functions.

3.6. THEOREM. Assume that $\mathcal{C} \subseteq \mathcal{O}$ is a precomplete clone, $\mathcal{O}^{(1)} \subseteq \mathcal{C}$, and $\mathcal{C}^{(2)} \subseteq \mathcal{U}^{(2)}$. Then $\mathcal{C} = \text{Pol } \mathcal{U}^{(2)}$.

We will prove this theorem below. We start by investigating which coordinates are responsible for a function having a large range.

3.7. DEFINITION. Let $g \in \mathcal{O}^{(n)}$. We define a set S_g of subsets of $\{1, \dots, n\}$ as follows:

$$S_g = \{s \subseteq \{1, \dots, n\} : \exists \bar{a} \in X^{\{1, \dots, n\} \setminus s} : |\{g(\bar{a} \cup \bar{x}) : \bar{x} \in X^s\}| = |X|\}$$

(Here we write X^s for the set of all functions from s to X .)

3.8. LEMMA. Assume $\text{cl}(g)^{(2)} \subseteq \mathcal{U}$. Then

$$\forall r, t \in S_g : r \cap t \neq \emptyset.$$

Proof. Choose r and t in S_g with $r \cap t = \emptyset$. Using unary functions, we will construct a binary function in $\text{cl}(g)$ which is not in $\mathcal{U}^{(2)}$.

Let $s := \{1, \dots, n\} \setminus (r \cup t)$, so $\{1, \dots, n\} = r \dot{\cup} s \dot{\cup} t$. So there is some $\bar{a} \in X^{s \cup t}$ and a sequence $(\bar{x}^\alpha : \alpha \in X)$ of elements of X^r such that all values $g(\bar{a} \cup \bar{x}^\alpha)$ are different. Similarly, there is some $\bar{b} \in X^{r \cup s}$ and a sequence $(\bar{y}^\beta : \beta \in X)$ of elements of X^t such that all values $g(\bar{b} \cup \bar{y}^\beta)$ are different.

Now for $l = 1, \dots, n$ define functions h_l as follows: Fix some element $0 \in X$ and set

$$h_l(\alpha, \beta) = \begin{cases} \bar{x}_\alpha(l) & \text{if } l \in r, \alpha \neq 0, \\ \bar{b}(l) & \text{if } l \in r, \alpha = 0, \\ \bar{a}(l) & \text{if } l \in s, \alpha \neq 0, \\ \bar{b}(l) & \text{if } l \in s, \alpha = 0, \\ \bar{y}_\beta(l) & \text{if } l \in t, \beta \neq 0, \\ \bar{a}(l) & \text{if } l \in t, \beta = 0. \end{cases}$$

Formally, the functions h_l are in $\mathcal{O}^{(2)}$, but each of them is essentially unary: $h_l(\alpha, \beta)$ depends only on α for $l \in r \cup s$, and only on β for $l \in t$. This implies that $h_l \in \mathcal{O}^{(1)}$.

Now $F = g(h_1, \dots, h_n)$, i.e., $F(\alpha, \beta) = g(h_1(\alpha, \beta), \dots, h_n(\alpha, \beta))$, will be a binary function in $\text{cl}(g)$ but not in $\mathcal{U}^{(2)}$, since the values $F(\alpha, 0) = g(\bar{x}_\alpha \cup \bar{a})$ are all different, as are the values $F(0, \beta) = g(\bar{y}_\beta \cup \bar{b})$. ■

The previous lemma will allow us to relate any “almost unary” clone to $\text{Pol } \mathcal{U}^{(2)}$:

3.9. LEMMA. *Assume $\mathcal{O}^{(1)} \subseteq \mathcal{C}$ and $\mathcal{C}^{(2)} \subseteq \mathcal{U}^{(2)}$. Then $\mathcal{C} \subseteq \text{Pol } \mathcal{U}^{(2)}$. That is: whenever $d_1, \dots, d_n \in \mathcal{U}^{(2)}$ and $g \in \mathcal{C}^{(n)}$, then $f := g(d_1, \dots, d_n) \in \mathcal{U}^{(2)}$.*

Proof. Since each $d_l \in \mathcal{U}^{(2)}$, we can find a decomposition $\{1, \dots, n\} = r \cup t$, $r \cap t = \emptyset$, and a function D mapping each $\alpha \in X$ to a small subset $D(\alpha) \subseteq X$ such that:

- For all $l \in r$, all $\alpha, \beta \in X$: $d_l(\alpha, \beta) \in D(\alpha)$.
- For all $l \in t$, all $\alpha, \beta \in X$: $d_l(\alpha, \beta) \in D(\beta)$.

By the previous lemma, we cannot have both r and t in S_g , so assume $t \notin S_g$.

Now fix any element $0 \in X$. We will show that the set $\{f(0, \beta) : \beta \in X\}$ is small.

Let us consider $f(0, \beta) = g(d_1(0, \beta), \dots, d_n(0, \beta))$. Identifying X^n with $X^{\{1, \dots, n\}}$, we can write the tuple $(d_1(0, \beta), \dots, d_n(0, \beta))$ as $a_\beta \cup y_\beta$, where $a_\beta \in X^r$ and $y_\beta \in X^t$. Now note that for $l \in r$ we have $d_l(0, \beta) \in D(0)$, so $a_\beta \in D(0)^r$, which is a small set.

Hence

$$\{f(0, \beta) : \beta \in X\} \subseteq \{g(a \cup y) : a \in D(0)^s, y \in X^t\}.$$

For each fixed $a \in D(0)^s$ the set $\{g(a \cup y) : y \in X^t\}$ is small (since $t \notin S_g$), so, since $D(0)$ is small, also

$$\{g(a \cup y) : a \in D(0)^s, y \in X^t\} = \bigcup_{a \in D(0)^r} \{g(a \cup y) : y \in X^t\}$$

is small. ■

Proof of Theorem 3.6. Assume that $\mathcal{O}^{(1)} \subseteq \mathcal{C}$, $\mathcal{C}^{(2)} \subseteq \mathcal{U}^{(2)}$, and \mathcal{C} is precomplete. Then by Lemma 3.9, we have $\mathcal{C} \subseteq \text{Pol } \mathcal{U}^{(2)}$. But since \mathcal{C} is precomplete, we must have $\mathcal{C} = \text{Pol } \mathcal{U}^{(2)}$. ■

4. Successors of regulars. We fix a set X of regular cardinality λ , and for simplicity we write $X = \lambda$. We fix a pairing function $\text{pr} : \lambda \times \lambda \rightarrow \lambda$ as in 1.2.

We will use the following combinatorial principle $\text{Pr}(\lambda, \mu)$:

- There is a symmetric function $c : \lambda \times \lambda \rightarrow \mu$ with the following anti-Ramsey property: For all sequences $(a_i : i < \lambda)$ of pairwise disjoint finite subsets of λ , and for all $c_0 \in \mu$, there are $i < j < \lambda$ such that $c \upharpoonright (a_i \times a_j)$ is constant with value c_0 .

(See Section 6 for background.)

We fix a function c witnessing the above statement.

4.1. DEFINITION. For any $A \subseteq \mu$ we define a function $F_A : \lambda \times \lambda \rightarrow \lambda$ as follows:

$$F_A(\alpha, \beta) = \begin{cases} \max(\alpha, \beta) & \text{if } \alpha = 0 \text{ or } \beta = 0 \text{ or } \alpha = \beta, \\ \text{pr}(\alpha, \beta) & \text{if } c(\alpha, \beta) \in A, \\ 0 & \text{otherwise.} \end{cases}$$

4.2. FACT. If $A \cup B = \mu$, then $\text{cl}(F_A, F_B) = \mathcal{O}$.

Proof. We will show how to construct a pairing function from F_A and F_B . Define

$$\text{pr}'(\alpha, \beta) = F_A(F_A(\alpha, \beta), F_B(\alpha, \beta)).$$

We claim that $\text{pr}'(\alpha, \beta) = \text{pr}(\alpha, \beta)$ for all distinct $\alpha, \beta > 0$. Indeed, if $c(\alpha, \beta) \in A \cap B$, then

$$\text{pr}'(\alpha, \beta) = F_A(\text{pr}(\alpha, \beta), \text{pr}(\alpha, \beta)) = \text{pr}(\alpha, \beta),$$

if $c(\alpha, \beta) \in A \setminus B$, then

$$\text{pr}'(\alpha, \beta) = F_A(\text{pr}(\alpha, \beta), 0) = \text{pr}(\alpha, \beta),$$

and if $c(\alpha, \beta) \in B \setminus A$, then

$$\text{pr}'(\alpha, \beta) = F_A(0, \text{pr}(\alpha, \beta)) = \text{pr}(\alpha, \beta).$$

Hence pr' is a pairing function. ■

4.3. MAIN LEMMA. Assume that $A \not\subseteq B_1 \cup \dots \cup B_k$. Then

$$F_A \notin \text{cl}(F_{B_1}, \dots, F_{B_k}).$$

We will prove this lemma below, but first we will show how it can be used.

4.4. DEFINITION. We say that $\mathcal{A} = (A_i : i \in I)$ is an *independent* family of subsets of X if every nontrivial Boolean combination of sets from \mathcal{A} is nonempty, i.e., whenever J_0 and J_1 are finite disjoint subsets of I , then

$$\bigcap_{i \in J_0} A_i \cap \bigcap_{i \in J_1} (X \setminus A_i) \neq \emptyset.$$

The following theorem of Hausdorff is well known:

4.5. THEOREM. *If $|X| = \mu$, then there is an independent family $\mathcal{A} = (A_i : i \in I)$ of subsets of X with $|I| = 2^\mu$.*

Proof. See [Ku83, Chapter VIII, Exercise A6] or [Ko89, Example 9.21]. ■

4.6. THEOREM. *Assume $\text{Pr}(\lambda, \mu)$. Then there are at least 2^{2^μ} precomplete clones above the unary functions on the set λ . (Hence: If $\lambda = \kappa^+$, κ regular uncountable, then there are 2^{2^λ} precomplete clones above $\mathcal{O}^{(1)}$.)*

Proof. Let $(A_i : i \in 2^\mu)$ be an independent family of subsets of μ . Write $-A_i$ for $\mu \setminus A_i$. For each $J \subseteq 2^\mu$ we let

$$\mathcal{C}_J = \text{the clone generated by } \{F_{A_i} : i \in J\} \cup \{F_{-A_i} : i \notin J\} \cup \mathcal{O}^{(1)}.$$

We will now show that

- (1) $\mathcal{C}_J \neq \mathcal{O}$ for all $J \subseteq 2^\mu$.
- (2) Whenever $J_1 \neq J_2$, then $\mathcal{C}_{J_1} \cup \mathcal{C}_{J_2}$ already generates \mathcal{O} .

This will conclude the proof, because (1) together with Fact 4.2 implies that each \mathcal{C}_J can be extended to a precomplete clone, and (2) implies that no single precomplete clone can contain $\mathcal{C}_{J_1} \cup \mathcal{C}_{J_2}$ for distinct J_1, J_2 .

Proof of (1): We can assume there is some $i \notin J$. By independence, A_i cannot be covered by any finite union from $\{A_j : j \in J\} \cup \{-A_j : j \notin J\}$. So by the main lemma, F_{A_i} is not in the clone \mathcal{C}_J .

Proof of (2): If $J_1 \neq J_2$, then there is some $i \in J_1 \setminus J_2$, say. Now $F_{A_i} \in \mathcal{C}_{J_1}$, $F_{-A_i} \in \mathcal{C}_{J_2}$, and by Fact 4.2, $\{F_A, F_{-A}\}$ generates \mathcal{O} . ■

We now prepare for the proof of the main lemma 4.3. Our situation is the following: We have a function c witnessing $\text{Pr}(\lambda, \mu)$. Using c and our fixed pairing function pr we have defined functions $F_A : \lambda \times \lambda \rightarrow \lambda$ for every $A \subseteq \mu$ in 4.1. We are given sets $A, B_1, \dots, B_k \subseteq \mu$, $A \not\subseteq B_1 \cup \dots \cup B_k$. Pick $c_0 \in A \setminus (B_1 \cup \dots \cup B_k)$.

We want to show that $F_A \notin \text{cl}(F_{B_1}, \dots, F_{B_k})$, i.e., F_{B_1}, \dots, F_{B_k} , together with all unary functions, do not generate F_A .

4.7. DEFINITION. *Terms over λ* are defined inductively as follows:

- (1) The formal variables \mathbf{x}, \mathbf{y} are terms, as well as every element of λ .
- (2) If σ is a term and $f : \lambda \rightarrow \lambda$ a unary function, then (f, σ) is a term.
- (3) If σ_1 and σ_2 are terms and $1 \leq i \leq k$, then $(F_{B_i}, \sigma_1, \sigma_2)$ is a term.

Every term τ induces (in the obvious way) a function $\tau : \lambda \times \lambda \rightarrow \lambda$ in $\text{cl}(\{F_{B_1}, \dots, F_{B_k}\})$. Conversely, every binary function in $\text{cl}(\{F_{B_1}, \dots, F_{B_k}\})$ is represented by a term.

We call a term *constant* if it is an element of λ , and *\mathbf{x} -unary* if \mathbf{y} does not appear in it; similarly for *\mathbf{y} -unary*. A term is *unary* if it is \mathbf{x} -unary or \mathbf{y} -unary. (By definition, constant terms are both \mathbf{x} -unary and \mathbf{y} -unary.)

For the following discussion, fix a term τ_0 . Our aim is to find a large set on which all subterms of τ_0 behave like unary functions. We will first explain how to find (terms for) these unary functions, and then we show they are indeed realized on some large set.

4.8. DEFINITION. Let $S \subseteq \lambda$. For any term τ we will try to define a unary term τ^S . Whenever σ^S is undefined for a subterm σ of τ , then also τ^S will be undefined. Our definition proceeds by induction on the structure of τ . “ B ” will stand for any of the sets B_1, \dots, B_n .

(1) $\tau = \mathbf{x}$ or $\tau = \mathbf{y}$ or $\tau = c \in \lambda$. In this case, $\tau^S = \tau$.

(2) $\tau = (f, \sigma)$, and $\sigma^S = c \in \lambda$. In this case, τ^S is also a constant, namely $f(c)$.

(3) $\tau = (f, \sigma)$, $\sigma^S = (g, \mathbf{x})$. If $f \circ g$ is 1-1 on S , then $\tau^S := (f \circ g, \mathbf{x})$. If $f \circ g$ is constant with value d on S , then $\tau^S := d$. If $f \circ g$ is neither 1-1 nor constant, then τ^S will be undefined.

(4) $\tau = (F_B, \sigma_1, \sigma_2)$, and σ_1^S and σ_2^S are constant (say, with values c_1 and c_2). In this case we let $\tau^S := F_B(c_1, c_2)$.

(5) $\tau = (F_B, \sigma_1, \sigma_2)$, and $\sigma_1^S = (f, \mathbf{x})$, $\sigma_2^S = d$ (a constant). If the function $h : x \mapsto F_B(f(x), d)$ is 1-1 or constant (say, with value c) on S , then we let τ^S be (h, \mathbf{x}) or c , respectively. (If h is neither constant nor 1-1 on S , then τ^S is again undefined.)

(6) $\tau = (F_B, \sigma_1, \sigma_2)$, and $\sigma_1^S = (f_1, \mathbf{x})$, $\sigma_2^S = (f_2, \mathbf{x})$. If the function $h : x \mapsto F_B(f_1(x), f_2(x))$ is 1-1 or constant (say, with value d), then we let τ^S be (h, \mathbf{x}) or d , respectively. (Otherwise, τ^S is again undefined.)

(7) $\tau = (F_B, \sigma_1, \sigma_2)$, and $\sigma_1^S = (f_1, \mathbf{x})$, $\sigma_2^S = (f_2, \mathbf{y})$. We let $\tau^S := 0$. **This is the crucial case of our definition.**

(8) Repeat all the above items with \mathbf{x} and \mathbf{y} interchanged, and/or σ_1 and σ_2 interchanged.

4.9. FACT. *Whenever τ^S is defined, then τ^S is either constant, or of the form (f, \mathbf{x}) or (f, \mathbf{y}) , where f is 1-1 on S .*

4.10. FACT. (1) *If τ^S is defined and $S' \subseteq S$, then $\tau^{S'}$ is defined.*

(2) *Fix a finite set T of terms which is closed under subterms. Then for every set S of regular infinite cardinality there a set $S' \subseteq S$ of the same cardinality such that for all $\tau \in T$, $\tau^{S'}$ is well defined.*

Proof. Proceed by induction on the complexity of the terms. We have to thin out the set S finitely many times in order to make finitely many functions 1-1 or constant. ■

4.11. LEMMA. *Assume that τ^S is defined and $|S| = \lambda$. Then there are $\alpha < \beta$ in S such that $\tau(\alpha, \beta) = \tau^S(\alpha, \beta)$ and $c(\alpha, \beta) = c_0$.*

Proof. Let T be the set of subterms of τ (including τ itself). Collect all the 1-1 functions appearing in σ^S for $\sigma \in T$, i.e.,

$$\mathcal{F} := \{f : \exists \sigma \in T \sigma^S = (f, \mathbf{x}) \text{ or } \sigma^S = (f, \mathbf{y})\}.$$

The set \mathcal{F} is finite, the identity function is in \mathcal{F} , and all functions in \mathcal{F} are 1-1. We may thin out the set S so that the family

$$(\{f(\alpha) : f \in \mathcal{F}\} : \alpha \in S)$$

is pairwise disjoint. So since c witnesses $\text{Pr}(\lambda, \mu)$, we can find $\alpha < \beta$ such that for all $f, g \in \mathcal{F}$, $c(f(\alpha), g(\beta)) = c_0$ (and $f(\alpha) \neq g(\beta)$). This implies $F_{B_i}(f(\alpha), g(\beta)) = 0$.

Now we can prove by induction on the complexity of the subterms σ of τ that $\sigma^S(\alpha, \beta) = \sigma(\alpha, \beta)$. ■

Proof of Lemma 4.3. Let $c_0 \in A \setminus (B_1 \cup \dots \cup B_k)$, and let τ be a term. We will find α, β such that $\tau(\alpha, \beta) \neq F_A(\alpha, \beta)$.

We can find a set S such that τ^S is defined. Let \mathcal{F} be again the finite set of 1-1 functions used in defining τ^S . We can thin out the set S so that for all $f \in \mathcal{F}$,

$$\forall \alpha, \beta \in S : \alpha \neq \beta \Rightarrow f(\alpha) \neq \text{pr}(\alpha, \beta) \neq f(\beta).$$

[Why? For each $f \in \mathcal{F}$ define a partial function \bar{f} such that $\bar{f}(\alpha) = \beta$ whenever $f(\alpha) = \text{pr}(\alpha, \beta)$, $\alpha \neq \beta$. \bar{f} is well defined, since pr is a pairing function. We can thin out S to get $\bar{f}(\alpha) \notin S$ for all $\alpha \in S$. This is sufficient.]

Now thin out S so that for all $\alpha \in S$, $f(\alpha) \notin S$ or $f(\alpha) = \alpha$, and that none of the finitely many constants appearing as τ^S is equal to $\text{pr}(\alpha, \beta)$ for $\alpha, \beta \in S$.

By Lemma 4.11, we can find $\alpha < \beta$ with $\tau(\alpha, \beta) = \tau^S(\alpha, \beta)$ and $c(\alpha, \beta) = c_0$. Now we have $F_A(\alpha, \beta) = \text{pr}(\alpha, \beta)$ (as $c(\alpha, \beta) = c_0 \in A$). On the other hand, τ^S is either constant or of the form (f, \mathbf{x}) or (f, \mathbf{y}) for some $f \in \mathcal{F}$. So $\tau^S(\alpha, \beta) \neq F_A(\alpha, \beta)$.

This concludes the proof of Lemma 4.3 and hence also of Theorem 4.6. ■

5. Weakly compact cardinals. In this section we deal with clones on infinite sets whose cardinality λ satisfies $\lambda \rightarrow (\lambda)_2^2$ (so either $\lambda = \aleph_0$ or λ is weakly compact).

Our main result here is that (in addition to the precomplete clone from Section 3) there is only one more precomplete clone containing all unary functions. This result was proved for the case $\lambda = \aleph_0$ already by Gavrilov [Ga65].

The proof we give here is self-contained and does not rely on Gavrilov's paper; however, many of the notions we use (in particular the case distinction depending on the behaviour of $f \upharpoonright \Delta$, $f \upharpoonright \nabla$) have obvious parallels in [Ga65].

Recall that $\lambda \rightarrow (\lambda)_2^2$ implies $\lambda \rightarrow (\lambda)_k^n$ for all $n, k < \omega$, i.e., whenever $h : [\lambda]^n \rightarrow \{1, \dots, k\}$, then there is a subset $S \subseteq \lambda$ with $|S| = \lambda$ such that $h \upharpoonright [S]^n$ is constant.

5.1. DEFINITION. Let $H : \lambda^n \rightarrow \lambda$.

(1) We say that H depends on the k th coordinate if there is an $(n - 1)$ -tuple $(a_1, \dots, a_{k-1}, a_{k+1}, \dots, a_n)$ such that the set

$$\{H(a_1, \dots, a_{k-1}, x, \dots, a_n) : x \in \lambda\}$$

has more than one element. In this case we may also write H symbolically as $H(\mathbf{x}_1, \dots, \mathbf{x}_n)$ and say H depends on \mathbf{x}_k . For $n = 2$ we may also say $H(\mathbf{x}, \mathbf{y})$ depends on \mathbf{x} or on \mathbf{y} .

(2) We say that $H(\mathbf{x}_1, \dots, \mathbf{x}_n)$ depends heavily on the k th coordinate (or: on \mathbf{x}_k) if there is an $(n - 1)$ -tuple $(a_1, \dots, a_{k-1}, a_{k+1}, \dots, a_n)$ such that the set

$$\{H(a_1, \dots, a_{k-1}, x, \dots, a_n) : x \in \lambda\}$$

has λ elements.

(3) We say that $\mathcal{C} \subseteq \mathcal{O}$ is heavily binary if there exists $H(\mathbf{x}, \mathbf{y}) \in \mathcal{C}$ which depends heavily on \mathbf{x} and which also depends heavily on \mathbf{y} .

Thus, the functions which are not heavily binary are exactly the almost unary functions of Definition 3.1, and the heavily binary clones are exactly those $\mathcal{C} \subseteq \mathcal{O}$ which satisfy $\mathcal{C}^{(2)} \not\subseteq \mathcal{U}$.

5.2. EXAMPLE. Let $H : \lambda \times \lambda \rightarrow \lambda$ be a function satisfying

$$(*) \quad \forall \alpha > 0 : \quad H(\alpha, 1) = \alpha = H(0, \alpha).$$

[E.g., the max function has this property.] Then H depends heavily on \mathbf{x} and \mathbf{y} .

Conversely, if \mathcal{C} is a clone containing all unary functions and at least one heavily binary function, then \mathcal{C} contains a function H satisfying $(*)$ above.

The following example shows that there are nontrivial heavily binary clones above $\mathcal{O}^{(1)}$.

5.3. EXAMPLE. We will write $[X]^{<n}$ for the family of subsets of X of size $< n$, and $[X]^{<\aleph_0}$ for the family of finite subsets of X .

(1) Let \mathcal{Q} be the set of all functions $f \in \mathcal{O}$ such that $f : X^n \rightarrow X$ for some n , and there is a function $Q : X \rightarrow [X]^{<\aleph_0}$ satisfying

$$\forall x_1, \dots, x_n : \quad f(x_1, \dots, x_n) \in Q(x_1) \cup \dots \cup Q(x_n).$$

(2) Let \mathcal{P} be the set of all functions $f \in \mathcal{O}$ such that $f : X^n \rightarrow X$ for some n , and there is some k and a function $P : X \rightarrow [X]^{<k}$ satisfying

$$\forall x_1, \dots, x_n : \quad f(x_1, \dots, x_n) \in P(x_1) \cup \dots \cup P(x_n).$$

Then:

- (A) \mathcal{P} and \mathcal{Q} are clones.
- (B) $\mathcal{P} \subseteq \mathcal{Q} \subseteq \mathcal{O}$.
- (C) \mathcal{P} contains a heavily binary function, as well as all unary functions.
- (D) If X is finite, then trivially $\mathcal{P} = \mathcal{Q} = \mathcal{O}$.
- (E) If X is countably infinite, then $\mathcal{P} \subsetneq \mathcal{Q} = \mathcal{O}$.
- (F) If X is uncountable, then $\mathcal{P} \subsetneq \mathcal{Q} \subsetneq \mathcal{O}$.

We leave the verification of this fact to the reader.

5.4. THEOREM. *Assume that $\lambda \rightarrow (\lambda)_2^2$, i.e., λ is weakly compact or $\lambda = \aleph_0$. Then there is a unique precomplete clone which contains all unary functions and is heavily binary.*

By Example 5.3 there are nontrivial heavily binary clones above $\mathcal{O}^{(1)}$, so by Fact 1.3(3) there must be at least one precomplete such clone. So it is enough to show the following: Whenever $\mathcal{C}_1, \mathcal{C}_2$ are heavily binary clones on λ and $\mathcal{O}^{(1)} \subseteq \mathcal{C}_1 \cap \mathcal{C}_2$, then $\text{cl}(\mathcal{C}_1 \cup \mathcal{C}_2) = \mathcal{O}$ implies $\mathcal{C}_1 = \mathcal{O}$ or $\mathcal{C}_2 = \mathcal{O}$.

In 5.11 we will give an explicit description of this precomplete clone.

To make the proof clearer, we need a few definitions and lemmas.

5.5. DEFINITION. For $S \subseteq \lambda$, let

$$\Delta_S = \{(\alpha, \beta) \in S \times S : \alpha > \beta\}, \quad \nabla_S = \{(\alpha, \beta) \in S \times S : \alpha < \beta\}.$$

We let $\nabla\Delta_S := \nabla_S \cup \Delta_S = \{(\alpha, \beta) \in S \times S : \alpha \neq \beta\}$.

We will write Δ, ∇ and $\nabla\Delta$ as abbreviations for $\Delta_\lambda, \nabla_\lambda$ and $\nabla\Delta_\lambda$, respectively.

5.6. DEFINITION. For $\bar{\alpha} = (\alpha_1, \alpha_2, \alpha_3, \alpha_4) \in \lambda^4, \bar{\beta} = (\beta_1, \beta_2, \beta_3, \beta_4) \in \lambda^4$ we define $\bar{\alpha} \sim \bar{\beta}$ iff $\forall i, j \in \{1, 2, 3, 4\} : (\alpha_i < \alpha_j \Leftrightarrow \beta_i < \beta_j)$.

5.7. DEFINITION. Let $F : \nabla\Delta_S \rightarrow \lambda$. We say that F is *canonical* on S if for all $\bar{\alpha} \sim \bar{\beta}$: If $F(\alpha_1, \alpha_2) < F(\alpha_3, \alpha_4)$, then $F(\beta_1, \beta_2) < F(\beta_3, \beta_4)$.

[This also implies, for all $\bar{\alpha} \sim \bar{\beta}$: If $F(\alpha_1, \alpha_2) = F(\alpha_3, \alpha_4)$, then $F(\beta_1, \beta_2) = F(\beta_3, \beta_4)$.]

5.8. FACT. *Let $F : \nabla\Delta \rightarrow \lambda$. If $\lambda \rightarrow (\lambda)_2^2$, then*

$$\forall S \in [\lambda]^\lambda \exists S' \in [S]^\lambda : F \text{ is canonical on } S'.$$

The proof uses the partition relation $\lambda \rightarrow (\lambda)_n^4$ for some large n . We leave the details to the reader. See also Fact 6.2.

Given a finite set F_1, \dots, F_n of functions from $\nabla\Delta \rightarrow \lambda$, we can apply the lemma repeatedly and get a single set $S \in [\lambda]^\lambda$ on which all F_i are canonical. This yields Lemma 15 of II, §2 in [Ga65].

5.9. LEMMA. Assume that F is canonical on $\nabla\Delta_S$, where $S \in [\lambda]^\lambda$. Then:

(1) $F \upharpoonright \Delta_S$ has one of the following properties:

- $F \upharpoonright \Delta_S$ is 1-1 [typical examples: pr, pr_{Δ_S}].
- $F \upharpoonright \Delta_S$ depends injectively on the first coordinate: $F(x, y) = g(x)$ for some 1-1 function g [typical examples: π_1^2 , \max].
- $F \upharpoonright \Delta_S$ depends injectively on the second coordinate: $F(x, y) = g(y)$ for some 1-1 function g [typical examples: π_2^2 , \min].
- $F \upharpoonright \Delta_S$ is constant.

(2) Similarly for $F \upharpoonright \nabla_S$.

(3) If at least one of $F \upharpoonright \Delta_S$, $F \upharpoonright \nabla_S$ is 1-1, then $F[\Delta_S] \cap F[\nabla_S] = \emptyset$, or F is symmetrical on $S \times S$ (i.e., $F(x, y) = F(y, x)$).

Proof. We may assume $S = \lambda$.

(1) and (2) are easy. For (3), assume that $F(\alpha, \beta) = F(\delta, \gamma)$ with $\alpha < \beta$, $\gamma < \delta$. We have to distinguish several cases:

Case 1: $\alpha = \gamma$, $\beta = \delta$. Since $F(\alpha, \beta) = F(\beta, \alpha)$, and F is canonical, we have $F(x, y) = F(y, x)$ for all x, y , so F is symmetrical.

Case 2: $\alpha = \gamma < \beta < \delta$. So $F(\alpha, \beta) = F(\delta, \alpha)$. Pick any β', δ' with $\delta < \beta' < \delta'$. Then $(\alpha, \beta, \gamma, \delta) \sim (\alpha, \beta, \gamma, \delta')$, so $F(\alpha, \beta) = F(\delta, \alpha)$ implies $F(\alpha, \beta) = F(\delta', \alpha)$; this means $F(\delta', \alpha) = F(\delta, \alpha)$. Similarly we find $F(\alpha, \beta) = F(\alpha, \beta')$. So F is neither 1-1 on Δ nor 1-1 on ∇ .

Other cases: Similar to case 2. ■

5.10. LEMMA. Let \mathcal{C} be a clone containing all unary functions. If \mathcal{C} contains a heavily binary function H and also a canonical function F which is 1-1 on Δ (or even: canonical on some $\nabla\Delta_S$ and 1-1 on Δ_S , for some $S \in [\lambda]^\lambda$), then $\mathcal{C} = \mathcal{O}$.

Proof. By 1.3, it is enough to find a function $g \in \mathcal{C}$ which is 1-1 on $\nabla\Delta$.

If F is symmetrical and 1-1 on Δ (and also 1-1 on ∇ , of course), then we may assume (replacing F by $h \circ F$ for some appropriate $h \in \mathcal{O}^{(1)}$, if necessary), that $F(x, y) > \max(x, y)$ for all x, y . We claim that the function

$$(x, y) \mapsto F(x, F(x, y))$$

is 1-1 on $\nabla\Delta$. Indeed, if $F(x, F(x, y)) = F(x', F(x', y'))$, then

$$\text{either } x = x', F(x, y) = F(x', y'), \text{ or } x = F(x', y'), F(x, y) = x'.$$

In the first case we get either $y = y'$ directly, or $x = y'$, $y = x$, so again $y = y'$. The second case leads to a contradiction: $x = F(x', y') > x'$, $x < F(x, y) = x'$.

So we assume now that $F : \lambda \times \lambda \rightarrow \lambda$ is canonical but not symmetrical. By Lemma 5.9, we know that $F[\Delta] \cap F[\nabla] = \emptyset$. Replacing F by $h \circ F$ for an appropriate $h \in \mathcal{O}^{(1)}$, we may assume that

- $F \upharpoonright \Delta$ is constantly 0.
- $F \upharpoonright \nabla$ takes only even values > 0 , and is 1-1.

Since \mathcal{C} contains a heavily binary function, \mathcal{C} contains some function H with $H(0, x) = x = H(x, 1)$ for all $x > 0$. Now check that the map $(x, y) \mapsto H(Fxy, Fyx + 1)$ is a pairing function. ■

Proof of Theorem 5.4. Assume that τ is a term for a function in $\text{cl}(\mathcal{C}_1 \cup \mathcal{C}_2)$ representing a 1-1 function on $\lambda \times \lambda$. Find a set $S \subseteq \lambda$ of size λ such that $\tau \upharpoonright S$ is canonical (see Definition 5.7). Since \mathcal{C} contains all unary functions, \mathcal{C} also contains a monotone bijection between S and λ , so we will assume that τ , as well as every subterm of τ , is canonical on λ .

Let Θ be the set of subterms of τ .

Let $U_\Delta \subseteq \Theta$ (and $U_\nabla \subseteq \Theta$) be the set of those terms σ which induce unary functions on Δ (∇ , respectively), i.e.,

$$U_\Delta = \{\sigma \in \Theta : \exists f \in \lambda^\lambda, [\forall(\alpha, \beta) \in \Delta : \sigma(\alpha, \beta) = f(\alpha)] \text{ or} \\ [\forall(\alpha, \beta) \in \Delta : \sigma(\alpha, \beta) = f(\beta)]\}.$$

Let σ be a minimal subterm of Θ which is not in $U_\Delta \cap U_\nabla$, say $\sigma \notin U_\nabla$.

Let G be the outermost function in the term σ , say $G \in \mathcal{C}_1$, G n -ary. It remains to show that \mathcal{C}_1 contains a pairing function.

All proper subterms of σ represent unary functions, so there are n functions f_1, \dots, f_n and some $k \leq n$ with

$$\forall \alpha < \beta : \sigma(\alpha, \beta) = G(f_1(\alpha), \dots, f_{k-1}(\alpha), f_k(\beta), \dots, f_n(\beta)).$$

So the function induced by σ (which we again call σ) is in \mathcal{C}_1 . Now $\sigma \upharpoonright \Delta$ is not essentially unary. But σ is canonical, so by Lemma 5.10 we have a pairing function in \mathcal{C}_1 .

This concludes the proof of Theorem 5.4. ■

Gavrilov [Ga65] gave the following explicit description of the precomplete heavily binary clone for the case of a countable base set. It turns out that the same description works also on weakly compact cardinals.

5.11. EXAMPLE. Let $T_2 := \{g \in \mathcal{O}^{(2)} : \forall S \in [\lambda]^\lambda : \text{neither } g \upharpoonright \Delta_S \text{ nor } g \upharpoonright \nabla_S \text{ are 1-1}\}$. Equivalently, let T_2 be the set of all $g \in \mathcal{O}^{(2)}$ such that:

For all 1-1 functions h_1, h_2 : $g(h_1, h_2)$ is neither 1-1 on Δ nor 1-1 on ∇ .

Let $\mathcal{T} := \text{Pol } T_2$. Then:

- (1) \mathcal{T} is a clone, $\mathcal{T} \supseteq T_2$.
- (2) All unary functions are in \mathcal{T} .
- (3) The binary function pr_Δ (see 3.4) is not in \mathcal{T} .
- (4) The binary function max is in \mathcal{T} .
- (5) \mathcal{T} is a precomplete clone.

Hence, \mathcal{T} is the unique precomplete clone which is heavily binary and contains all unary functions.

Proof. (1)–(4) are easy (use Fact 5.8).

Proof of (5): Let \mathcal{C} be a clone properly containing \mathcal{T} . So \mathcal{C} contains a function which is 1-1 on some set Δ_S . Now use Lemma 5.10 to see that $\mathcal{C} = \mathcal{O}$. ■

6. Appendix: set-theoretic assumptions

6.1. DEFINITION. Let λ, μ, n, c be cardinals (usually: λ and μ infinite, n finite). The *partition symbol*

$$\lambda \rightarrow (\mu)_c^n$$

says: Whenever $[\lambda]^n$, the set of subsets of λ of cardinality n , is partitioned into c classes (i.e., whenever $f : [\lambda]^n \rightarrow C$, where $|C| = c$), then there is a subset $A \subseteq \lambda$ with at least μ elements such that all subsets of A of size n are in the same equivalence class (i.e., the restriction of f to $[A]^n$ is a constant function).

For example, the infinitary Ramsey theorem

$$\aleph_0 \rightarrow (\aleph_0)_2^2$$

says: whenever the edges of a complete (undirected) graph on countably many vertices are colored with 2 colors, then there is an infinite complete subgraph, all of whose edges have the same color.

We will be mainly interested in the situation $\lambda \rightarrow (\lambda)_2^2$. If $\lambda \rightarrow (\lambda)_2^2$, and λ is an uncountable cardinal, then λ is called *weakly compact*.

6.2. FACT. *If $\lambda \rightarrow (\lambda)_2^2$, then for all finite n, c we have $\lambda \rightarrow (\lambda)_c^n$.*

(In Fact 5.8, we use this property in the particular case of $n = 4$ and some large number c , approximately $c = 3^{256}$.)

The property $\lambda \rightarrow (\lambda)_2^2$ is a rather strong statement, i.e., it has many interesting consequences. Therefore, its mere negation,

$$\lambda \not\rightarrow (\lambda)_2^2$$

or explicitly:

- There is a map $f : [\lambda]^2 \rightarrow \{0, 1\}$ such that for any $A \subseteq \lambda$ of cardinality λ the function $f \upharpoonright [A]^2$ is not constant [i.e., is onto $\{0, 1\}$]

is a rather weak property of λ . There is, however, a strengthening of this negative partition relation which already yields interesting consequences.

6.3. DEFINITION. The statement $\lambda \not\rightarrow [\lambda]_\lambda^2$, the *negative square bracket partition relation*, means:

- There is a map $f : [\lambda]^2 \rightarrow \lambda$ such that for any $A \subseteq \lambda$ of cardinality λ the function $f \upharpoonright [A]^2$ is *onto* λ .

We will now consider an even stronger property of λ :

6.4. DEFINITION. Let $\lambda \geq \aleph_0$ and μ be cardinals. The statement (or, depending on your point of view, the “principle” or “axiom”) $\text{Pr}(\lambda, \mu)$ is defined as follows:

- There is a symmetric $c : \lambda \times \lambda \rightarrow \mu$ with the following property: For all $k \in \omega$, for all sequences $(a_i : i < \lambda)$ of pairwise disjoint subsets of λ of size k , and for all $c_0 \in \mu$,
 there are $i < j < \lambda$ such that $c \upharpoonright (a_i \times a_j)$ is constant with value c_0 .

Note that if we consider the case $\mu = \lambda$, and weaken the conclusion by allowing only $k = 1$, we get just $\lambda \not\rightarrow [\lambda]_\lambda^2$.

This statement as well as several variants of it are discussed in [Sh:g, III.4 and appendix 1]. What we call $\text{Pr}(\lambda, \mu)$ corresponds to $\text{Pr}_1(\lambda, \lambda, \mu, \aleph_0)$ there.

While the property $\text{Pr}(\lambda, \lambda)$ is quite strong (in particular: sufficiently strong to prove the result in Section 4), it turns out that it is not so rare: $\text{Pr}(\lambda, \lambda)$ holds for many successor cardinals already in ZFC without extra axioms. More general results (with proofs) can be found in Chapter III of [Sh:g], and also in [Sh 535] and [Sh 572].

(1) If there is a nonreflecting $S \subseteq \{\delta < \aleph_2 : \text{cf}(\delta) = \aleph_0\}$, then $\text{Pr}(\lambda, \lambda)$ holds. See [Sh:g, III.4.6C(6)].

(2) If $\kappa \geq \aleph_1$ is regular, then $\text{Pr}(\kappa^+, \kappa^+)$. See [Sh:g, III.4.8(1)], and [Sh 572, Theorem 1.1] for the proof of $\text{Pr}(\aleph_2, \aleph_2)$. For $\kappa > 2^{\aleph_0}$ there is a proof in [Sh 280].

(3) If κ is singular, and the set of Jonsson cardinals (= cardinals without a Jonsson algebra) is bounded in κ , then $\text{Pr}(\kappa^+, \kappa^+)$ holds. In particular, $\text{Pr}(\aleph_{\omega+1}, \aleph_{\omega+1})$ holds. See [Sh 535, 1.18].

References

- [CN2] W. W. Comfort and S. Negrepointis, *The Theory of Ultrafilters*, Grundlehren Math. Wiss. 211, Springer, Berlin, 1974.
- [DR85] R. O. Davies and I. G. Rosenberg, *Precomplete classes of operations on an uncountable set*, Colloq. Math. 50 (1985), 1–12.
- [Ga65] G. P. Gavrillov, *On functional completeness in countably-valued logic*, Probl. Kibernetiki 15 (1965), 5–64 (in Russian).
- [Ko89] S. Koppelberg, *Handbook of Boolean Algebras*, Vol. 1, North-Holland, 1989.
- [Ku83] K. Kunen, *Set Theory: An Introduction to Independence Proofs*, Stud. Logic Found. Math. 102, North-Holland, 1983.

- [Ro76] I. G. Rosenberg, *The set of maximal closed classes of operations on an infinite set A has cardinality $2^{2^{|A|}}$* , Arch. Math. (Basel) 27 (1976), 561–568.
- [Sh:g] S. Shelah, *Cardinal Arithmetic*, Oxford Logic Guides 29, Oxford Univ. Press, 1994.
- [Sh 280] —, *Strong negative partition above the continuum*, J. Symbolic Logic 55 (1990), 21–31.
- [Sh 535] —, *Further on colouring*, Arch. Math. Logic, submitted; math.LO/9808138 ⁽¹⁾.
- [Sh 572] —, *Colouring and non-productivity of \aleph_2 -cc*, Ann. Pure Appl. Logic 84 (1997), 153–174; math.LO/9609218.
- [Sz86] Á. Szendrei, *Clones in universal algebra*, Presses Univ. Montréal, Montréal, Qué., 1986.

Algebra
 TU Wien
 Wiedner Hauptstrasse 8-10/118.2
 A-1040 Wien, Austria
 E-mail: Martin.Goldstern@tuwien.ac.at
<http://www.tuwien.ac.at/goldstern/>

Department of Mathematics
 Hebrew University of Jerusalem
 91904 Jerusalem, Israel
 E-mail: shelah@math.huji.ac.il
<http://math.rutgers.edu/~shelah>

*Received 24 July 2000;
 in revised form 10 September 2001*

⁽¹⁾ References of the form math.XX/... refer to the arXiv.org preprint archive.