

CSCI-1200 Data Structures — Spring 2013

Lecture 7 — Templated Classes & Vector Implementation

Review from Lectures 5 & 6

- Arrays and pointers
- Different types of memory (“automatic”, static, dynamic)
- Dynamic allocation of arrays

1.1 Today’s Lecture

- Designing our own container classes:
 - Mimic the interface of standard library (STL) containers
 - Study the design of memory management.
 - Move toward eventually designing our own, more sophisticated classes.
- Vector implementation
- Templated classes
- Copy constructors, assignment operators, and destructors

Optional Reading: Ford & Topp, Sections 5.3-5.5; Koenig & Moo Chapter 11

1.2 Vector Public Interface

- In creating our own version of the STL vector class, we will start by considering the public interface:

```
public:
// MEMBER FUNCTIONS AND OTHER OPERATORS
T& operator[] (size_type i);
const T& operator[] (size_type i) const;
void push_back(const T& t);
void resize(size_type n, const T& fill_in_value = T());
void clear();
bool empty() const;
size_type size() const;
```

- To implement our own generic (a.k.a. templated) vector class, we will implement all of these operations, manipulate the underlying representation, and discuss memory management.

1.3 Templated Class Declarations and Member Function Definitions

- In terms of the layout of the code in `vec.h` (pages 5 & 6 of the handout), the biggest difference is that this is a *templated class*. The keyword `template` and the template type name must appear before the class declaration:

```
template <class T> class Vec
```

- Within the class declaration, `T` is used as a type and all member functions are said to be “templated over type `T`”. In the actual text of the code files, templated member functions are often defined (written) *inside the class declaration*.

- The templated functions defined outside the template class declaration must be preceded by the phrase: `template <class T>` and then when `Vec` is referred to it must be as `Vec<T>`. For example, for member function `create` (two versions), we write:

```
template <class T> void Vec<T>::create(...
```

1.4 Syntax and Compilation

- Templated classes and templated member functions are not created/compiled/instantiated until they are needed. Compilation of the class declaration is triggered by a line of the form: `Vec<int> v1;` with `int` replacing `T`. This also compiles the default constructor for `Vec<int>` because it is used here. Other member functions are not compiled unless they are used.
- When a different type is used with `Vec`, for example in the declaration: `Vec<double> z;` the template class declaration is compiled again, this time with `double` replacing `T` instead of `int`. Again, however, only the member functions used are compiled.
- This is very different from ordinary classes, which are usually compiled separately and all functions are compiled regardless of whether or not they are needed.
- The templated class declaration and the code for all used member functions must be provided where they are used. As a result, member functions definitions are often included within the class declaration or defined outside of the class declaration but still in the `.h` file. If member function definitions are placed in a separate `.cpp` file, this file must be `#include-d`, just like the `.h` file, because the compiler needs to see it in order to generate code. (Normally we don't `#include .cpp` files!) See also diagram on page 7 of this handout.

Note: Including function definitions in the `.h` for ordinary non-templated classes may lead to compilation errors about functions being “multiply defined”. Some of you have already seen these errors.

1.5 Member Variables

Now, looking inside the `Vec<T>` class at the member variables:

- `m_data` is a pointer to the start of the array (after it has been allocated). Recall the close relationship between pointers and arrays.
- `m_size` indicates the number of locations currently in use in the vector. This is exactly what the `size()` member function should return,
- `m_alloc` is the total number of slots in the dynamically allocated block of memory.

Drawing pictures, which we will do in class, will help clarify this, especially the distinction between `m_size` and `m_alloc`.

1.6 Typedefs

- Several types are created through `typedef` statements in the first `public` area of `Vec`. Once created the names are used as ordinary type names. For example `Vec<int>::size_type` is the return type of the `size()` function, defined here as an `unsigned int`.

1.7 operator[]

- Access to the individual locations of a `Vec` is provided through `operator[]`. Syntactically, use of this operator is translated by the compiler into a call to a function called `operator[]`. For example, if `v` is a `Vec<int>`, then:

```
v[i] = 5;
```

translates into:

```
v.operator[](i) = 5;
```

- In most classes there are two versions of `operator[]`:
 - A non-const version returns a reference to `m_data[i]`. This is applied to non-const `Vec` objects.
 - A const version is the one called for `const Vec` objects. This also returns `m_data[i]`, but as a const reference, so it can not be modified.

1.8 Default Versions of Assignment Operator and Copy Constructor Are Dangerous!

- Before we write the copy constructor and the assignment operator, we consider what would happen if we didn't write them.
- C++ compilers provide default versions of these if they are not provided. These defaults just copy the values of the member variables, one-by-one. For example, the default copy constructor would look like this:

```
template <class T>
Vec<T> :: Vec(const Vec<T>& v)
    : m_data(v.m_data), m_size(v.m_size), m_alloc(v.m_alloc)
    {}
```

In other words, it would construct each member variable from the corresponding member variable of `v`. This is dangerous and incorrect behavior for the `Vec` class. We don't want to just copy the `m_data` pointer. We really want to create a copy of the entire array! Let's look at this more closely...

1.9 Exercise

Suppose we used the default version of the assignment operator and copy constructor in our `Vec<T>` class. What would be the output of the following program? Assume all of the operations **except** the copy constructor behave as they would with a `std::vector<double>`.

```
Vec<double> v(4, 0.0);
v[0] = 13.1; v[2] = 3.14;
Vec<double> u(v);
u[2] = 6.5;
u[3] = -4.8;
for (unsigned int i=0; i<4; ++i)
    cout << u[i] << " " << v[i] << endl;
```

Explain why this happens by drawing a picture of the memory of both `u` and `v`.

1.10 Classes With Dynamically Allocated Memory

- For `Vec` (and other classes with dynamically-allocated memory) to work correctly, each object must do its own dynamic memory allocation and deallocation. We must be careful to keep the memory of each object instance separate from all others.
- All dynamically-allocated memory for an object should be released when the object is finished with it or when the object itself goes out of scope (through what's called a *destructor*).
- To prevent the creation and use of default versions of these operations, we must write our own:

- *Copy constructor*
- *Assignment operator*
- *Destructor*

1.11 Copy Constructor

- This constructor must dynamically allocate any memory needed for the object being constructed, copy the contents of the memory of the passed object to this new memory, and set the values of the various member variables appropriately.
- **Exercise:** In our `Vec` class, the actual copying is done in a private member function called `copy`. Write the private member function `copy`.

1.12 Aside (1): the “this” pointer

- All class objects have a special pointer defined called `this` which simply points to the current class object, and it may not be changed.
- The expression `*this` is a reference to the class object.
- The `this` pointer is used in several ways:
 - Make it clear when member variables of the current object are being used.
 - Check to see when an assignment is self-referencing.
 - Return a reference to the current object.

1.13 Aside (2): Assignment operators, generally speaking

- Assignment operators of the form: `v1 = v2;`
are translated by the compiler as: `v1.operator=(v2);`
- Cascaded assignment operators of the form: `v1 = v2 = v3;`
are translated by the compiler as: `v1.operator=(v2.operator=(v3));`
- Therefore, the value of the assignment operator (`v2 = v3`) must be suitable for input to a second assignment operator. This in turn means the result of an assignment operator ought to be a reference to an object.

1.14 Assignment operator for Vec

- The implementation of an assignment operator usually takes on the same form for every class:
 - Do no real work if there is a self-assignment.
 - Otherwise, destroy the contents of the current object then copy the passed object, just as done by the copy constructor.
 - Return a reference to the (copied) current object, using the `this` pointer.

1.15 Destructor (the “constructor with a tilde”)

- The destructor is called implicitly when an automatically-allocated object goes *out of scope* or a dynamically-allocated object is *deleted*. It can never be called explicitly!
- The destructor is responsible for deleting the dynamic memory “owned” by the class.
- The syntax of the function definition is a bit weird. The `~` has been used as a logic negation in other contexts.

1.16 Increasing the Size of the Vec

- `push_back(T const& x)` adds to the end of the array, increasing `m_size` by one T location. But what if the allocated array is full (`m_size == m_alloc`)?
 1. Allocate a new, larger array. The best strategy is generally to double the size of the current array. Why?
 2. If the array size was originally 0, doubling does nothing. We must be sure that the resulting size is at least 1.
 3. Then we need to copy the contents of the current array.
 4. Finally, we must delete current array, make the `m_data` pointer point to the start of the new array, and adjust the `m_size` and `m_alloc` variables appropriately.
- Only when we are sure there is enough room in the array should we actually add the new object to the back of the array.

1.17 Exercises

- Finish the definition of `Vec::push_back`.
- Write the `Vec::resize` function.

1.18 Vec Declaration & Implementation (vec.h)

```
#ifndef Vec_h_
#define Vec_h_
// Simple implementation of the vector class, revised from Koenig and Moo. This
// class is implemented using a dynamically allocated array (of templated type T).
// We ensure that that m_size is always <= m_alloc and when a push_back or resize
// call would violate this condition, the data is copied to a larger array.

template <class T> class Vec {

public:
    // TYPEDEFS
    typedef unsigned int size_type;

    // CONSTRUCTORS, ASSIGNMENT OPERATOR, & DESTRUCTOR
    Vec() { this->create(); }
    Vec(size_type n, const T& t = T()) { this->create(n, t); }
    Vec(const Vec& v) { copy(v); }
    Vec& operator=(const Vec& v);
    ~Vec() { delete [] m_data; }

    // MEMBER FUNCTIONS AND OTHER OPERATORS
    T& operator[] (size_type i) { return m_data[i]; }
    const T& operator[] (size_type i) const { return m_data[i]; }
    void push_back(const T& t);
    void resize(size_type n, const T& fill_in_value = T());
    void clear() { delete [] m_data; create(); }
    bool empty() const { return m_size == 0; }
    size_type size() const { return m_size; }

private:
    // PRIVATE MEMBER FUNCTIONS
    void create();
    void create(size_type n, const T& val);
    void copy(const Vec<T>& v);

    // REPRESENTATION
    T* m_data; // Pointer to first location in the allocated array
    size_type m_size; // Number of elements stored in the vector
    size_type m_alloc; // Number of array locations allocated, m_size <= m_alloc
};

// Create an empty vector (null pointers everywhere).
template <class T> void Vec<T>::create() {
    m_data = NULL;
    m_size = m_alloc = 0; // No memory allocated yet
}

// Create a vector with size n, each location having the given value
template <class T> void Vec<T>::create(size_type n, const T& val) {
    m_data = new T[n];
    m_size = m_alloc = n;
    for (T* p = m_data; p != m_data + m_size; ++p)
        *p = val;
}

// Assign one vector to another, avoiding duplicate copying.
template <class T> Vec<T>& Vec<T>::operator=(const Vec<T>& v) {
    if (this != &v) {
        delete [] m_data;
        this -> copy(v);
    }
    return *this;
}
}
```

```

// Create the vector as a copy of the given vector.
template <class T> void Vec<T>::copy(const Vec<T>& v) {

}

// Add an element to the end, resize if necessary.
template <class T> void Vec<T>::push_back(const T& val) {
    if (m_size == m_alloc) {
        // Allocate a larger array, and copy the old values

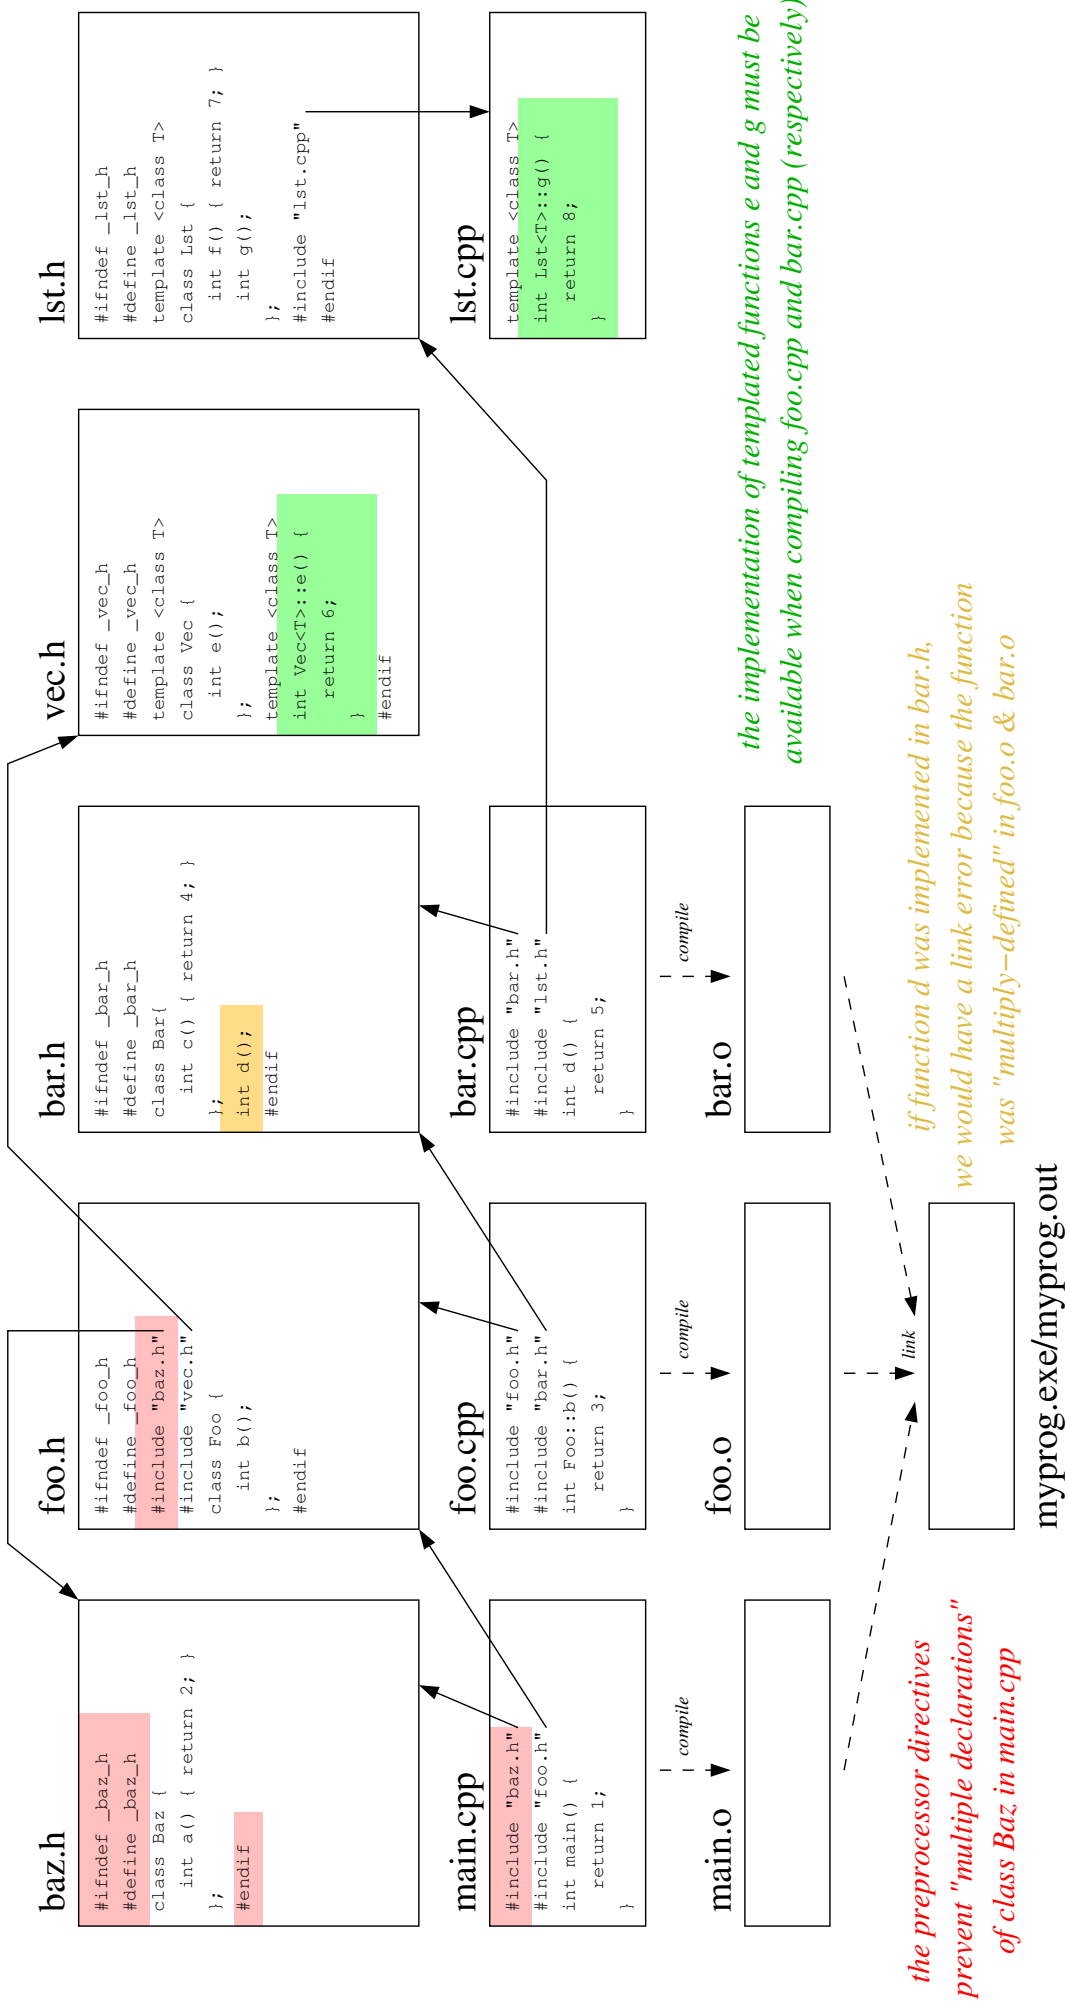
    }
    // Add the value at the last location and increment the bound
    m_data[m_size] = val;
    ++ m_size;
}

// If n is less than or equal to the current size, just change the size. If n is
// greater than the current size, the new slots must be filled in with the given value.
// Re-allocation should occur only if necessary. push_back should not be used.
template <class T> void Vec<T>::resize(size_type n, const T& fill_in_value) {

}

#endif

```



1.19 Valgrind Memory Checker

We recommend the memory debugging tool “Valgrind”, <http://valgrind.org/>, which is available for Linux, FreeBSD, and MacOSX. (Unfortunately, Valgrind *is not* available for Windows or Cygwin. But you can use Dr. Memory, which is described in the next section.) Valgrind’s Memcheck tool can detect the following problems:

- Use of uninitialized memory
- Reading/writing memory after it has been free’d
- Reading/writing off the end of malloc’d blocks
- Reading/writing inappropriate areas on the stack
- Memory leaks - where pointers to malloc’d blocks are lost forever
- Mismatched use of malloc/new/new [] vs free/delete/delete []
- Overlapping src and dst pointers in memcpy() and related functions

For HW3, the homework submission server is configured to run your code with Valgrind to search for memory problems and present the output with the submission results.

1.20 Sample Buggy Program

Can you see the errors in this program?

```
1 #include <iostream>
2
3 int main() {
4
5     int *p = new int;
6     if (*p != 10) std::cout << "hi" << std::endl;
7
8     int *a = new int[3];
9     a[3] = 12;
10    delete a;
11
12 }
```

1.21 Using Valgrind

To use Valgrind, you should compile your program with debugging options (-g) enabled:

```
g++ -g valgrind_test.cpp -o a.out
```

Then, you run your program as usual, with the addition of `valgrind` and the memory leak checking options to the front of the command line.

```
valgrind --leak-check=yes --show-reachable=yes a.out
```

And this is what Valgrind has to say about the errors:

```
==85391== Memcheck, a memory error detector for x86-linux.
==85391== Copyright (C) 2002-2004, and GNU GPL'd, by Julian Seward.
==85391== Using valgrind-2.1.0.352, a program supervision framework for x86-linux.
==85391== Copyright (C) 2000-2004, and GNU GPL'd, by Julian Seward.
==85391== Estimated CPU clock rate is 1694 MHz
==85391== For more details, rerun with: -v
==85391==
==85391== Conditional jump or move depends on uninitialised value(s)
==85391==    at 0x8048A2A: main (valgrind_test.cpp:6)
hi
==85391==
==85391== Invalid write of size 4
==85391==    at 0x8048A68: main (valgrind_test.cpp:9)
==85391==    Address 0x3C231094 is 0 bytes after a block of size 12 alloc'd
```



```

==85391== at 0x3C03847B: operator new[](unsigned) (in /software/valgrind-2.1.0.352-0/pkg/lib/valgrind/vgpreload_memch
==85391== by 0x8048A5B: main (valgrind_test.cpp:8)
==85391==
==85391== Mismatched free() / delete / delete []
==85391== at 0x3C0386F7: operator delete(void*) (in /software/valgrind-2.1.0.352-0/pkg/lib/valgrind/vgpreload_memch
==85391== by 0x8048A78: main (valgrind_test.cpp:10)
==85391== Address 0x3C231088 is 0 bytes inside a block of size 12 alloc'd
==85391== at 0x3C03847B: operator new[](unsigned) (in /software/valgrind-2.1.0.352-0/pkg/lib/valgrind/vgpreload_memch
==85391== by 0x8048A5B: main (valgrind_test.cpp:8)
==85391==
==85391== ERROR SUMMARY: 3 errors from 3 contexts (suppressed: 0 from 0)
==85391== malloc/free: in use at exit: 4100 bytes in 2 blocks.
==85391== malloc/free: 3 allocs, 1 frees, 4112 bytes allocated.
==85391== For counts of detected errors, rerun with: -v
==85391== searching for pointers to 2 not-freed blocks.
==85391== checked 2398936 bytes.
==85391==
==85391==
==85391== 4 bytes in 1 blocks are definitely lost in loss record 1 of 2
==85391== at 0x3C038253: operator new(unsigned) (in /software/valgrind-2.1.0.352-0/pkg/lib/valgrind/vgpreload_memch
==85391== by 0x8048A1D: main (valgrind_test.cpp:5)
==85391==
==85391== LEAK SUMMARY:
==85391== definitely lost: 4 bytes in 1 blocks.
==85391== possibly lost: 0 bytes in 0 blocks.
==85391== still reachable: 0 bytes in 0 blocks.
==85391== suppressed: 4096 bytes in 1 blocks.

```

And here's what it looks like after fixing those bugs:

```

==85444== Memcheck, a memory error detector for x86-linux.
==85444== Copyright (C) 2002-2004, and GNU GPL'd, by Julian Seward.
==85444== Using valgrind-2.1.0.352, a program supervision framework for x86-linux.
==85444== Copyright (C) 2000-2004, and GNU GPL'd, by Julian Seward.
==85444== Estimated CPU clock rate is 1696 MHz
==85444== For more details, rerun with: -v
==85444==
hi
==85444==
==85444== ERROR SUMMARY: 0 errors from 0 contexts (suppressed: 0 from 0)
==85444== malloc/free: in use at exit: 4096 bytes in 1 blocks.
==85444== malloc/free: 3 allocs, 2 frees, 4112 bytes allocated.
==85444== For counts of detected errors, rerun with: -v
==85444== searching for pointers to 1 not-freed blocks.
==85444== checked 2398936 bytes.
==85444==
==85444== LEAK SUMMARY:
==85444== definitely lost: 0 bytes in 0 blocks.
==85444== possibly lost: 0 bytes in 0 blocks.
==85444== still reachable: 0 bytes in 0 blocks.
==85444== suppressed: 4096 bytes in 1 blocks.

```

Note that because the STL string class uses its own allocator, there may be a warning about memory that is “still reachable” even though you’ve deleted all your dynamically allocated memory. The newest version of Valgrind automatically suppresses this specific error, so you will see this listed as a “suppressed leak”.

1.22 Dr. Memory

On Windows, we recommend the memory debugging tool, which detects the same classes of errors as Valgrind. “Dr. Memory”. Dr. Memory is also available on Linux.

<http://dynamorio.org/drmemory.html>

Dr. Memory will work with executables compiled with either the Visual Studio compiler or the MinGW g++ compiler (which can be run from the Cygwin shell). *Unfortunately, because the Cygwin g++ compiler automatically adds the Cygwin libraries to your executable, you cannot use Dr. Memory with executables compiled with the Cygwin g++ compiler.* Detailed instructions for Dr. Memory installation and various compilation instructions are available here:

http://www.cs.rpi.edu/academics/courses/spring13/ds/memory_debugging.html

To run the code above under Dr. Memory:

```
drmemory -brief -batch -- foo.exe arg1 arg2+
```

And here’s how Dr. Memory reports the errors:

```
~~Dr.M~~ Dr. Memory version 1.4.3
~~Dr.M~~ Running "E:\src\dr\test\drmemory_test.exe"
~~Dr.M~~
~~Dr.M~~ Error #1: UNINITIALIZED READ: reading 4 byte(s)
~~Dr.M~~ # 0 main [drmemory_test.cpp:6]
hi
~~Dr.M~~
~~Dr.M~~ Error #2: UNADDRESSABLE ACCESS: writing 4 byte(s)
~~Dr.M~~ # 0 main [drmemory_test.cpp:9]
~~Dr.M~~ Note: refers to 1 byte(s) beyond last valid byte in prior malloc
~~Dr.M~~
~~Dr.M~~ Error #3: LEAK 4 bytes
~~Dr.M~~ # 0 operator new
~~Dr.M~~ # 1 main [drmemory_test.cpp:5]
~~Dr.M~~
~~Dr.M~~ ERRORS FOUND:
~~Dr.M~~ 1 unique, 1 total unaddressable access(es)
~~Dr.M~~ 1 unique, 1 total uninitialized access(es)
~~Dr.M~~ 0 unique, 0 total invalid heap argument(s)
~~Dr.M~~ 0 unique, 0 total warning(s)
~~Dr.M~~ 1 unique, 1 total, 4 byte(s) of leak(s)
~~Dr.M~~ 0 unique, 0 total, 0 byte(s) of possible leak(s)
~~Dr.M~~ Details: E:\src\drmemory\git\build_drmemory/logs/DrMemory-drmemory_test.exe.5432.000/results.txt
```

And the fixed version:

```
~~Dr.M~~ Dr. Memory version 1.4.3
~~Dr.M~~ Running "E:\src\dr\test\drmemory_fixed.exe"
hi
~~Dr.M~~
~~Dr.M~~ ERRORS FOUND:
~~Dr.M~~ 0 unique, 0 total unaddressable access(es)
~~Dr.M~~ 0 unique, 0 total uninitialized access(es)
~~Dr.M~~ 0 unique, 0 total invalid heap argument(s)
~~Dr.M~~ 0 unique, 0 total warning(s)
~~Dr.M~~ 0 unique, 0 total, 0 byte(s) of leak(s)
~~Dr.M~~ 0 unique, 0 total, 0 byte(s) of possible leak(s)
```

Note that Dr. Memory does not report the Mismatched free() / delete / delete [] error, unless it is associated with indirect memory leaks.