Carrying On with Infinite Decimals

by Jim Propp (University of Massachusetts, Lowell),
with the members of the MIT Chips and Rotors Research Initiative
(Giuliano Giacaglia, Holden Lee, Delong Meng,
Henrique Oliveira Pinto, XiaoLin (Danny) Shi, and Linda Zayas-Palmer)

(preliminary draft, dated 3/9/10;
not for distribution outside of CARRI)

We all know how to perform carries for *finite* decimals: start at the rightmost position and perform carries, progressing to the left, until no more carries are possible.

But how can we do this in the case of *infinite* decimals, where there's no rightmost position at which we can start?

When I ask "how can we do this", I am using the word "do" in an idealized and totally impractical sense. After all, if there are infinitely many digits in each of the numbers being added, then even if there are no carries to perform, we'll have to do infinitely many single-digit additions, which under most reasonable models of computation would take an infinite amount of time. In fact, just writing down the answer would take forever!

Since we would never be able to declare such a computation to be "done", in what sense could we talking about "doing" it?

This article provides one possible answer. Under the model of carrying described here, the sum of the numeral 0.999... with itself is emphatically 1.999... and not 2.000.... Indeed, I'll show that whenever you add two or more infinite decimal numerals, the only way the sum can end in an infinite string of 0's is if every summand ends in an infinite string of 0's.

Before I explain the general set-up or prove anything, let's discuss the sum $0.999\ldots + 0.999\ldots$. To prepare for what's to come, we'll write this as $0.9,9,9,\ldots + 0.9,9,9,\ldots$, using commas to separate the 1's place from the .1's place from the .01's place etc. We first add without doing any carrying, obtaining "digits" that exceed 9; specifically, we get the sum $0.18,18,18,\ldots$. We refer to the numbers between the commas as "generalized digits". A generalized digit can be any non-negative integer. If you like, you can think of this kind of generalized decimal numeral as representing an infinite version of the kind of primitive abacus in which counters are kept in different compartments, with each compartment representing a different place-value.

In our particular situation, all the compartments to the left of the decimal point are empty and all the compartments to the right of the decimal point have 18 counters. The carry-operation is the process of removing 10 counters from some compartment and replacing them by a single counter in the next compartment to the left; or, if you are working with the numerical representation, you reduce some generalized digit by 10 and increase the preceding generalized digit by 1. Of course, it isn't permitted to reduce a generalized digit by 10 if it is less than 10, since we aren't allowing our generalized digits to be negative.

In some ways in might be conceptually helpful (though wasteful of ink) to write zero.eighteen,eighteen,eighteen,... instead of 0.18,18,18,..., to discourage us from thinking that the use of the decimal notation for the finite positive integers between the commas has anything to do with the generalized decimal system that we are using as intermediate stages in calculations involving infinite decimals. But since most readers of this article will more readily ascertain the truth value of a proposition like "nine plus nine equals eighteen" when it is presented in the form $9+9 = 18$, I'll use decimal notation for the generalized digits.

The general procedure for adding k decimal numerals is to start by adding the numbers decimal place by decimal place, without any carries, obtaining a generalized decimal numeral with digits between 0 and $9k$. From this point on, we can forget about where the generalized decimal numeral (which we will hereafter call simply a "numeral") came from; it represents the sum we're looking for, but in a non-standard way, and our goal is to standardize it by doing carries, so that the sum is re-expressed as a numeral in which no digit exceeds 9, which we will call a *stable* numeral.

Returning to our example, if we start from the left and start carrying to the right, we successively turn 0.18,18,18,18,... into 1.8,18,18,18,... and 1.9,8,18,18,... and 1.9,9,8,18,... etc., where each position eventually becomes a 9 and thereafter stays a 9.

In each position, we have one and only one chance to perform a carry; so as long as we eventually perform a carry in every position (that is, a carry from every position to the position to its left), it doesn't matter in what order we do the carries. More specifically, consider the kth position to the right of the decimal point (the 10^{-k} place). Once there's been a carry from the kth position to the $k-1$st (causing a decrement of 10 in the kth position) and a carry from the $k+1$st position to the kth (causing an increment of 1 in the kth position), the kth position will contain the generalized digit

$18 - 10 + 1 = 9$, and this will remain true forever after. Hence, the last time the kth position changes, it contains a 9, so we may say that its state at "time infinity" is still 9.

Note that if we stabilize .18,18,18,... by performing carries in successive positions 2,1,4,3,6,5,..., each generalized digit to the right of the decimal point will eventually become (and remain) a 9. However, if we perform carries in successive positions 2,3,4,5,6,... and never get around to performing a carry in position 1, we obtain the generalized decimal numeral 0.19,9,9,9,..., not the desired answer 1.9,9,9,9,.... This merely confirms what was said above: if you want to arrive at the correct final answer, you need to make sure that any position that *can* execute a carry eventually *does* execute a carry.

Things are a little more subtle for resolving numerals like 0.9,10,9,10,...; here, performing a carry in one position can cause the generalized digit to the left of that position to become greater than or equal to 10, making possible a carry that wasn't possible before. For other examples, a chain reaction of such carries can change digits far to the left of the site of the original carry. So our stipulation "any position that can execute a carry eventually does" requires some clarification. Suppose at some stage in our carrying process a generalized digit in some particular position, say position k, becomes greater than 9; then we require that at some later time position k must execute a carry. We call this the Zero Infinite Procrastination (or ZIP) rule, and we call a sequence of carries "zippy" if this condition is satisfied.

You can check that one particular protocol for performing carries in a zippy way is to always perform a carry in the leftmost position that exceeds 9. However, it is helpful to consider the totality of all zippy carry-histories, and not just the history that arises from the "leftmost first" protocol.

Much worse than 0.9,10,9,10,... is the numeral 0.10,100,1000,.... If we look ahead to the fact that this numeral represents the divergent infinite sum $10/10 + 100/100 + 1000/1000 + \ldots$, we know we're heading for trouble. And indeed, the leftmost first protocol applied to this numeral causes each position to execute a carry infinitely many times. In fact, *any* zippy carry-history will have this feature. So the digits don't stabilize, and there's no way our theory can assign a stable value to this numeral. We could of course perform carries in a non-zippy way that would cause each position to carry only finitely often, but the resulting numeral would have no significance.

We can now state a major organizing principle for the study of carrying in infinite numerals: given a starting numeral, either *no* zippy carry-history leads to stabilization, or *every* zippy carry-history leads to stabilization and moreover leads to the *same* stable numeral.

This will be a consequence of the following more specific dichotomy: in any position k, every zippy carry-history causes the same number of carries in position k (where the number of carries is either a finite non-negative integer or infinity).

Note that if we are in the *tame* case where every position is the site of only finitely many carries, then for all k the eventual value of the kth site is the original value of the kth site, minus the number of carries from the kth site to the $k-1$st site, plus ten times the number of carries from the $k+1$st site to the kth site.

In the rest of the article we will prove three easy results:

Proposition 1: Every zippy carry-history causes the same number of carries in each position.

Proposition 2: If the initial numeral has all its digits less than B for some finite bound B, then every zippy carry-history causes only finitely many carries in each position, so that there's a well-defined stabilization. (Note that when you add together k stable infinite decimals place-by-place without carries, each generalized digit is at most $9k$, in which case you can take $B = 9k$.)

Proposition 3: If the initial numeral doesn't end in an infinite string of zeroes then neither does its stabilization.

Before proving Proposition 1, we need a lemma.

Lemma 1: If there is a carry-history that causes m or more carries to occur in the kth position, then every zippy carry-history must cause m or more carries to occur in the kth position.

Proof of Lemma 1: To avoid notational complexities, I'll sketch the proof in a concrete case that displays all the ideas a fully general proof requires. (This is the proof-by-example method that was fashionable until a few centuries ago. The main advantage of the method is that it is concrete; the main disadvantage is that, carelessly used, it can lead you to believe things that aren't true.)

Suppose that there is a carry-history H' that causes two carries in the 17th position. More specifically, there is a carry history that begins with a carry in the 17th position, followed sometime later by a carry in the 18th position, followed sometime later by a carry in the 17th position. We wish to show that in every zippy carry-history, there must be two carries in the 17th position.

Let H be some zippy carry-history. We start by noticing that since H' begins with a carry in the 17th position, the 17th position starts out as a legal site for carrying in H. Since H is zippy, there must eventually be a carry in the 17th position in H. Once this has happened, the 18th position is ripe for carrying, unless there's already been a carry there. In the former case, the 18th position must eventually have a carry in H (since H is zippy); in the latter case, the 18th position has already had a carry in H. In either case, we see that there must occur a time in H when sites 17 and 18 have both had carries. Once this has happened, the 17th position is ripe for a second carry operation (as is witnessed by H'), unless there has already been a second carry there. In the former case, the 17th position must eventually have a second carry in H (since H is zippy); in the latter case, the 17th position has already had a second carry in H. In either case, we see that there must occur a time in H when site 17 has had at least two carries and site 18 has had at least one.

This style of argument (technically, a proof by mathematical induction) is enough to prove the Lemma.

Proof of Proposition 1: Fix k, and define M as the maximum over all histories of the number of carries at site k (if the maximum doesn't exist, write $M = \infty$). If $M = \infty$, then for all m, there exists a carry-history in which site k has m carries, so by Lemma 1, in every zippy carry-history site k has at least m carries; since this is true for all m, site k has infinitely many carries. If $M < \infty$, then there exists a carry-history in which site k has M carries, so by the Lemma, in every zippy carry-history site k has at least M carries. But by the definition of M, no carry-history causes site k to carry more than M times, so no zippy carry-history causes site k to carry more than M times. Combining the two previous sentences, we see that every zippy carry-history causes site k to carry exactly M times.

To prove Proposition 2, we need another Lemma:

Lemma 2: If a numeral is tame, then any numeral that it dominates on a digit-by-digit basis is tame as well.

Proof: Let $\mathbf{n}$ be a tame numeral, and $\mathbf{n}'$ be another numeral such that for all k, the kth digit of $\mathbf{n}'$ is less than or equal to the kth digit of $\mathbf{n}$. Fix k, and let m be the number of carries in $\mathbf{n}$ in position k arising from any zippy carry-history (this is well-defined by Proposition 1). If there were a carry-history for $\mathbf{n}'$ causing more than m carries in position k, we could lift this carry-history to a carry-history for $\mathbf{n}$ by simply "ignoring the extra counters". But no carry history for $\mathbf{n}$ can cause more than m carries in position k. Hence no carry-history for $\mathbf{n}'$ causes more than m carries in position k. This is true for each position k, so we conclude that $\mathbf{n}'$ is tame.

Proof of Proposition 2: Without loss of generality, assume B is a power of 10; say $B = 10^r$. By Lemma 2, it's enough to show that the numeral $0.10^r,10^r,10^r,\ldots$ is tame. But this is clear, since we can do the carries explicitly, obtaining a stable numeral consisting of a half-infinite string of 0's followed by a half-infinite string of 1's.

Proof of Proposition 3: If only finitely many carries take place in the stabilization of the numeral, then there's nothing to prove, since a finite number of operations can neither create nor destroy an infinite tail of 0's. Now suppose that there are infinitely many carries in every zippy carry-history for $\mathbf{n}$ (but only finitely many in each position). Let k_1 be the first site that undergoes a carry. Site k_1 undergoes only finitely many carries; wait until the last one has taken place. Let k_2 be a site to the right of k_1 that subsequently undergoes a carry. (Such a site must exist, since there are only finitely many sites that aren't to the right of k_1, and each of them undergoes only finitely many carries.) Wait until the last carry at site k_2 has taken place. Let k_3 be a site to the right of k_2 that subsequently undergoes a carry. And so on. Then the final stabilized version of the original numeral must contain a non-zero digit between positions k_1 and k_2 (more precisely, in at least one of the positions $k_1, k_1 + 1, \ldots, k_2 - 1$), and must contain a non-zero digit between positions k_2 and k_3, and must contain a non-zero digit between positions k_3 and k_4, and so on, which rules out the possibility of an infinite tail of 0's.

If this development of the properties of decimal numerals seems like overkill, it should be remarked that once the theoretical apparatus sketched above is in place, it isn't difficult to extend the results to less obvious problems, such as the behavior of multiplication of infinite decimals. (Exercise: Show that if you multiply two positive infinite decimals, the only way in which the product can end in infinitely many zeroes is if all the factors do.)

Of course, the generalized decimal numeral $0.a,b,c,\ldots$ is just the value of the infinite series $a/10 + b/100 + c/1000 + \ldots$, and if we are willing to use the theory of infinite series, we can get a much more direct proof of Theorem 3 using analytical ideas instead of combinatorial ones (where the word "analysis" here refers to the branch of mathematics that grows out of the calculus). For instance, suppose a decimal numeral with infinitely many non-zero digits happens to represent an infinite series whose value is the real number 1, and we want to show that the stabilization of this numeral is $0.999\ldots$ rather than $1.000\ldots$. No finite carrying process could lead cause the units place to change from 0 to 1, because at any given stage in the carrying history, only finitely many counters have moved, and these counters together represent a number strictly less than $1.000\ldots$. Since the units digit can never become 1, the sum must be $0.999\ldots$ rather than $1.000\ldots$.

The advantage of the combinatorial approach taken here is that it doesn't presuppose the theory of the real numbers. Indeed, with the ideas that were presented here (and a few more that weren't), one can *build* the theory of real numbers from the ground up, using infinite decimals rather than more traditional but more arcane constructions such as Dedekind cuts or Cauchy sequences. Most mathematicians assume that constructing the reals via their base ten (or base two) expansions must be an ugly procedure, but if one takes our approach to carries, the proofs aren't all that bad. In fact, this way of building up the real numbers is a variation on the unjustly neglected article "The real numbers as a wreath product" by Faltin, Metropolis, Ross, and Rota (*Advances in Mathematics* **16** (1975), pages 278–304).

When we construct the reals via their decimal expansions, one extra wrinkle that turns up is that we need a rationale for treating $0.9,9,9\ldots$ and $1.0,0,0\ldots$ as the same number; this comes from the fact that there exists a numeral $\mathbf{n}$ such that $0.9,9,9\ldots + \mathbf{n}$ is the same numeral as $1.0,0,0\ldots + \mathbf{n}$. (Indeed, $0.9,9,9\ldots + \mathbf{n}$ and $1.0,0,0,\ldots + \mathbf{n}$ are the same numeral if $\mathbf{n}$ is any numeral that doesn't end in an infinite string of 0's.)

A more serious problem is that once we have agreed to call two numerals $\mathbf{m}$ and $\mathbf{m}'$ the same if there exists a numeral $\mathbf{n}$ such that $\mathbf{m}+\mathbf{n} = \mathbf{m}'+\mathbf{n}$, we need to show that the whole number system doesn't collapse into triviality; that is, we need to show that all numerals don't end up representing the *same* number. To see why this could be an issue, note that $1.0,0,0,\ldots = 0.10,0,0,\ldots = 0.0,100,0,0,\ldots = \ldots$. In a naive sense, this sequence of numerals might appear to be converging to $0.0,0,0,\ldots$ (since converges to zero digit-by-digit). When we build up the real numbers via their decimal expansions, the “hardest” theorem (or rather the one theorem whose proof requires the use of a genuinely analytic idea) turns out to be the proof that $1.000\ldots$ isn't equal to zero!

Acknowledgments: Thanks to Henry Cohn, Lionel Levine, Scott Sheffield, and Peter Shor for helpful conversations.