

Implementing Mathematical Objects in Set Theory

Thomas Forster

July 7, 2007

In general little thought is given to the *general* question of how to implement mathematical objects in set theory. It is clear that—at various times in the past—people have gone to considerable lengths to devise implementations with nice properties. There is a literature on the evolution of the Wiener-Kuratowski ordered pair, and a discussion by Quine of the merits of an ordered-pair implementation that makes every set an ordered pair. The implementation of ordinals as Von Neumann ordinals is so attractive that it is universally used in all set theories which have enough replacement to prove Mostowski’s collapse lemma. I have frequently complained in the past about the widespread habit of referring to implementations of pairs (ordinals etc) as *definitions* of pairs (etc). My point here is a different one: generally little attention has been paid to the question of what makes an implementation a good implementation. In most cases of interest the merits of the candidates are uncontroversial. What I want to examine here is an example where there are competing implementations for ordered pairs, and—although it is clear to the *cognoscenti* and also (with a bit of arm-waving) plausible to the logician in the street that some of the impossible candidates are impossible, nobody has ever given a satisfactory explanation of why this is so.

The example I have in mind is the implementation of ordered pairs in Quine’s NF.¹ The complications attending implementations of mathematical entities in NF all arise from the failure in NF of unstratified replacement. This is highly significant, and for quite general reasons. In general, the successful implementation of a math-

¹This topic has never been given a thorough treatment in the literature, though it was discussed briefly in the closing pages of Lake’s Ph.D. Thesis [3].

emathical gadget into set theory will generate a typing discipline. For example, when one is implementing pairing into set theory, one does not generally care whether or not x should ever be equal to $\langle x, y \rangle$. There are exceptions to this: for example both the Hailperin [1] axiomatisation of NF and Gödel’s F -functions for generating L trade on the fact that ordered pairs are Wiener-Kuratowski ordered pairs and have particular set-theoretic structure) but generally expressions like “ $x = \langle x, y \rangle$ ” are regarded as syntactically aberrant, and—at that—aberrant in a fairly straightforward way. It is an important fact (one perhaps not sufficiently widely appreciated) that expressions which respect the obvious syntactic constraints turn out to be invariant under choice of implementation, and conversely. Indeed, the syntactic discipline wouldn’t be much use if this were not so! Equally important is that this equivalence implies the axiom scheme of replacement! For example, Mathias [unpublished] has shown that if we assume that $x \times y$ exists for all x and y irrespective of our choice of implementation of ordered pair then the axiom scheme of replacement holds. However in NF we know that unstratified replacement fails. This warns us that in NF it might really matter how we implement ordered pairs, in the sense that the truth-value of certain assertions about relations or functions or cartesian products (the feature common to these topics being the need for an implementation of ordered pair) will vary with our choice of implementation of ordered pair.

What is an implementation of ordered pair anyway? At the very least it must be a three-place relation $P(x, y, z)$ satisfying

$$(1) (\forall xy)(\exists!z)(P(z, y, z))$$

and

$$(2) (\forall z)(\forall x)(\forall x')(\forall y)(\forall y')(P(x, y, z) \wedge P(x', y'z) \rightarrow x = x' \wedge y = y')$$

Any P satisfying this will be said to be a *pairing relation*. What else can we insist that a pairing function should do? There are some things that it clearly cannot be asked to do. Pairs cannot be required to have any particular set-theoretic structure. There is a natural type-theoretic discipline proper to any use of ordered pairs (and this type discipline has nothing whatever to do with stratification à la NF(!)) and according to it expressions like ‘ $x \in \langle y, z \rangle$ ’ are

ill-typed. It doesn't mean that they are illformed or will lack truth-values once pairing has been implemented: clearly they will have truth-values. The point is merely that it is no part of the job of the implementation to give them one truth-value rather than another.

The kind of thing we could reasonably insist that an implementation of pairing should reproduce would be uncontroversial banalities of utterly elementary theories of things that require ordered pairs. One such theory is elementary (binary) relational algebra. This theory has operations like composition ($R \circ G$), inverse (R^{-1}) and boolean operations on relations (thought of as their graphs) over any fixed domain. This theory contains assertions like

$$\begin{aligned} R \subseteq S &\rightarrow R^{-1} \subseteq S^{-1} \\ R \subseteq S &\rightarrow R \circ T \subseteq S \circ T. \end{aligned}$$

Agreeing to reproduce truths like these commits us to having a notion of ordered pair that means that the composition of two (graphs of) relations is the (graph of) a relation, and so on.

(Bearing in mind that it is implementations in NF that we are considering) let us make at this stage the observation that if the composition of two relations R and S is to be a relation then $R \circ S$ which is of course

$$\{z : (\exists x \in R)(\exists y \in S)(\exists abc)(P(a, b, x) \wedge P(b, c, y) \wedge P(a, c, z))\} \quad (1)$$

had better be a stratified set abstract. That is to say, ' $(\exists x \in R)(\exists y \in S)(\exists abc)(P(a, b, x) \wedge P(b, c, y) \wedge P(a, c, z))$ ' must be stratified. This requires that ' $P(-, -, -)$ ' be stratified and that ' a ', ' b ' and ' c ' all receive the same type.

That is to say that in ' $P(-, -, -)$ ' the first two variables must receive the same type. (2)

It doesn't tell us anything about the type of the third variable. Similarly uncontroversial will be the expectation that every relation should have an inverse. However this won't tell us anything new. Consideration of the expression

$$R^{-1} = \{z : (\exists z' \in R)(\exists ab)(P(a, b, z') \wedge P(b, a, z))\}$$

will tell us that the first two arguments to ‘ $P(-, -, -)$ ’ must receive the same type in any stratification. Again, it tells us nothing about the type of the third argument.

This insight enables us to answer the point often made by people encountering Cantor’s theorem in NF for the first time. If we try to prove that a map $f : X \rightarrow \mathcal{P}(X)$ is not onto we find ourselves considering the diagonal set

$$\{x \in X : (\forall w \in f)(\forall X' \subseteq X)(P(x, X', w) \rightarrow x \notin X')\} \quad (3)$$

For us to be confident that the diagonal set is genuinely a set we would need ‘ $P(x, X', w)$ ’ to be stratified with ‘ x ’ one type lower than ‘ X' ’ and this of course we do *not* have.

However we can prove an analogue, which for many purposes is just as good: in some sense it will enable us to recover the same mathematics. Recall that ι is the singleton function, so that ι “ x is $\{\{y\} : y \in x\}$. (This notation does not presuppose that the graph of ι is a set!). Clearly $\{\{y\} : y \in x\}$ is a set, being the denotation of a stratified set abstraction. Next we attempt to prove that no function $f : \iota$ “ $X \rightarrow \mathcal{P}(X)$ can be surjective. This time the diagonal set is

$$\{x \in X : (\forall w \in f)(\forall X' \subseteq X)(P(\{x\}, X', w) \rightarrow x \notin X')\}$$

which can be seen (even before we eliminate the curly brackets in ‘ $\{x\}$ ’!) to be stratified. So we seem to have proved that there are fewer singletons than sets. But what about the singleton function—surely it is a bijection between ι “ V and V ? Yes, but its graph isn’t a set. And this is because, as we saw earlier, the two components of the ordered pair must be given the same type.

It may be worth thinking a little bit about what would happen were we prepared to change our definition of ordered pair so that ‘ $P(x, y, z)$ ’ were stratified with ‘ x ’ one type higher than ‘ y ’. Then the set abstract in (3) would be a set and the proof would succeed. We would have shown that X is indeed smaller than $\mathcal{P}(X)$. But what does “smaller than” mean with this definition of ordered pair? Since our definition no longer ensures that the composition of (the graphs of) two relations is a (graph of a) relation we find that equinumerosity no longer appears to be transitive.

This is the explanation that NF-istes offer to non NF-istes for the decision to opt for ordered pair functions that give their two inputs the same type. The explanation convinces most questioners. Or perhaps one should say that it *silences* them. People who come to introductory talks on NF generally want to know about how mathematics is done in NF and are correspondingly willing to refrain from picking fights over definitions of ordered pairs if such restraint on their part enables them to get on with what they came for.

The choice of pairing functions that that tradition has made for NF has resulted in our perforce making certain choices about which assertions about relations and functions we wish to come out true. Faced with a choice between making every set the same size as its set of singletons and ensuring that equinumerosity was an equivalence relation we decided to go for the pairing that makes equinumerosity an equivalence relation. This was certainly the correct thing to do, but can we explain why?

I believe we can, and that it is as follows. There are various banalities about pairing, relational algebra and functions that we can express in a strongly typed system that regards the components of the ordered pairs as having no internal structure. Nothing must be allowed to override the requirement on an implementation that it respect those banalities about pairing, relational algebra and functions that can be captured in this way. For example, the assertions that the composition of two relations always exists does not require us to look inside the components of the ordered pairs, as contemplation of formula (1) above will confirm. Similarly equinumerosity. The assertion that

If x and y are equinumerous and y and z are equinumerous,
then x and z are equinumerous

—although it requires us to look inside x , y and z —does not require us to look inside any of the components of the ordered pairs that we mention.

Contrast this with the desideratum for an ordered pair function of making x and ιx turn out to be the same size. We will find that if we state this properly we will be looking inside one of the components of an ordered pair—specifically to state that it is a singleton. It is worth making the point here that the expectation that x and ιx are the same size relies on an appeal to an instance

of the axiom of replacement. The failure of the singleton function to be a set according to all implementations of ordered pair satisfying (2) is in fact exactly what we want. We do not want to include in our spec for the implementation of the pairing function that it should make x and ιx appear to be the same size. That is not a job for the implementation of pairing; that is a job for the set existence axioms.

The point is not that all well-typed banalities should be accommodated. It should be conceded that the existence of compositions and converses of relations does depend on set existence axioms—albeit fairly trivial ones. The formula asserting existence of transitive closures of relations is also well-typed, in that it does not require us to look inside the components of the ordered pairs it discusses. However, it does require a bit more set theory—enough to perform inductive definitions—and so one should not expect an implementation of ordered pair to automatically deliver the existence of transitive closures. The point is rather that no well-typed banality should be sacrificed as part of an attempt to accommodate a less-strictly-typed assertion (such as the existence of the graph of the singleton function) which might be thought desirable.

I think the consideration I invoked a few paragraphs ago—that we cannot require of our pairing function that it deliver the truth (or falsehood) of any general assertion about sets, functions and relations that involve looking into the internal structure of components of ordered pairs—is completely general in the sense that analogous considerations apply to implementations of other mathematical entities.

However, these general considerations have left some points open. We have decided that the formula $P(x, y, z)$ (whichever formula it should turn out to be) that says that z is the ordered pair of x and y must be stratified with ‘ x ’ and ‘ y ’ receiving the same type. It doesn’t tell us what type ‘ z ’ should be given relative to ‘ x ’ and ‘ y ’. For example, the Wiener-Kuratowski ordered pair is perfectly acceptable in NF. We have to be more careful with Wiener-Kuratowski triples and n -tuples for higher n . The usual definition of ordered triples in the Wiener-Kuratowski style makes $\langle w, x, y \rangle$ the Wiener-Kuratowski pair $\langle w, \langle x, y \rangle \rangle$ where the embedded pair is Wiener-Kuratowski. This triple is unsatisfactory, since it makes ‘ x ’ and ‘ y ’ two types higher than ‘ w ’. A much better solution is to take

$\langle w, x, y \rangle$ to be $\langle \{\{w\}\}, \langle x, y \rangle \rangle$. where once again the two pairs are Wiener-Kuratowski: this makes ‘ x ’, ‘ y ’ and ‘ w ’ all the same type. A similar manoeuvre can be used for quadruples and higher types. This is the implementation used by Hailperin [1].

We should note that in NF we can actually prove that there is no pairing relation $P(x, y, z)$ where ‘ z ’ is one type lower than ‘ x ’ and ‘ y ’. Suppose there were; then the map $x \mapsto \{\langle x, x \rangle\}$ is an injection from V into ιV contradicting the fact that there are more sets than singletons.²

However, the pair that is always used in NF is the Quine pair. I shall not explain it here, since there are already adequate discussions of it in the literature. It has two quite desirable features. The first is that it makes everything into a pair. The second is that the formula $P(x, y, z)$ (that says that z is the Quine pair of x and y) makes ‘ x ’, ‘ y ’ and ‘ z ’ all the same type. We noticed that the considerations earlier did not constrain the type of ‘ z ’, but there is no doubt that having ‘ x ’, ‘ y ’ and ‘ z ’ all the same type makes life easier. It means that when we proceed to triples and quadruples etc as in the previous paragraph we do not have to wrap curly brackets around variables to ensure that all components of tuples are the same type.

Finally, some quite subtle considerations. We have resigned ourselves to the graph of the singleton function not being a set. Let us now consider the natural numbers: by defining \mathbb{N} to be the intersection of all sets containing the singleton of the empty set and closed under succ where $\text{succ}(x) =: \{w : (\exists y \in w)(w \setminus \{y\} \in x)\}$ we make no use of pairing functions.

Cogitations on stratifications like those in the previous paragraph will convince us that for a natural number n there is in general no reason to suppose that there will be a bijection between an arbitrary set of size n and the set $[0, n]$ of natural numbers less than n . This set, $[0, n]$, is finite and its cardinal is a natural number, and we notate this cardinal ‘ T^2n ’. Why ‘ T^2 ’? Why don’t we define this T function so that $Tn = |[0, n]|$? The point is that (check it!) $|[0, n]|$ is two types higher than n not one. For a variety of technical reasons it is more sensible to have as our defined term something that raises by one type than something that raises by two. Note that, although the assertion that each natural number counts the set of its predecessors is not stratified, there is no good reason to suppose it is refutable.

²And there is no difficulty proving Schröder-Bernstein!

Something similar happens with ordinals. If α is an ordinal, the set of ordinals strictly less than α is naturally wellordered, and therefore has a length which is an ordinal. What is this ordinal? For stratification reasons this ordinal will not be α but will turn out to be the result of applying a T -like function to α . Ward Henson, who was the first person to consider this function applied to ordinals rather than cardinals (see [2]), was sensitive to the difference between ordinals and cardinals in this respect, and he wrote the operation on ordinals with a ‘ U ’ not a ‘ T ’.³ $|[0, n]|$ is two types higher than n . How many types higher than ‘ α ’ is the ordinal of the set of ordinals below α ordered by magnitude? Let’s calculate it. Ordinals are implemented as isomorphism classes (which turn out to be sets) of wellorderings. So we consider the set of ordinals below α , and we wellorder it by magnitude. This gives us a set (‘ A ’ for the moment) of ordered pairs of ordinals, and we take its equivalence class under isomorphism, and this is the ordinal we want. It will of course be one type higher than A . But what is the type of A relative to the type of α ? The answer to this will depend on our choice of ordered pair! If we are using Quine pairs it will be one type higher than α , but if we are using Wiener-Kuratowski pairs the difference will be three!

We can of course also implement wellorderings not by means of ordered pairs, but as the set of their initial (or for that matter, their terminal) segments. One then implements ordinals as isomorphism classes of wellorderings as before. The fact that under any sensible implementation of ordered pair (or even without it, by using the initial segment coding) the collection of all ordinals is a set has the consequence that there must always be a nontrivial appearance of the T (or, if you are Ward Henson, the U) function to enable us to say that

$$T^k\alpha \text{ is the length of the ordinals below } \alpha: \tag{4}$$

If α counted the length of the ordinals below α we would be able to prove the Burali-Forti paradox. Therefore any *true* (4)-like assertion about the length of an initial segment of ordinals *must* involve a T -function. (This is sharp contrast to the case with natural numbers where the assertion that each natural number counts the set of its predecessors appears to be consistent—albeit strong.) The appear-

³Nowadays it usually written with a ‘ T ’, using overloading.

ance of the T function here is therefore not an artefact of our choice of implementation for ordered pairs or wellorderings: it is a genuine manifestation of the underlying mathematics associated with having a set of all ordinals.

Despite the inevitability of the appearance here of a T -function, there is nothing in the underlying mathematics to tell us what the exponent on it must be in formula (4)! This fact is generally known to *nfistes* but its significance seems not be understood even by them. The most helpful remark in this connection is probably the observation of Dana Scott's (personal communication) that NF is really a type theory not a set theory. It bears thinking about.

References

- [1] Hailperin, T. [1944] A set of axioms for logic. *Journal of Symbolic Logic* **9** pp. 1–19.
- [2] Henson, C.W. [1973] Type-raising operations in NF. *Journal of Symbolic Logic* **38** pp. 59–68.
- [3] Lake, J. [1974] Some topics in set theory. Ph.D. thesis, Bedford College, London University.

Thomas Forster
Department of Pure Mathematics and Mathematical Statistics
Centre For Mathematical Sciences
Wilberforce Road
Cambridge CB4 3NE