

Object-Oriented Programming in Smalltalk

COMP 105 Assignment

Due Sunday, December 10, 2017 at 11:59PM

Contents

Overview	2
Setup	2
Interactive debugging with μ Smalltalk	2
Reading comprehension (10 percent)	2
Individual Problem	5
Pair Problems: Bignum arithmetic	6
Big picture of the solution	6
Unit testing bignums in Smalltalk	7
Details of all the problems	8
More advice about testing natural numbers	10
Other hints and guidelines	11
Avoid common mistakes	12
Extra credit	12
What and how to submit: Individual problems	13
What and how to submit: Pair problems	14
How your work will be evaluated	14

Overview

The purpose of this assignment is to help you get acquainted with pure object-oriented programming. Even simple algorithms tend to be implemented by sending lots of messages back and forth among a cluster of cooperating, communicating objects. The assignment is divided into two parts.

- In the first part, you do a few small warmup problems to get you used to pure object-oriented style and to acquaint you with μ Smalltalk's large initial basis.
- In the second part you implement arbitrary-precision arithmetic, also known as *bignums*, in μ Smalltalk. Bignums are useful in their own right, and they show how to use *double dispatch* when it is needed to expose representation selectively.

This assignment is time-consuming. Many students have experience in languages called object-oriented, but very few students have experience with the extensive, pervasive inheritance that characterizes idiomatic Smalltalk programs.

Setup

The μ Smalltalk interpreter is in `/comp/105/bin/usmalltalk`. Many useful μ Smalltalk sources are included in the book's git repository, which you can clone by

```
git clone homework.cs.tufts.edu:/comp/105/build-prove-compare
```

Sources that can be found in the `examples` directory includes copies of predefined classes, collection classes, financial history, and other examples from the textbook.

Interactive debugging with μ Smalltalk

Smalltalk is a small language with a big initial basis. There are lots of predefined classes and methods. Here are two tools that can help you get oriented and debug your code:

- Every *class* understands the messages `protocol` and `localProtocol`. These methods are shown in Figure 10.8 on page 802. They are a quick way to remind yourself what messages an object understands and how the message names are spelled.
- The interpreter can help you debug by emitting a *trace* of message sends and answers to depth `n`. Just enter the definition

```
(val &trace n)
```

at the read-eval-print loop.

Reading comprehension (10 percent)

These problems will help guide you through the reading. We recommend that you complete them before starting the other problems below. You can download the questions¹.

¹./cqs.small.txt

1. Read the first seven pages of chapter 10, which starts on page 777. For each of the parts, pick exactly one of the numbered responses in that part.
 - (a) What is a “receiver”? What is an “argument”?
 1. A receiver is the method that the message dispatches to, and an argument is a value that is bound to a formal parameter in a method body.
 2. A receiver is the object that a message is sent to, and an argument is a value that is bound to a formal parameter in a method body.
 3. A receiver is a message that is sent to some object, and an argument is also a message that is sent to some object.
 - (b) What is the relationship between a *message* and a *protocol*?
 1. A protocol says what messages can be sent
 2. A protocol is an object, and a message is a value sent to that object
 3. A protocol is the same as a class, and a message is something that can be sent to instances of that class
 - (c) What is the difference between *instance* protocol and *class* protocol?
 1. Messages defined for a class protocol can be sent to both instances of the class and the class itself, but messages defined for an instance protocol can be sent to only to the instances
 2. Messages defined for a class protocol can only be sent to the class itself, but messages defined for an instance protocol can be sent to both instances of the class and the class itself
 3. Messages defined for a class protocol can be sent to the class itself, and messages defined for an instance protocol can be sent to only to the instances
 - (d) In the name of a message, what is the significance of a colon or colons?
 1. The colons make the message code go faster; i.e. if a message has a large number of colons the message runs faster
 2. The colons indicate the number of arguments expected; i.e. if a message has n number of colons it expects n actual parameters
 3. The colon indicates that the message is only part of a class protocol; i.e. if a message has more than 0 colons it should only be sent to a class and not to an instance of the class
2. Read Section 10.3.4 on message send and dynamic dispatch. Using the class definitions in that section, message `m1` is sent to an object of class `C`: What method *definitions* are dispatched to in what order? An answer looks something like “method `m5` on class `H`, then method `m12` on class `E`, ...”.
3. Continuing with method dispatch, now study the *class* method `new` defined on class `List` just after page 847. The definition sends message `new` to `super`. (Keep in mind: because `new` is a *class* method, both `super` and `self` stand for the class, not for any instance.)

- (a) When *class* method *new* is executed, what three messages are sent by the method body, in what order?
- (b) What does each message send accomplish?
- (c) If we changed the body so instead of `(new super)` it said `(new self)`, what would happen? Pick one of the numbers below to answer this part.
- 1) The new message would indefinitely be dispatched to `self`'s class, because that's what we're implementing. The program would then fail to terminate.
 - 2) The sentinel would point to the wrong value
 - 3) The sentinel gets initialized correctly; so, there's no difference between `(new super)` and `(new self)` here.
4. Find the implementation of class `Number` starting around page 853. Now study the implementation of class `Fraction` starting around page 857.
- Message `-` (minus) is sent to `Fraction (/ 1 2)` with argument `Fraction (/ 1 3)`. What methods are executed in what order? Which class implements `-`?
- You are getting ready for exercise 39(a).
5. Read about the Magnitude protocol on pages 814 and 815. Study the implementation of class `Magnitude` on page 853. Then study the implementations of classes `Integer` and `SmallInteger` starting on page 855.
- Message `<=` is sent to `small integer 3` with argument `9`. What methods are executed in what order? Which class implements `<=`?
- You are ready for exercise 39(a), and you are ready to implement class `Natural` (exercise 42).
6. In section 10.10.1 on page 900, you will find a short definition of an “abstract class.” What do you think is the purpose of an abstract class? Pick one of the lettered responses below.
- (a) An abstract class' purpose is to hide the implementation of instances from clients so that the creator of the class can change internal details without affecting clients
 - (b) An abstract class' purpose is to define methods that other classes inherit. It's a way to provide default implementations that subclasses could use.
 - (c) An abstract class' purpose is no different from a (non-abstract) class' purpose.
7. One example of an abstract class is class `Number`. Look up its definition starting around page 853. We've copied many of the methods defined on `Number` below. To answer this question, please write down the corresponding letter for each of method that you must implement on the `LargeInteger` class. Some methods you will have to implement may not be listed below.
- (a) `+`
 - (b) `*`
 - (c) `negated`
 - (d) `asFraction`
 - (e) `asFloat`
 - (f) `coerce`
 - (g) `negative`
 - (h) `nonnegative`

- (i) `strictlyPositive`
- (j) `=`
- (k) `<`
- (l) `>`
- (m) `<=`

You are getting ready to implement large integers.

8. Read Section 10.6.5, which starts on page 852.

Of the methods listed in comprehension question 7 (subclass responsibilities of `Number`), list each one that needs to know whether its *argument* (not the receiver) is a large or small integer, and if large, whether it is positive or negative. These methods will require double dispatch.

- (a) List the methods by their corresponding letter. For example, `+` is such a method, so you would at least write (a)
- (b) For each method, list what method it will dispatch to on its argument if the received is a `LargePositiveInteger`. For example, method `+` will dispatch to `addLargePositiveIntegerTo`: so at least write (a) will dispatch to `addLargePositiveIntegerTo`:

You are ready to implement large integers (Exercise 43).

9. Read about coercion in section Section 10.4.6 on page 815. Look at the protocol on page 814.

Explain the roles of the methods `asInteger`, `asFraction`, `asFloat`, and `coerce`. If you are unsure, look at the implementations of these methods on class `Integer` starting on page 855.

You are ready to implement mixed arithmetic, with coercions, in Exercise 44.

Individual Problem

Working on your own, please solve Exercise 39(a) on page 918 of Ramsey. This exercise is a warm-up designed to prepare you for the Bignum problems in the pair portion of the assignment.

39(a). *Interfaces as Abstraction Barriers.* Do Exercise 39(a) on page 918 of Ramsey. When the problem says “Arrange the `Fraction` and `Integer` classes”, the text means to revise one or both of these classes or define a related class. Think about *protocols*, not implementation.

Put your solution in file `frac-and-int.smt`. At minimum, your solution should support addition and subtraction, so include at least one unit test for each of these operations.

Hints:

- In a system with abstract data types, you can’t easily mix integers and fractions; they have different types. But in an object-oriented system with behavioral subtyping, you just have to get one object to “behave like” another—which means implementing its protocol. In some cases, this might include implementing private methods.
- Once you know what protocol to implement, you need to either define a related class or revise some existing class/classes. If you go with the latter, take a look at the description for exercise 44, which appears later in this document, if you haven’t already. The idiom described there is how you should revise any existing classes.

- If you change `Integer`, this change doesn't affect `SmallInteger`, which continues to inherit from the original version of `Integer`. So if you change `Integer`, count on changing `SmallInteger` as well.

Related reading:

- For an overview of the `Magnitude` class and its relationship to numbers, read the first page of Section 10.4.6, which starts on page 815. Also in that section, read about `Integer` and `Fraction`.
- For the implementation of `Integer`, see page 855. For the time being, you can ignore class `SmallInteger`.
- For the implementation of `Fraction`, see page 857. Study the implementation of method `+`, and observe how it relies on the exposure of representation through private methods `num` and `den`.
- If nothing comes to you, try reading about how we get access to multiple representations in the object-oriented way: Section 10.6.5, which starts on page 852. You will need to read this section later anyway.

How big is it? To solve these problems, you shouldn't need to add or change more than 10 lines of code in total. The optimal solution is no more than a few lines long.

Pair Problems: Bignum arithmetic

For these problems, you may work with a partner. Please solve Exercise 42 on page 919, Exercise 43 on page 921, and Exercise 44 on page 922 of Ramsey, and Exercise **T** below. You do not need to implement multiplication or division.

Sometimes you want to do computations that require more precision than you have available in a machine word. Full Scheme, Smalltalk, and Icon all provide “bignums.” These are integer implementations that automatically expand to as much precision as you need. Because of their dynamic-typing discipline, these languages make the transition transparently—you can't easily tell when you're using native machine integers and when you're using bignums. In this portion of the homework, you will reimplement infinite-precision natural numbers in Smalltalk, use these natural numbers to define infinite-precision integers, and then modify the built-in `SmallInteger` class so that when more precision is needed, operations seamlessly roll over to infinite precision.

Big picture of the solution

Smalltalk code sends lots of messages back and forth among lots of methods that are defined on different classes. This model shows both the power of Smalltalk—you get a lot of leverage and code reuse—and the weakness of Smalltalk—every algorithm is smeared out over half a dozen methods defined on different classes, making it hard to debug. But this is the object-oriented programming model that lends power not just to Smalltalk but also to Ruby, Objective-C, Swift, Self, JavaScript, and to a lesser extent, Java and C++. To work effectively in any of these languages, one needs the big picture of which class is doing what. A good starting point is the Smalltalk bignums handout²

²../handouts/bignums.pdf

Unit testing bignums in Smalltalk

Like the other bridge languages, μ Smalltalk has check-expect. But in Smalltalk, checking for equality is often not useful—because many Smalltalk abstractions are mutable, equality is normally defined as *object identity*. For example, two lists compare equal only if they are the same list. If they are different lists containing the same elements, then the = method says they are different.

To test equality of bignums in Smalltalk, we will convert each bignum to a list of digits, which we will compare for similarity using the has:to: *class* method on class Similarity:

```
(class Similarity Object
  ()
  (class-method has:to: (ms ns)
    (locals i n equivalent)
    (set n (size ms))
    (if (≠ n (size ns))
      {false}
      {(set i 0)
       (set equivalent true)
       (while {(and: equivalent {( < i n)}})
         {(ifTrue:ifFalse: (≠ (at: ms i) (at: ns i))
           {(set equivalent false)}
           {(set i (+ i 1))}}})
       equivalent}}))
```

With this class method, we can even compare a list and an array for similarity, successfully. Here's an example of how it works.

```
-> (val ns (new List))
List( )
-> (add: ns 1)
-> (add: ns 2)
-> (add: ns 3)
-> ns
List( 1 2 3 )
-> (= ns ns)
<True>
-> (= ns '( 1 2 3 ))
<False>
-> (has:to: Similarity ns '( 1 2 3 ))
<True>
```

If you want to test the results of *arithmetic*, your check-expect should send message has:to: to class Similarity and will compare the result with true. You can download class Similarity³. If you want to test *comparisons*, your check-expect can simply refer to true and false as usual.

³./similarity.smt

Details of all the problems

42. Implementing arbitrary precision natural numbers. Do Exercise 42 on page 919 of Ramsey. You do not need to implement `*` (multiplication) or `div:` (division).

As the book says, you are free to use a list or an array to represent the sequence of digits. Just be aware that unlike Standard ML, μ Smalltalk has no *immutable* list abstraction. If you want immutable lists, you will have to simulate them yourself. For example, you could try using the `Cons` and `Nil` classes that are already defined, along with the `isNil` message that is defined on every object.

How big is it? My solution for the `Natural` class, using the array representation and the hints in the book, is over 100 lines of μ Smalltalk code.

Related reading:

- There is a detailed implementation guide in the `bignums` handout⁴.
- In a system with abstract data types, binary operations like `+` and `*` are easy: you automatically have access to both representations. In a system with objects, not so! To learn how to get access to multiple representations in the object-oriented way, read Section 10.6.5, which starts on page 852
- Class `Natural` is a subclass of `Magnitude`. Study the `Magnitude` protocol in Section 10.4.6. For information about the implementation of `Magnitude`, which should provide useful ideas about `Natural`, as well as the “subclass responsibilities,” study the implementation of `Magnitude` on page 853.⁵
- For the interface to a Smalltalk array, study the `Collection` protocol in Section 10.4.5, which starts on page 805. You have access to the protocol in Figure 10.10 on page 808, but you are more likely to use the `KeyedCollection` protocol in Figure 10.11 on page 810, especially at `at:` and `at:put:`. Don’t overlook the **Arrays** section on pages 812 and 813, including its description of the `Array` class methods `new:` and `from:`.
- For list construction, which you will need for the `decimal` method, look at the `List` protocol in Section 10.4.5, especially Figure 10.13 on page 812.

43. Implementing arbitrary precision integers. Do Exercise 43 on page 921 of Ramsey. You need not implement the `div:` method. But in order to facilitate testing, **you must implement a decimal method** like the one you implemented on class `Natural`. An implementation is shown in the `bignums` handout⁶.

Because you build large integers on top of `Natural`, you don’t have to think about array or list representations any more. Instead you must focus on dynamic dispatch and on getting information from where it is to where it is needed.

How big is it? My solutions for the large integer classes are 22 lines apiece.

Related reading: This problem is all about dynamic dispatch, including double dispatch. Read Section 10.6.5, which starts on page 852. (You’ll also have a chance to practice double dispatch in recitation.)

44. Modifying `SmallInteger` so operations that overflow roll over to infinite precision. Do Exercise 44 on page 922 of Ramsey.

⁴../handouts/bignums.pdf

⁵Note: an object of class `Natural` is not a `Number` as Smalltalk understands it. In particular, class `Natural` does not support methods `negated` or `reciprocal`.

⁶../handouts/bignums.pdf

You must modify `SmallInteger` *without* editing the source code of the μ Smalltalk interpreter. To do so, you will *redefine* class `SmallInteger` using the idiom on page 922:

```
(class SmallInteger SmallInteger
  ()
  ... new method definitions ...
)
```

This idiom modifies the existing class `SmallInteger`; it can both change existing methods and define new methods. **This code changes the basic arithmetic operations that the system uses internally.** If you have bugs in your code, the system will behave erratically. At this point, you must restart your interpreter and fix your bugs. Then use the idiom again.

How big is it? My modifications to the `SmallInteger` class are about 25 lines.

Related reading: Everything about dispatch and double dispatch still applies. In addition, you need to know how overflow is handled using “exception blocks.”

- Review the presentation of blocks, especially the *parameterless* blocks (written with curly braces) in Section 10.4.3, which starts on page 803.
- Read the description of `at:ifAbsent:` in the keyed-collection protocol in Figure 10.11 on page 810. Now study this expression:

```
(at:ifAbsent: '(0 1 2) 3 {0})
```

This attempts to access element 3 of the array (0 1 2), which is out of bounds because the array only has only 3 elements (remember, the first index in a μ Smalltalk is 0). When given an out of bounds index `at:ifAbsent:` sends `value` to the “exception block” `{0}`, which ultimately answers zero.

- Study the implementation of the `at:` method in chunk 843d, which uses `at:ifAbsent:` with an “exception block” that causes a run-time error if `value` is sent to it.
- Finally, study the overflow-detecting primitive methods in Exercise 44 on page 922, and study the implementation of `addSmallIntegerTo:` in the code chunk immediately below. That is the technique you must emulate.

T. Testing Bignums. You will write 9 tests for bignums:

- 3 tests will test only class `Natural`.
- 3 tests will test the large-integer classes, which are built on top of class `Natural`.
- 3 tests will test mixed arithmetic and comparison involving both small and large integers.

Each test will have the same structure:

1. The test will begin with a **summary characterization** of the test in at most 60 characters, formatted on a line by itself as follows:

```
; Summary: .....
```

The summary must be a simple English phrase that describes the test. Examples might be “Ackermann’s function of (1, 1),” “sequence of powers of 2,” or “combinations of +, *, and - on random numbers.”

2. Code will compute a result of class `Natural`, `LargePositiveInteger`, or `LargeNegativeInteger`. The code may appear in a method, a class method, a block, or wherever else you find convenient.
3. The expected result will be written as a literal array, like one of the following:


```
'( 1 2 3 9 4 9 )
'(- 6 5 5 3 6 )
```
4. The test itself will send the `decimal` message to the computed result, then compare the decimal representation with the literal array. Comparison will use class method `has:to:` on class `Similarity`.
5. Finally, the comparison will appear in a `check-expect`, expecting the result of `has:to:` to be `true`.⁷

Each test must take less than 2 CPU seconds to evaluate.

Here is a complete example:

```
; Summary: 10 to the tenth power, mixed arithmetic
(class Test10Power Object
  ()
  (class-method run: (power)
    [locals n 10-to-the-pred-n 10-to-the-n]
    (set n 0)
    (set 10-to-the-pred-n 0) ; not quite right
    (set 10-to-the-n 1)
    (whileTrue: {(< n power)}
      {(set n (+ n 1))
       (set 10-to-the-pred-n 10-to-the-n)
       (set 10-to-the-n (+ 10-to-the-n 10-to-the-n))
       (set 10-to-the-n (+ 10-to-the-n 10-to-the-n))
       (set 10-to-the-n (+ 10-to-the-n 10-to-the-pred-n))
       (set 10-to-the-n (+ 10-to-the-n 10-to-the-n))}}
     10-to-the-n)
  )
  (check-expect (has:to: Similarity (decimal (run: Test10Power 10))
                '( 1 0 0 0 0 0 0 0 0 0))
    true)
```

Related reading: No special reading is recommended for the testing problem. As long as you understand the examples above, that should be enough.

More advice about testing natural numbers

Although we don't ask you to implement multiplication, for testing purposes it may make your life easier if you have a way to at least multiply by 10. One way to do that with just addition is demonstrated by the code below, which accumulates $10n$ in `p`.

⁷In μ Smalltalk, truth is `true`, not `#t`.

```
(set p n) ; p == n
(set p (+ p p)) ; p == 2n
(set p (+ p p)) ; p == 4n
(set p (+ p n)) ; p == 5n
(set p (+ p p)) ; p == 10n
```

If you have bugs in your addition code, fix those before trying to multiply by 10 like above.

Once you are confident that addition works, you can test subtraction of natural numbers by generating a long random sequence, then subtracting the same sequence in which all digits except the most significant are replaced by zero.

You can create more ambitious tests of subtraction by generating random natural numbers and using the algebraic law $(m + n) - m = n$. You can also check to see that unless n is zero, $m - (m + n)$ causes a run-time error on class `Natural`.

If you want more effective tests of subtraction and so on, compare your results with a working implementation of bignums. The languages Scheme, Icon, and Haskell all provide such implementations. (Be aware that the real Scheme `define` syntax is slightly different from what we use in `uScheme`.) We recommend you use `ghci` on the command line; standard infix syntax works. If you want something more elaborate, use Standard ML of New Jersey (command `sm1`), which has an `IntInf` module that implements bignums.

Other hints and guidelines

Start early. Seamless arithmetic requires in-depth cooperation among about eight different classes (those you write, plus `Magnitude`, `Number`, `Integer`, and `SmallInteger`). This kind of cooperation requires aggressive message passing and inheritance, which you are just learning. There is a handout online⁸ with suggestions about which methods depend on which other methods and in what order to tackle them.

For your reference, Dave Hanson's book⁹ discusses bignums and bignum algorithms at some length. It should be free online to Tufts students. You can think about borrowing code from Hanson's implementation¹⁰ (see also his distribution¹¹). Be aware though that your assignment differs significantly from his code and unless you have read the relevant portions of the book, you may find the code overwhelming.

- In Hanson's code, `XP_add` does add with carry. `XP_sub` does subtract with borrow. `XP_mul` does $z := z + x * y$, which is useful, but is not what we want unless z is zero initially.
- Hanson passes all the lengths explicitly, which would not be idiomatic in `μSmalltalk`.
- Hanson's implementation uses mutation extensively, but the class `Natural` is an immutable type. Your methods must *not* mutate existing natural numbers; you can mutate only a newly allocated number that you are sure has not been seen by any client.

If you do emulate Hanson's code, acknowledge him in your README file.

⁸../handouts/bignums.pdf

⁹../readings/indexbody.html#cii

¹⁰xp.c

¹¹<http://www.cs.princeton.edu/software/cii>

Avoid common mistakes

Below you will find some common mistakes to avoid.

It is common to overlook class methods. They are a good place to put information that doesn't change over the life of your program.

It's not common, but it's a **terrible mistake** to try to implement mixed operations by interrogating an object about its class—a so-called “run-time type test.” Run-time type tests destroy behavioral subtyping. If you find yourself wanting to ask an object what its class is, seek help immediately.

It is surprisingly common for students to submit code for the small problems without ever even having run the code or loaded it into an interpreter. If you run even one test case, you will be ahead of the game.

It is too common to submit bignum code without having tested all combinations of methods and arguments. Your best plan is to write a program, in the language of your choice, that loops over operator and both operands and generates at least one test case for every combination. Because μ Smalltalk is written using S-expressions, you could consider writing this program in μ Scheme—but any language will do.

It is relatively common for students' code to make a false distinction between two flavors of zero. In integer arithmetic, there is only one zero, and it always prints as “0”.

It's surprisingly common to fail to tag the test summary with the prefix Summary:, or to forget it altogether.

Extra credit

Seamless bignum arithmetic is an accomplishment. But it's a long way from industrial. The extra-credit problems explore some ideas you would deploy if you wanted everything on a more solid foundation.

Base variations. A key problem in the representation of integers is the choice of the base b . Today's hardware supports $b = 2$ and sometimes $b = 10$, but when we want bignums, the choice of b is hard to make in the general case:

- If $b = 10$, then converting to decimal representation is trivial, but storing bignums requires lots of memory.
- The larger b is, the less memory is required, and the more efficient everything is.
- If b is a power of 10, converting to decimal is relatively easy and is very efficient. Otherwise, as long as $b \geq 10$, conversion requires only short division.
- If $(b-1) * (b-1)$ fits in a machine word, then you can implement multiplication in high-level languages without difficulty. (Serious implementations pick the largest b such that one digit fits in a machine word, e.g., 2^{64} on modern machines. Unfortunately, to work with such large values of b requires special machine instructions to support “add with carry” and 128-bit multiply, so serious implementations have to be written in assembly language.)
- If b is a power of 2, bit-shift can be very efficient, but conversion to decimal is expensive. Fast bit-shift can be important in cryptographic and communications applications.

If you want signed integers, there are more choices: signed-magnitude and b's-complement. Knuth volume 2¹² is pretty informative about these topics.

For extra credit, try the following variations on your implementation of class `Natural`:

1. Implement the class using an internal base $b = 10$. Measure the time needed to compute the first 50 factorials.
2. Determine the largest possible base that is still a power of 10 and explain your reasoning. Change your class to use that base internally. (If you are both careful and clever, you should be able to change only the class method `base` and not any other code.) Measure the time needed to compute **and print** the first 50 factorials.
3. Determine the largest possible base that is a power of 2 and explain your reasoning. Change your class to use that base internally. Measure the time needed to compute **and print** the first 50 factorials. Does having fewer digits recoup the higher cost of converting to decimal?

Because Smalltalk hides the representation from clients, a well-behaved client won't be affected by a change of base. If we wanted, we could take more serious measurements and pick the most efficient representation.

Remember that the private `decimal` method must return a list of **decimal** digits, even if base 10 is not what is used in the representation. Suppress leading zeroes unless the value of `Natural` is itself zero.

Write up your arguments and your measurements in your README file.

Multiplication. Add multiplication to your implementation of class `Natural`.

Follow the protocol described in 42, and do use the built-in multiplication; don't simply multiply by repeated addition.

Long division. Implement long division for `Natural` and for large integers. If this changes your argument for the largest possible base, explain how. This article¹³ by Per Brinch Hansen describes how to implement long division.

Mixed Comparisons. Make sure comparisons work, even with mixed kinds of integers. So for example, make sure comparisons such as `(< 5 (* 1000000 1000000))` produce sensible answers.

Space costs. Instrument your `Natural` class to keep track of the size of numbers, and measure the space cost of the different bases. Estimate the difference in garbage-collection overhead for computing with the different bases, given a fixed-size heap.

Pi (hard). Use a power series to compute the first 100 digits of pi (the ratio of a circle's circumference to its diameter). Be sure to cite your sources for the proper series approximation and its convergence properties. *Hint: I vaguely remember that there's a faster convergence for pi over 4. Check with a numerical analyst.*

What and how to submit: Individual problems

Submit these files:

¹²../readings/indexbody.html#knuth2

¹³../readings/indexbody.html#division

- A README file containing
 - The names of the people with whom you collaborated
 - The numbers of the problems that you solved
 - The number of hours you worked on the assignment.
- A file `cqs.small.txt` containing your answers to the reading-comprehension questions
- A file `frac-and-int.smt` showing whatever definitions you used to do Exercise 39(a). It probably includes new definitions (or redefinitions) of one or more of these classes: `Fraction`, `Integer`, and `SmallInteger`. And it most definitely includes at least four unit tests.

Please identify your solutions using *conspicuous* comments, e.g.,

```

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;;;
;;; Solution to Exercise XXX
(class Array ...
)

```

As soon as you have the files listed above, run `submit105-small-solo` to submit a preliminary version of your work. Keep submitting until your work is complete; we grade only the last submission.

What and how to submit: Pair problems

Submit these files:

- A README file containing
 - A description of how you tested your bignum code
 - The names of the people with whom you collaborated
 - The numbers of the problems that you solved (including any extra credit)
 - Narrative and measurements to accompany your extra-credit answers, if any
- A file `bignum.smt` showing your solutions to Exercises 42, 43, and 44. This file **must** work with an *unmodified* `usmalltalk` interpreter. Therefore if you use results from Exercise 39(a), or any other problem, you will need to duplicate those modifications in `bignum.smt`.
- A file `bigtests.smt` containing your solution to Exercise T.

As soon as you have the files listed above, run `submit105-small-pair` to submit a preliminary version of your work. Keep submitting until your work is complete; we grade only the last submission.

How your work will be evaluated

All our usual expectations for **form**, **naming**, and **documentation** apply. But in this assignment we will focus on **clarity** and **structure**. To start, we want to be able to understand your code.

	Exemplary	Satisfactory	Must Improve
Clarity	<ul style="list-style-type: none"> • Course staff see no more code than is needed to solve the problem. • Course staff see how the structure of the code follows from the structure of the problem. 	<ul style="list-style-type: none"> • Course staff see somewhat more code than is needed to solve the problem. • Course staff can relate the structure of the code to the structure of the problem, but there are parts they don't understand. 	<ul style="list-style-type: none"> • Course staff see roughly twice as much code as is needed to solve the problem. • Course staff cannot follow the code and relate its structure to the structure of the problem.

Structurally, your code should hide information like the base of natural numbers, and it should use proper method dispatch, not bogus techniques like run-time type checking.

	Exemplary	Satisfactory	Must Improve
Structure	<ul style="list-style-type: none"> • The base used for natural numbers appears in exactly one place, and all code that depends on it consults that place. • <i>Or</i>, the base used for natural numbers appears in exactly one place, and code that depends on either consults that place or assumes that the base is some power of 10 • No matter how many bits are used to represent a machine integer, overflow is detected by using appropriate primitive methods, not by comparing against particular integers. • Code uses method dispatch instead of conditionals. • Mixed operations on different classes of numbers are implemented using double dispatch. • <i>Or</i>, mixed operations on different classes of numbers are implemented by arranging for the classes to share a common protocol. • <i>Or</i>, mixed operations on different classes of numbers are implemented by arranging for unconditional coercions. • Code deals with exceptional or unusual conditions by passing a suitable <code>exnBlock</code> or other block. • Code achieves new functionality by reusing existing methods, e.g., by sending messages to <code>super</code>. • <i>Or</i>, code achieves new functionality by adding new methods to old classes to respond to an existing protocol. • An object's behavior is controlled by dispatching (or double dispatching) to 	<ul style="list-style-type: none"> • The base used for natural numbers appears in exactly one place, but code that depends on it knows what it is, and that code will break if the base is changed in any way. • Overflow is detected only by assuming the number of bits used to represent a machine integer, but the number of bits is explicit in the code. • Code contains one avoidable conditional. • Mixed operations on different classes of integers involve explicit conditionals. • Code protects itself against exceptional or unusual conditions by using Booleans. • Code contains methods that appear to have been copied and modified. • An object's behavior is influenced by interrogating it to learn something about its class. 	<ul style="list-style-type: none"> • The base used for natural numbers appears in multiple places. • Overflow is detected only by assuming the number of bits used to represent a machine integer, and the number of bits is <i>implicit</i> in the value of some frightening decimal literal. • Code contains more than one avoidable conditional. • Mixed operations on different classes of integers are implemented by interrogating objects about their classes. • Code copies methods instead of arranging to invoke the originals. • Code contains case analysis or a conditional that depends on the class of an object.

Exemplary	Satisfactory	Must Improve
------------------	---------------------	---------------------