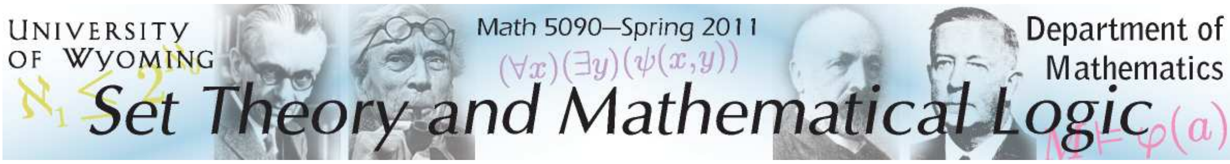

Ultraproducts

(Handout April, 2011)

G. Eric Moorhouse, University of Wyoming

Let $\{M_\xi\}_{\xi\in X}$ be an indexed collection of structures over a common language L. The product $\Pi = \prod_{\xi\in X} M_\xi$ is also a structure over this same language L.

Example. Suppose each M_ξ is a field. We take for L the language of rings (which serves also as the common language of fields). Thus L has constants 0 and 1, and two binary operations: $+$ and $\times$. (Possibly other symbols as well: a unary function symbol $-$ for additive inverses, and a unary function for multiplicative inverses, perhaps represented by the superscript symbol $^{-1}$. Let's not worry about these for now.) Then Π is a structure over the same language L: addition and multiplication are defined componentwise. We interpret the symbol 0 in Π as the element having 0 in each component (or more correctly, in component ξ we take the zero element of M_ξ, as the interpretation of the symbol 0 may vary from component to component. Similarly 1 is interpreted in Π as the element whose ξ-component is interpreted as the identity element of M_ξ. If each M_ξ is a field, then it is easily checked that $\Pi = \prod_{\xi\in X} M_\xi$ is a ring; but not a field for $|X| > 1$, since Π has zero divisors. To correct this we need ultrafilters.

Let $\mathfrak{U}$ be an ultrafilter on X. Then $M = \Pi/\mathfrak{U}$ is also a structure over L; this is known as an *ultraproduct* of the M_ξ's. Elements of $M = \Pi/\mathfrak{U}$ are equivalence classes of Π where two elements $x = (x_\xi)_{\xi\in X}$ and $y = (y_\xi)_{\xi\in X}$ in Π are equivalent if 'almost all' coordinates agree (relative to $\mathfrak{U}$). This means that $x_\xi = y_\xi$ for all $\xi \in U$, where $U \in \mathfrak{U}$. (Check, using the definition of an ultrafilter, that this is an equivalence relation.) All functions and relations on Π are defined coordinatewise, and are well-defined (i.e. their value is independent of the choice of representatives for equivalence classes).

For example if f is a binary function symbol in L and $f(x_\xi, y_\xi) = z_\xi$ in M_ξ, then $f(x,y) = z$ in M where $x = (x_\xi)_{\xi\in X}$, $y = (y_\xi)_{\xi\in X}$, $z = (z_\xi)_{\xi\in X}$. (Although one function symbol f from L is used here, its interpretation in each M_ξ may be different for each ξ.) Altering the coordinates of x or y on a 'small' set of ξ's (i.e. for ξ in some non-ultrafilter subset of X) preserves the equivalence class of $f(x,y)$. If R is a unary relation symbol in L and $x = (x_\xi)_{\xi\in X} \in \Pi$, then either $R(x_\xi)$ holds for almost all x, in which case we say that $R(x)$ holds in M; or $R(x_\xi)$ fails for almost all ξ, in which case we say that $R(x)$ does not hold in M. Again, the question is whether or not the subset $\{\xi \in X : R(x_\xi)\} \subseteq X$ is an ultrafilter set; and the answer to this question does not change if we alter a small subset of the coordinates x_ξ.

If the structure $M_\xi = M_0$ is independent of ξ, then $\Pi = M_0^X$ and the ultraproduct $M = \Pi/\mathfrak{U} = M_0^X/\mathfrak{U}$ is called simply an *ultrapower.*

If Σ is a set of first-order statements in the language of L, and $M_\xi \vDash \Sigma$ for each ξ, then the ultraproduct $M \vDash \Sigma$. (However, Π itself does not model Σ in general; the ultrafilter plays an indispensible role in this construction of new models from old.)

If the ultrafilter $\mathfrak{U}$ is principal, generated by the singleton $\{\xi_0\}$, say (i.e. $\mathfrak{U} = \{A \subseteq X : \xi_0 \in A\}$), then it is easy to see that the ultraproduct $M \cong M_{\xi_0}$ so the ultraproduct construction does not give anything new. Our interest is therefore in *nonprincipal* ultrafilters, which bear the promise of yielding new and interesting models.

1. Example: Hyperreals

Let $X = \omega = \{0, 1, 2, \ldots\}$ and let $\mathfrak{U}$ be a nonprincipal ultrafilter on ω. Then the ultrapower $\widehat{\mathbb{R}} = \mathbb{R}^\omega/\mathfrak{U}$ is the field of hyperreal numbers as described in the previous handout. Here L is the language of rings (and of fields). Each factor $M_\xi = \mathbb{R} \vDash \{\text{axioms for fields}\}$. Although $\Pi = \mathbb{R}^\omega$ is only a ring (not a field), the ultrapower $\widehat{\mathbb{R}} = \mathbb{R}^\omega/\mathfrak{U}$ is a field.

2. Example: Ultraproducts of Finite Fields

Consider the set of primes $X = \{2, 3, 5, 7, 11, \ldots\}$ and let $\mathfrak{U}$ be a nonprincipal ultrafilter on X. For each prime $p \in X$, denote by $\mathbb{F}_p$ the field of order p. Then $R = \prod_{p \in X} \mathbb{F}_p$ is a ring, and the ultraproduct $F = R/\mathfrak{U}$ is a field. For each positive integer n, the statement $n \neq 0$ holds in $F = R/\mathfrak{U}$ (since it fails for only a small set of primes $p \in X$, namely those primes dividing n). So F has characteristic zero, i.e. F is an extension of $\mathbb{Q}$. We may check that F is uncountable, so the degree of the extension $[F : \mathbb{Q}]$ is also uncountable. However, F has many of the same algebraic properties as finite fields, as we proceed to describe.

Given a finite field K of order q and a positive integer r, the field K has a unique extension field of degree r (up to isomorphism), namely the field of order q^r. The property that a field has an extension of degree r which is unique up to isomorphism, is expressible as a statement θ_r in the first order theory of fields. Since $\mathbb{F}_p \vDash \theta_r$ for every $p \in X$ and $r \geqslant 1$, it follows that $F \vDash \theta_r$ also, where $F = R/\mathfrak{U}$ as above. It is very difficult to find explicit examples of fields of characteristic zero with this property! Note that $\mathbb{C} \vDash \theta_r$ iff $r = 1$ (the field of complex numbers has no nontrivial finite extensions since it is algebraically closed). Also $\mathbb{R} \vDash \theta_r$ iff $r \in \{1, 2\}$ (the only nontrivial finite extension of $\mathbb{R}$ is $\mathbb{C}$). Although $\mathbb{Q}$ has finite extensions of every degree $r \geqslant 2$, it has infinitely many such extensions; for example the quadratic extensions of $\mathbb{Q}$ have the form $\mathbb{Q}[\sqrt{d}\,]$ where d is a squarefree integer; and for different choices of d, these extensions are nonisomorphic. The same is true of every algebraic number field, i.e. finite extension field $K \supseteq \mathbb{Q}$: we have $K \vDash \theta_r$ iff $r = 1$.

3. Ultraproducts of Projective Planes

Consider the language of point-line incidence geometry. Objects are distinguished (using a unary relation) as either points or lines. There is a single binary relation called incidence; pairs which satisfy this relation necessarily consist of one point and one line. Consider a structure P in this language, i.e. P is a system of points and lines, some pairs of which are incident. We say P is a *projective plane* if it satisfies three first-order axioms, which embody the following requirements:

(P1) Any two distinct points are incident with exactly one line.
(P2) Any two distinct lines are incident with exactly one point.
(P3) There exist four points, of which no three are collinear.

Given a family of projective planes P_ξ indexed by $\xi \in X$, and an ultrafilter $\mathfrak{U}$ on X, the ultraproduct $P = (\prod_{\xi \in X} P_\xi)/\mathfrak{U}$ is a projective plane.

There are many known examples of projective planes. The best-known examples, which we define here, are the *classical projective planes*, constructed as follows. Given a field F, the structure $P^2(F)$ has as its points and lines the subspaces of F^3 of dimension 1 and 2 respectively; here F^3 denotes a 3-dimensional vector space over F. If F is finite of order q, then $P^2(F)$ has exactly q^2+q+1 points and the same number of lines; if F is infinite, then the set of points (or lines) of $P^2(F)$ is also infinite, of the same cardinality as F. Since any two distinct 1-spaces of F^3 span a 2-space, $P^2(F) \vDash$ (P1). Since any two distinct 2-subspaces of F^3 intersect in a 1-space, $P^2(F) \vDash$ (P2). It is a routine exercise to check that $P^2(F) \vDash$ (P3), so $P^2(F)$ is a projective plane.

An important theorem in projective plane geometry is that a projective plane P is classical (i.e. isomorphic to $P^2(F)$ for some field F) iff it satisfies a first-order condition (P4) asserting that for any configuration of points and lines in P as shown, the points R_0, R_1 and R_2 must be collinear:

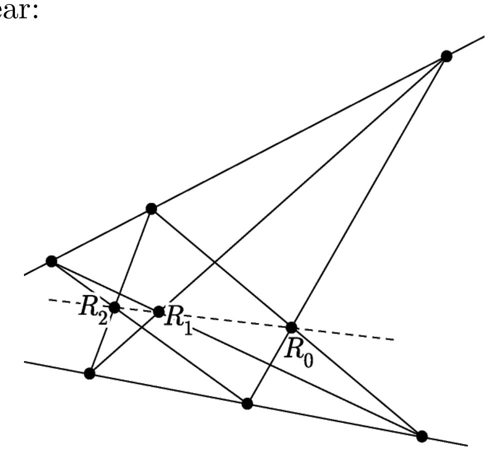

There are many projective planes known which do not satisfy this condition, and so are *nonclassical.*

An ultraproduct of classical projective planes $P = \left(\prod_{\xi\in X} P^2(F_\xi)\right)/\mathfrak{U}$, must also be classical. This is because $P^2(F_\xi) \vDash$ (P4) for each ξ, which forces $P \vDash$ (P4). In particular, an ultraproduct of the finite classical planes $P^2(\mathbb{F}_p)$ must be isomorphic to $P^2(F)$ for some field F. It should come as no surprise that F is actually the ultraproduct of the finite fields $\mathbb{F}_p$ as described in Section 2 above. The resulting plane, while infinite (indeed, uncountable) shares many properties with the classical finite planes.

4. Ultraproducts of Graphs

The language of graphs requires a single binary relation symbol for adjacency. The axioms for graph theory require that this relation be symmetric (at least for ordinary graphs, where edges are not directed) and irreflexive (if we do not want loops). If $\{\Gamma_\xi\}_{\xi\in X}$ is an indexed family of graphs, and $\mathfrak{U}$ is an ultrafilter on X, then the ultraproduct $\Gamma = (\prod_{\xi\in X}\Gamma_\xi)/\mathfrak{U}$ is also a graph. If Σ is a set of first-order sentences in the language of graphs, and $\Gamma_\xi \vDash \Sigma$ for all $\xi \in X$, then $\Gamma \vDash \Sigma$. For example, if every Γ_ξ is triangle-free (i.e. having no 3-cycles), then Γ is also triangle-free. It is promising to look for new and interesting graphs this way, in particular infinite graphs which share many properties of known families of finite graphs.

An ultraproduct of graphs of degree k has degree k. Here we can replace 'degree k' by 'degree at most k' or by 'degree $\in K$ for any finite set K of natural numbers. Or we can replace 'degree' by 'diameter' or 'girth' throughout. In each case the property in question is expressible in first order logic.

However, connectedness is not preserved by ultraproducts, as this is not a first-order property. For example, if Γ_n is a path of length n for each $n \geqslant 1$, then a (nonprincipal) ultraproduct $\Gamma = (\prod_{n\geqslant 1}\Gamma_n)/\mathfrak{U}$ has uncountably many infinite paths as its connected components, each path having 0 or 1 (but never 2) endpoints. Although the property of having diameter 3 (or diameter k for any other fixed k) is expressible in first order logic; but we cannot quantify over k, only over vertices.

A (nonprincipal) ultraproduct of bipartite graphs need not be bipartite, as the property of being bipartite is not expressible in the first order theory of graphs. (One can of course use instead the first order theory of bipartite graphs, in which a unary relation is added to the language in order to distinguish the two parts of the partition.) More generally, the property of having chromatic number k is not a first-order property of graphs, and so this property need not be preserved by ultraproducts. (The case of chromatic number 2 is equivalent to the property of bipartiteness.) For example if K_2 is the graph on two vertices with one edge, then K_2 has chromatic number 2 (i.e. is bipartite), whereas a nonprincipal ultrapower of K_2 is a complete graph on an uncountable vertex set, so its chromatic number is uncountable.

5. The Compactness Theorem

We obtained the Compactness Theorem for first order logic as an easy consequence of the Soundness and Completeness Theorem; however, we did not work through the details of the proof of the Soundness and Completeness Theorem. Here, for the first time in class this semester, we give a reasonably self-contained proof of the Compactness Theorem for first order logic.

Let Σ be a set of first order sentences in some language L. Let X be the collection of all finite subsets of Σ. We suppose that for each $A \in X$, there exists a model $M_A \vDash A$. We must show that $M \vDash \Sigma$ for some L-structure M. (Without loss of generality, Σ is infinite; otherwise the desired conclusion follows immediately.) We will obtain M as an ultraproduct over X. However, we must be careful in the choice of ultrafilter $\mathfrak{U}$ on X; it is not sufficient for $\mathfrak{U}$ to be nonprincipal. Also note that we require an ultrafilter on X rather than on Σ itself.

For each $A \in X$, consider the collection of all finite supersets of A:

$$A^+ = \{B \in X \,:\, B \supseteq A\} \subseteq X.$$

Note that for any $A_1, A_2, \ldots, A_n \in X$, we have $A_1 \cup A_2 \cup \cdots \cup A_n \in X$ since this is a finite union of finite subsets of X, so

$$A_1^+ \cap A_2^+ \cap \cdots \cap A_n^+ = (A_1 \cup A_2 \cup \cdots \cup A_n)^+ \neq \varnothing.$$

Since the collection $\{A^+ \,:\, A \in X\}$ of subsets of X satisfies the finite intersection property, it generates a filter $\mathfrak{F}$ on X. Extend $\mathfrak{F}$ to an ultrafilter $\mathfrak{U} \supseteq \mathfrak{F}$ on X, and consider the ultraproduct $M = \big(\prod_{A\in X} M_A\big)/\mathfrak{U}$.

We must show that $M \vDash \Sigma$. Equivalently, for every finite subset $B \subseteq \Sigma$, we show that $M \vDash B$. Clearly $M_A \vDash B$ for all $A \in B^+$, since in this case $M_A \vDash A \supseteq B$. Note that $B^+ \in \mathfrak{U}$ since $B^+ \in \mathfrak{F}$. We have $M_A \vDash B$ for almost all $A \in X$ (with respect to the ultrafilter $\mathfrak{U}$) so $M \vDash B$ as required.