

COMP 105 Assignment: Smalltalk

Due **Friday**, April 15, at 5:59PM.

Extended to **Sunday**, April 17, at 11:59PM.

The purpose of this assignment is to help you get acquainted with pure object-oriented programming. 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 uSmalltalk's large initial basis.
- In the second part, you implement *bignums* in uSmalltalk. Bignums are useful in their own right, and they illustrate the important object-oriented technique of *double dispatch*.

Setup

The uSmalltalk interpreter for the course is in `/comp/105/bin/usmalltalk`. To help with debugging, you can instruct the interpreter to emit a trace of message sends and answers to depth `n` by entering:

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

Useful sources are in the git repository, which you can clone by

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

The `examples` directory in the repository includes copies of the initial basis, collection classes, financial history, and other examples from the textbook.

Individual Problems

Working on your own, please solve Exercise 9 on page 674 and Exercise 29a on page 679 of Ramsey. These exercises are warm-ups designed to prepare you for the Bignum problems in the pair portion of the assignment.

Problem Details

9. Collection Protocols. Do Exercise 9 on page 674 of Ramsey.

29a. Interfaces as Abstraction Barriers. Do Exercise 29a on page 679 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.

To solve these problems, you shouldn't need to add or change more than 10 lines of code in total.

Pair Problems

For these problems, you may work with a partner. Please solve Exercise 31 on page 680, Exercise 32 on page 682, and Exercise 33 on page 683 of Ramsey, and Exercise **T below**. You do not need to implement long division.

Problem Details

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 implement infinite precision natural numbers, use these natural numbers to define infinite precision integers, and then modify the built-in `SmallInteger` class so that operations seamlessly roll over to infinite precision when overflows

COMP 105 Assignment: Smalltalk

occur.

31. Implementing arbitrary precision natural numbers. Do Exercise 31 on page 680 of Ramsey. The text of the exercise contains significant information about how to implement the required arbitrary-precision arithmetic. We recommend you choose base $b = 10$ to represent your bignums. The course solution for the `Natural` class is over 100 lines of uSmalltalk code.

32. Implementing arbitrary precision integers. Do Exercise 32 on page 682 of Ramsey. The course solution for the large integer classes are 22 lines apiece.

33. Modifying `SmallInteger` so that operations that overflow roll over to infinite precision. Do Exercise 33 on page 683 of Ramsey. The idiom

```
(class SmallInteger SmallInteger
  ...modifications...
)
```

modifies the existing class `SmallInteger` according to the code in `...modifications...`. You should use this idiom to make your changes. Note that once you enter this code into the interpreter, **you have changed the 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 should restart your interpreter and fix your bugs before re-modifying the `SmallInteger` class. The modifications to the `SmallInteger` class in the course solution are about 25 lines.

T. Testing Bignums. Write a test for bignums and explain in comments what way the submitted test is superior to the factorial test. The question you should ask yourself devising your test is: *What will constitute an adequate test of bignums?*

Your test file should be formatted as follows:

1. It should begin with whatever definitions you need to run the test.
2. It should contain a **summary characterization** of the test in at most 60~characters, formatted on a line by itself as follows:

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

Your summary should be a simple English phrase that describes your test. Examples might be ``Ackermann's function of (1, 1),'' ``sequence of powers of 2,'' or ``combinations of +, *, and - on random numbers.''

3. It should define a class `Test105` with a class method `run` that actually runs the test.
4. It should end with **comments** that contain just the output from sending the `run` method to the `Test105` class, and nothing else. If the output is a single line, write a one-line comment. If the output takes multiple lines, put each line of output in a comment on its own line.
5. The expression `(run Test105)` **must take less than 2 CPU seconds to evaluate**.

Your test file should *not* include any code that *implements* bignums.

Here is a trivial example:

<example bigtests.smt>=

```
; Comments explaining why test is good.
; Summary: 10 to the tenth power
(class Test105 Object
  ()
  (class-method run ()
    (locals n 10-to-the-n)
    (set n 0)
    (set 10-to-the-n 1)
    (whileTrue: [(< n 10)]
      [(set n (+ n 1))
       (set 10-to-the-n (* 10 10-to-the-n))])
    10-to-the-n)
```

```
)
; 10000000000
```

If this test is run in an unmodified interpreter, it breaks with an arithmetic overflow and a stack trace.

Extra-credit: 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 it requires (possibly long) 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 $a[i]$ is guaranteed to fit 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 2](#) 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. Make an argument for the largest possible base that is still a power of 10. 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 the first 50 factorials. Note both your measurements and your argument in your README file.

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.

More Extra-credit problems

Division. Implement long division for `Natural` and for large integers. If this changes your argument for the largest possible base, explain how. This [article](#) by Per Brinch Hansen describes how to implement long division.

Largest base. Change the base to the largest reasonable base, not necessarily a power of 10. You will have to re-implement `decimal` using long division. *Measure* the time needed to compute *and print* the first 50 factorials. Does the smaller number of digits recoup the higher cost of converting to decimal?

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.

What to submit: Individual Problems

You should submit three files:

- 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 `finhist.smt` showing your solution to Exercise 9.
- A file `basis.smt` showing whatever changes you had to make to the initial basis to do Exercise 29a.

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

```

////////////////////////////////////
;;;
;;; Solution to Exercise Exercise 9
(class Array ...
)

```

How to submit: Individual Problems

When you are ready, run `submit105-small-solo` to submit your work. Note you can run this script multiple times; we will grade the last submission.

What to submit: Pair Problems

You should submit three 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)
 - ◆ the number of hours you worked on the assignment.
- A file `bignum.smt` showing your solutions to Exercises 31, 32, and 33. This file **must** work with an *unmodified* `usmalltalk` interpreter. Therefore if you use results from Exercise 29a, or any other problem, you will need to duplicate those modifications in `bignum.smt`.
- A file `bigtests.smt` containing your solution to Exercise T.

How to submit: Pair Problems

When you are ready, run `submit105-small-pair` to submit your work. Note you can run this script multiple times; we will grade the last submission.

Hints and Guidelines

- Don't overlook the `protocol` and `localProtocol` methods, which are defined *on every class*, as shown in Figure 12.6 on page 577. These class methods will return the methods that are defined in a class.
- **Start early.** The bignum problems are hard; they involve writing a lot of complicated, interrelated methods using language idioms that you are just learning. There is a [handout online](#) with suggestions about which methods depend on which other methods and in what order to tackle them.
- For your reference, [Dave Hanson's book](#) 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 implementation](#) (see also his [distribution](#)). 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 the 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.

COMP 105 Assignment: Smalltalk

- ◆ Hanson passes all the lengths explicitly, which would not be idiomatic in uSmalltalk.
- ◆ 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.

Testing bignum arithmetic

To help you test your work, here is code that computes and prints factorials:

<fact.smt>=

```
(define factorial (n)
  (if (strictlyPositive n)
      [(* n (value factorial (- n 1)))]
      [1]))

(class Factorial Object
  ()
  (classMethod printUpto: (limit) (locals n nfac)
    (begin
      (set n 1)
      (set nfac 1)
      (while [(<= n limit)]
        [(print n) (print #!) (print space) (print #=) (print space) (println nfac)
         (set n (+ n 1))
         (set nfac (* n nfac))])))
```

You might find it useful to test your implementation with the following table of factorials:

```
1! = 1
2! = 2
3! = 6
4! = 24
5! = 120
6! = 720
7! = 5040
8! = 40320
9! = 362880
10! = 3628800
11! = 39916800
12! = 479001600
13! = 6227020800
14! = 87178291200
15! = 1307674368000
16! = 20922789888000
17! = 355687428096000
18! = 6402373705728000
19! = 121645100408832000
20! = 2432902008176640000
21! = 51090942171709440000
22! = 112400072777607680000
23! = 25852016738884976640000
24! = 620448401733239439360000
25! = 15511210043330985984000000
```

Be warned that **this test by itself is inadequate**. You will want other tests. Here is some advice on testing:

- Try testing your class `Natural` by generating a long, random string of digits, then computing the corresponding number using a combination of addition and multiplication by 10.
- If you don't have multiplication working yet, you can use the following sequence to multiply by 10:

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

COMP 105 Assignment: Smalltalk

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

This idea will test only your addition; if you have bugs there, fix them before you go on.

- You can write, in `μSmalltalk` itself, a method that uses the techniques above to convert a sequenceable collection of decimal digits into a natural number.
- For easy comparison, you can even define a *subclass* of `Array` or `List` that prints just the digits, with no spaces.
- 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.
- It is harder to test multiplication, but you can at least use repeated addition to test multiplication by *small* values. The `timesRepeat` method is defined on any integer.
- You can also easily test multiplication by large powers of 10.
- You can use similar techniques to test large integers.

If you want more effective tests of multiplication 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 `sml`), which has an `IntInf` module that implements bignums.

Avoid common mistakes

Here are 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 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 simple shell script that has three nested loops (over operator and both operands) and can emit test code either for your own interpreter or for an interpreter that has bignums built in, like `ghci`.
- It's a common mistake to submit a file `bigtests.smt` that has junk at the end or is otherwise badly formatted and thereby earns little credit.
- 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`.

How your work will be evaluated

All our usual expectations for **form**, **naming**, and **documentation** apply. But in this assignment we will focus on **structure** and **clarity**.

	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	<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	<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.

COMP 105 Assignment: Smalltalk

	<p>power of 10</p> <ul style="list-style-type: none"> • 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 an appropriate method of its class. 	<p>the number of bits is explicit in the code.</p> <ul style="list-style-type: none"> • 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"> • 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.
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 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.
Documentation	<ul style="list-style-type: none"> • Private methods are documented with contracts. 		<ul style="list-style-type: none"> • Private methods are neither documented nor consistent with the <code>bignums</code> handout.

COMP 105 Assignment: Smalltalk

	<ul style="list-style-type: none">• <i>Or</i>, private methods use exactly the contracts suggested in the bignums handout.
Correctness	<ul style="list-style-type: none">• Code that works with collections works with <i>any</i> <code>Collection</code> class.

	<ul style="list-style-type: none">• Code that is supposed to work with all collections works only with some subclasses of collections.

Back to the [class home page](#)