

The MATH Library

(Version alpha 3-22-2002)

César A. R. Crusius

March 22, 2002 at 23:26

1. Introduction. The MATH library is my attempt to finally bring to the world the wonderful sensation of the power of object-oriented programming to numerical mathematics. The library is completely templated, so that adding a new type of matrix (such as a tridiagonal matrix) should not only be straightforward but old functions could still be used with it. The library can in principle be used with any system with STL installed and a C++ compiler that supports templates.

When designing the library I had the following philosophy in mind: when you deal with matrices, that's what you do. You deal with matrices. A matrix is a matrix, and only after that it is a sparse matrix, a symmetric matrix and so on. So the first decision I made was to provide only *one* matrix type, the **matrix** class. This is, in my opinion, the correct, intuitive and natural thing to do, and many libraries fail in this point because if you want to add a new matrix type you will actually have to define a whole new matrix class (when in fact you want to redefine the structure or whatever).

Next, I was concerned about an interesting issue: in most matrix libraries, there is *no* distinction between storage and structure. I'll explain better: you can have sparse matrices, or you can have symmetric matrices. But can you have sparse symmetric matrices? The concepts are really independent, but most systems do not address this issue correctly: we have to separate these concepts, and that's what I did. A matrix has as attributes a *structure* and a *storage method*. I'll now explain both in more detail:

- The *storage method* of a matrix determines how the entries are stored and how they are assigned values. For example, in a sparse matrix assigning an element to zero will actually remove it from memory. Note that this has nothing to do with the matrix structure!
- The *structure* of a matrix determines the relationship between the entries, and decides on *how* the storage will be used. I'll give examples: in a diagonal matrix, for example, the data could be stored in an one-column vector. In a symmetric matrix, the (i, j) element is equal to the (j, i) element, thus the structure can decide that assignment to (i, j) where $j > i$ will actually not take place. Note that the structure does not need to know *how* the entries are stored!

In fact, a symmetric structure actually has to deny assignment to half of the matrix. Consider, for example, the algorithm

$$A(i, j) = 2 * A(i, j).$$

The first assignment works fine, but when you try $A(j, i) = 2 * A(j, i)$ you will actually be multiplying the *new* $A(i, j)$ by two, and you'll have a wrong result. Due to this fact, we must make the following assumption: all algorithms compute all the elements of the matrix, unless you *really* know what you are doing. In the symmetric matrix case, for example, you can specialize an algorithm but you need to be sure which elements the symmetric structure really assign a value.

Now some words about performance. Suppose you are a performance freak, who likes to make nice graphics comparing how long does it take to make an SVD decomposition using various libraries. Probably the MATH library will loose. My main concern is to define a philosophically correct library in the programming sense. There is room, however, for performance improvements: you can always specialize the matrix type you are concerned with. The most obvious case is the dense and unstructured (or symmetric) matrix, which is used everywhere.

Brief remark: actually there are two correct philosophies: the template approach I just described and an "storage and structure hierarchy" approach, which would enable lists of different types of matrices but no specialization. Hence, I was guided by performance in some sense.

2. Basic definitions. The natural place to start is with the basic definitions, which are the the **index** and the **error** types: together, they provide us with the necessary tools to begin the longer **matrix** definition. These basic definitions, along with the matrix definition, are declared in the **math.h** file. As customary, the definitions are stated in the header file, and some bigger function bodies are defined in **math.cc**. As will happen with all MATH classes and functions, these definitions are declared inside the **math** namespace. The next three sections follow a pattern on this document: first we define the beginning of the source and header files, and then proceed with the code.

```

1  #include "math.h"
   <Big definitions 112>

3. <math.h 3> ≡
2  #ifndef __MATH__
3  #define __MATH__ 1.0
4  <Include files math 6>
5  <Preprocessor definitions>
6  namespace math {
7  <Basic definitions 4>
8  <Element definition 65>
9  <Structure definition 11>
10 <Storage definition 10>
11 <Matrix definition 12>
12 <Submatrix definition 81>
13 <Basic algebraic operations 105>
14 <Specializations 106>
15 <export-waiting big definitions 33>
16 }
17 #endif

```

4. The first thing we do is to define the C type that will be used for indexing elements of matrices. We define it as an **unsigned int**, which means, for one thing, that you can not define behaviors based on an **index** being negative. Since the entire library is based on the (1, 1) origin default, you can always use a zero value as an indicator. Note that an **index** definition already exists – it is defined in the standard library’s **string** class. Therefore, a safe use of the type will be **math::index**, even when you declare that you are using namespace **math**.

```

18 <Basic definitions 4> ≡
   typedef unsigned int index;

```

See also sections 5, 40, and 82.

This code is used in section 3.

5. In order to be able to differentiate MATH errors from others, we will define a some classes under the namespace `math::error` which we will use to signal anomalies. Our objective here is to provide a consistent and flexible means of passing errors from the library. The first generic error class will contain a string that can be used to describe the cause of the errors. Common errors can be derived from this class.

```

19  namespace error {
20      class generic {
21          string theMessage;
22      public:
23          <Generic error class methods 7>
24          virtual ~generic() {} /* A base class it is. */
25      };
26      <Predefined error types 9>
27  }
```

6. <Include files `math` 6> ≡
28 `#include <string>`

See also sections 31, 113, and 114.

This code is used in section 3.

7. The `generic` error class is intended to be used in `try-catch` mechanisms, so when something goes wrong you simply `throw` an error. In this way, it doesn't make sense to declare an `error` variable and update it during the program. Hence, the only provided way to modify an `error` is at the time of construction, and we will then need the appropriate constructors: the two defined below are used when the error we want to signal is not predefined. In that case, we create an unknown error with or without an explaining message.

```

<Generic error class methods 7> ≡
29     generic(void):theMessage("unknown") {}
30     generic(const char *msg):theMessage(msg) {}
```

See also section 8.

This code is used in section 5.

8. Of course, you may want to check out the message that was passed. We provide one method to retrieve the friendly error message. Depending on the error type we could even not consider it an error (see, for example, the `det` function).

```

<Generic error class methods 7> +≡
31     const string &message(void) const { return theMessage; }
```

9. Next we deal with predefined errors. What we do is to derive some classes that implement common errors.

⟨Predefined error types 9⟩ ≡

```

32  class singular : public generic {
33  public: singular()
34      : generic("Matrix_is_singular_to_working_precision.") {} };
35  class filerr : public generic {
36  public: filerr()
37      : generic("Generic_file_error.") {} };
38  class infeasible : public generic {
39  public: infeasible()
40      : generic("Problem_is_infeasible.") {} };
41  class nonsquare : public generic {
42  public: nonsquare()
43      : generic("Matrix_should_be_square") {} };
44  class nonpositivedef : public generic {
45  public: nonpositivedef()
46      : generic("Matrix_should_be_positive_definite.") {} };
47  class dimension : public generic {
48  public: dimension()
49      : generic("Wrong_matrix_dimensions.") {} };
50  class notimplemented : public generic {
51  public: notimplemented()
52      : generic("Feature_not_yet_implemented.") {} };
53  class maxiterations : public generic {
54  public: maxiterations()
55      : generic("Maximum_number_of_iterations_reached.") {} };
56  class unboundedbelow : public generic {
57  public: unboundedbelow()
58      : generic("Problem_is_unbounded_below.") {} };
59  class domain : public generic {
60  public: domain()
61      : generic("Domain_violation.") {} };
62  class rankdeficient : public generic {
63  public: rankdeficient()
64      : generic("Matrix_is_rank_deficient.") {} };

```

This code is used in section 5.

10. Matrix basics. Let us now begin the fun part. A matrix is a template on the type of its elements, its structure and its storage. Being that way, a matrix can be symmetric and sparse at the same time. The matrix is, then, a container for a storage of elements that are to be used according to some structure. The *function* of a matrix consists in coordinating the use of its storage by asking the structure what to do. In what follows we will define this protocol as necessity arrives by using the standard **dense** storage and **unstructured** structure.

⟨Storage definition 10⟩ ≡

```

65     template⟨class T⟩
66     class dense {
67         typedef T element_type;
68         ⟨Dense storage internal variables 28⟩
69     public:
70         ⟨Dense storage methods 29⟩
71     };

```

This code is used in section 3.

11. ⟨Structure definition 11⟩ ≡

```

72     template⟨class T⟩
73     class unstructured {
74         typedef T element_type;
75     public:
76         ⟨Unstructured structure methods 27⟩
77     };

```

This code is used in section 3.

12. Now to the user's matrix type. The first thing we need to realize is that we will inevitably need to return matrices from some functions (it suffices to think of a function to return the identity matrix). In order to minimize the overhead in passing matrices around, we separate the matrix *representation* from the matrix *handle*. The **matrix** type is, then, a handle to a representation that can be shared among many matrices. An assignment of the type $A = B$ will be equivalent to “matrix A will share matrix B representation.” Of course, we need to be careful to not modify matrices we don't want to. For example, $A(1,1) = 1$ can not modify B . As we proceed we will take care of all these cases.

#define *matrix_simple_template*

```

78         class T ← double ,
79         template⟨class⟩ class structure ← unstructured ,
80         template⟨class⟩ class storage ← dense

```

⟨Matrix definition 12⟩ ≡

```

81     ⟨Matrix representation definition 13⟩
82     template⟨matrix_simple_template⟩
83     class matrix {
84     public:
85         ⟨Matrix internal types 14⟩
86     private:
87         ⟨Matrix internal variables 15⟩
88     public:
89         ⟨Matrix methods 19⟩
90     };

```

This code is used in section 3.

13. As aligned before, a matrix is simply a handle to a matrix representation. It is the representation who holds the storage and has the structure information. Hence, a representation is a template of all of them.

⟨Matrix representation definition 13⟩ ≡

```

91     template⟨matrix_simple_template⟩
92     class representation {
93         typedef T element_type;
94         ⟨Matrix representation internal variables 17⟩
95     public:
96         ⟨Matrix representation methods 18⟩
97     };

```

This code is used in section 12.

14. A matrix knows its representation through a pointer to a it. In this way a single representation can be shared among various matrices. This also means that it is the matrix itself who must coordinate creation of new representations and destruction of them, as we shall see later.

⟨Matrix internal types 14⟩ ≡

```

98     typedef representation⟨T, structure, storage⟩ rep_type;

```

See also sections 16 and 88.

This code is used in section 12.

15. ⟨Matrix internal variables 15⟩ ≡

```

99     rep_type *theRepresentation;

```

This code is used in section 12.

16. Before going on, we note that we must have a means to obtain the type of elements, structure and storage of a matrix. Since a matrix is a template, we provide a means of accessing the type of the matrix as internal types in the same way as the Standard Template Library does.

⟨Matrix internal types 14⟩ +≡

```

100     typedef T element_type;
101     typedef storage⟨T⟩ storage_type;
102     typedef structure⟨T⟩ structure_type;
103     typedef matrix⟨T, structure, storage⟩ matrix_type;

```

17. Now back to business. Since the representation is the owner of the storage, it is the representation who holds the size of the matrix, not the **matrix** type.

⟨Matrix representation internal variables 17⟩ ≡

```

104     index num_rows, num_cols;

```

See also sections 20, 21, and 50.

This code is used in section 13.

18. As will happen often, the representation provides methods to access the information about the matrix, and the **matrix** type provides the interface to them. Often the representation methods don't perform any error checking, since the user will in the end use only the **matrix** type. The definition of two types of methods (one **const** and one not) is necessary for specializations (for example, the *reshape* function will modify the number of rows directly).

```

105  <Matrix representation methods 18> ≡
      index &rows(void) { return num_rows; }
106  index &cols(void) { return num_cols; }
107  index rows(void) const { return num_rows; }
108  index cols(void) const { return num_cols; }

```

See also sections 22, 25, 26, 47, 51, 54, 59, 63, 72, and 77.

This code is used in section 13.

```

109  19. <Matrix methods 19> ≡
      index rows(void) const { return (theRepresentation ? theRepresentation->rows() : 0); }
110  index cols(void) const { return (theRepresentation ? theRepresentation->cols() : 0); }

```

See also sections 23, 24, 34, 36, 37, 38, 42, 44, 45, 46, 53, 57, 58, 62, 71, 73, 74, 76, 89, 92, 94, 96, 99, 102, 108, and 115.

This code is used in section 12.

20. Also, as aligned before, it is the representation who has the storage and structure information. It will be the job of the representation to use them appropriately.

```

111  <Matrix representation internal variables 17> +≡
      storage<T> *theStorage;
112  structure<T> *theStructure;

```

21. Creating and copying a matrix. We can now begin to handle the ways in which a representation is handled by the matrix, and we will start by the ways a matrix is created. First, we need to know how many matrices are sharing a single representation. Obviously, this information can only be kept by the representation itself, and the matrix must be able to retrieve this information.

⟨Matrix representation internal variables 17⟩ +≡

```
113     int num_instances;
```

22. When a matrix is emptied or some operation modifies it, it is often the case that a matrix will need to create a new representation. On the other hand, copying matrices simply means incrementing the number of shared instances of some representation. Since it is the matrix who deals with this protocol, we return a reference to the number of instances so that a matrix can modify it itself.

⟨Matrix representation methods 18⟩ +≡

```
114     inline int &instances(void) { return num_instances; }
```

23. Let us then enable a matrix to be created. The default constructor should create an empty matrix, and the dimension constructor creates a fresh representation.

⟨Matrix methods 19⟩ +≡

```
115     matrix(void):theRepresentation(0) {}
```

24. ⟨Matrix methods 19⟩ +≡

```
116     matrix(const index rows, const index cols) #ifndef __MATH_INSANE_DEBUG__
117         cout << "[math]:_matrix(" << rows << ', ' << cols << ")_oldrep=" << theRepresentation;
118     #endif
119     theRepresentation ← new rep_type(rows, cols) #ifndef __MATH_INSANE_DEBUG__
120     cout << "_newrep=" << theRepresentation << '\n';
121     #endif
122     }
```

25. A representation is never empty. It is always created with dimensions, in the way you just saw. In order to really forbid the empty representation creation we define a default constructor which throws an error.

⟨Matrix representation methods 18⟩ +≡

```
123     representation(void)
124     {
125         throw error::generic("Cannot_instantiate_empty_matrix_representation!");
126     }
```

26. It is the job of the representation constructor to create the structure and the storage. By defining the representation constructor we begin, then, to define the communication protocol between storages and structures. We will explain the protocol by defining the **dense** and **unstructured** classes.

⟨Matrix representation methods 18⟩ +≡

```
127     representation(index rows, index cols):num_rows(rows), num_cols(cols), num_instances(1)
128     {
129         theStructure ← new structure<element_type>(&rows, &cols);
130         theStorage ← new storage<element_type>(rows, cols);
131     }
```

27. The first part of the protocol consists in creating a matrix with enough storage for its elements. A symmetric matrix, for example, does not need to store $rows * cols$ elements. Hence we first create the structure, and rely on the constructor to modify its arguments so that the a posteriori storage creation works accordingly. For the unstructured type, however, we don't need to modify the arguments.

```
<Unstructured structure methods 27> ≡
unstructured(index *, index *) { }
```

132

See also sections 49, 60, and 78.

This code is used in section 11.

28. Having now the correct dimensions of the matrix we can create the storage. Let's do it for the dense case. A dense matrix is stored as a vector in *elements* by columns, that is, the first elements are from the first column and so on. The index to the first element of a column is indexed in *data*. Later on this can appear to be unintuitive, but this scheme saves memory for vectors, when there is only one element in *data*, and also it will facilitate the build of a LAPACK interface. We will also allow the storage to have storage for more rows and columns than what is needed. This will be very useful when resizing.

```
<Dense storage internal variables 28> ≡
```

```
element_type *elements;
element_type **data;
index num_rows, num_cols;
index max_rows, max_cols;
```

133

134

135

136

This code is used in section 10.

29. We define, as usual, a default constructor to zero everything.

```
<Dense storage methods 29> ≡
dense(:elements(0), data(0), num_rows(0), num_cols(0), max_rows(0), max_cols(0) { }
```

137

See also sections 30, 32, 48, 52, 55, 56, 61, 64, 75, and 79.

This code is used in section 10.

30. Now to the useful dense constructor. We set the initial array values to zero, and we use *memset* for doing this. One could argue that we should do a loop and assign **element_type**(0) to each element, but let's get real... Since the task of initializing the data is also useful outside the constructor, we define an *init* method to be used for initialization and call it from the constructor.

```
<Dense storage methods 29> +≡
```

```
dense(const index &rows, const index &cols)
: elements(0), data(0) {
  init(rows, cols);
  memset(elements, 0, num_rows * num_cols * sizeof(element_type));
}
```

138

139

140

141

142

```
31. <Include files math 6> +≡
```

```
#include <string.h> /* For memset (and memcpy). */
```

143

32. The *init* function allocates space and initialize the data pointers, but leaves the data contents unchanged. The method will be constructed so that it can be used to resize the data if the number of elements remains the same (if they are not, you'll have segfault somewhere later). The trick in this method is that the vectors are initialized in a way such that the (1,1) origin standard works, that is, *data*[1][1] is the first element of the matrix. The only catch is that, due to the way the data is stored, the element $A(i, j)$ is stored in *data*[*j*][*i*].

```
<Dense storage methods 29> +≡
```

```
void init(const index &rows, const index &cols);
```

144

```

33. <export-waiting big definitions 33> ≡
145  template<class T> void dense<T>::init(const index &rows, const index &cols)
146  {
147      num_rows ← max_rows ← rows;
148      num_cols ← max_cols ← cols;
149      if (¬elements) elements ← new element_type[rows * cols];
150      if (data) delete[] ++data;
151      data ← new element_type*[cols];
152      for (index i ← 0; i < cols; ++i) data[i] ← elements + i * rows - 1;
153      data--;
154  }

```

See also sections 35, 39, 43, and 80.

This code is used in section 3.

34. Voilà! We are now able to create an empty matrix and a matrix filled with zeros. Sometimes, however, it is useful to create a matrix filled with some specified value. We provide a method to fill a matrix with a specified value and a constructor to do the same thing.

```

<Matrix methods 19> +≡
155  void fillwith(const element_type &value);

35. <export-waiting big definitions 33> +≡
156  template<matrix_simple_template> void matrix<T, structure, storage>::fillwith(const element_type
      &value)
157  {
158      for (index i ← 1; i ≤ rows(); i++)
159          for (index j ← 1; j ≤ cols(); j++) theRepresentation->set(i, j, value);
160  }

```

```

36. <Matrix methods 19> +≡
161  matrix(const index rows, const index cols, const element_type &value)
162  {
163  #ifdef __MATH_INSANE_DEBUG__
164      cout << "[math]:_matrix(" << rows << ', ' << cols << ', ' << value << ")_oldrep=" <<
          theRepresentation;
165  #endif
166      theRepresentation ← new rep_type(rows, cols);
167      fillwith(value);
168  #ifdef __MATH_INSANE_DEBUG__
169      cout << "_newrep=" << theRepresentation << '\n';
170  #endif
171  }

```

37. Next we define the copy constructor. As noted before, the only thing we need to do is to share the representation and update the number of shared instances.

```

<Matrix methods 19> +≡
172  matrix(const matrix_type &source)
173  {
174      if (source.theRepresentation) source.theRepresentation->instances()++;
175      theRepresentation ← source.theRepresentation;
176  }

```

38. Now we have all the basic constructors we need. The next step is to define the copy operations, which are very similar. In order to be able to do that we define a method that reinitializes a matrix to a given size. This method is like a constructor, but it takes into account the fact that the matrix can already have a representation. If it has one and its not shared and it has the same dimensions we don't need to do anything. In the other case we need to create a new representation anyway, taking care of the old one (the deletion of a representation is the topic of the next section, but the details are not necessary here).

⟨Matrix methods 19⟩ +≡

```
177     void init(const index num_rows, const index num_cols);
```

39. ⟨export-waiting big definitions 33⟩ +≡

```
178     template<matrix_simple_template> void matrix<T, structure, storage>::init(const index
179         num_rows, const index num_cols)
```

```
179     {
180         if (theRepresentation ^ cols() ≡ num_cols ^ rows() ≡ num_rows ^ theRepresentation-instances() ≡ 1)
181             return;
182         if (theRepresentation ^ --theRepresentation-instances() ≡ 0) delete theRepresentation;
183         theRepresentation ← new rep_type(num_rows, num_cols);
184     }
```

40. We are now able to define the assignment operator. In our context the task is a simple matter of using the same representation. Sometimes, however, this is not what we want: we want the contents copied and a new representation created. This normally occurs in numerical code where all matrices are modified almost instantly anyway. We provide a way to control the behavior of the assignment operator through the global *fast_assignment*. If *fast_assignment* is *true*, then the assignment operator uses the faster but memorywise expensive *copyfrom* method, to be defined later.

⟨Basic definitions 4⟩ +≡

```
185     extern bool fast_assignment;
```

41. `bool math::fast_assignment ← false;`

42. Now to the method itself. In case *fast_assignment* is *false*, we share the representation, the only catch being that we need to avoid confusions when making ridiculous things like $A = A$.

⟨Matrix methods 19⟩ +≡

```
186     matrix_type &operator←(const matrix_type &source);
```

43. ⟨export-waiting big definitions 33⟩ +≡

```
187     template<matrix_simple_template> matrix<T, structure, storage>::matrix_type
188         &matrix<T, structure, storage>::operator←(const matrix_type &source)
```

```
188     {
189         if (fast_assignment) return copyfrom(source);
190         if (source.theRepresentation) source.theRepresentation-instances()++;
191         if (theRepresentation ^ --theRepresentation-instances() ≡ 0) delete theRepresentation;
192         theRepresentation ← source.theRepresentation;
193         return *this;
194     }
```

44. We may also want to perform assignments with different types of source matrices. In that case the only way is to copy all elements, one by one. This method requires the operator `()`, which retrieves an element from a matrix, and the methods `set`, which assigns a value to a particular element. The details are irrelevant for now.

```

195     #define matrix_template(M)
196         class T ##M ,
197         template<class> class STR##M ,
198         template<class> class STO##M
199
200     <Matrix methods 19> +=
201     template<matrix_template(A)>
202     matrix_type &operator<←(const matrix<TA, STRA, STOA> &source)
203     {
204         init(source.rows(), source.cols());
205         index i, j;
206         for (i ← 1; i ≤ rows(); ++i)
207             for (j ← 1; j ≤ cols(); ++j) set(i, j, source(i, j));
208         return *this;
209     }

```

```

45. <Matrix methods 19> +=
210     template<matrix_template(A)>
211     matrix(const matrix<TA, STRA, STOA> &source):theRepresentation(0)
212     { *this ← source; }

```

46. On occasion we may want to copy a matrix without sharing the representation (for example, when we know the matrix will be modified right away so that a new representation will be created anyway if we share the instance). So here's what we do: if we don't have a representation yet we simply create a new one. If we do have one and it has the same dimensions (we know it has the same type), we assign a copy of the old one to it (this can be very fast). If none of this happens we behave just like the assignment operator, except we always will create a new representation.

```

210 <Matrix methods 19> +=
211     matrix_type &copyfrom(const matrix_type &source)
212     {
213         if (theRepresentation ≡ source.theRepresentation) return *this;
214         if (theRepresentation ∧ theRepresentation~instances() ≡ 1 ∧ rows() ≡ source.rows() ∧ cols() ≡
215             source.cols()) {
216             *theRepresentation ← *(source.theRepresentation);
217             return *this;
218         }
219         if (theRepresentation ∧ --theRepresentation~instances() ≡ 0) delete theRepresentation;
220         theRepresentation ← new rep_type(*(source.theRepresentation));
221         return *this;
222     }

```

47. In order to define the *copyfrom* method we created a representation based on an existing one. We will need to do the same thing when setting elements. For this task we will define a copy representation constructor, which in turn will require copy constructors for both the dense and unstructured classes.

⟨Matrix representation methods 18⟩ +≡

```

221     representation(const representation &source)
222     {
223         theStructure ← new structure⟨T⟩(*source.theStructure);
224         theStorage ← new storage⟨T⟩(*source.theStorage);
225         num_instances ← 1;
226         num_rows ← source.rows();
227         num_cols ← source.cols();
228     }

```

48. ⟨Dense storage methods 29⟩ +≡

```

229     dense(const dense &source)
230     : elements(0), data(0) {
231         init(source.num_rows, source.num_cols);
232         memcpy(elements, source.elements, num_rows * num_cols * sizeof(element_type));
233     }

```

49. ⟨Unstructured structure methods 27⟩ +≡

```

234     unstructured(const unstructured &) { }

```

50. We also used assignment operators in *copyfrom*. In order to simplify things, we only allow assignment operators to be called when the sources have the same dimensions. This may change in the future, but for now it's good enough (since the user will never use these operators directly).

⟨Matrix representation internal variables 17⟩ +≡

```

235     typedef representation⟨T, structure, storage⟩ rep_type;

```

51. ⟨Matrix representation methods 18⟩ +≡

```

236     rep_type &operator←(const rep_type &source)
237     {
238         if (source.rows() ≠ rows() ∨ source.cols() ≠ cols())
239             throw error::generic("Can_only_assign_representation_to_same_dimension.");
240         *theStorage ← *(source.theStorage);
241         *theStructure ← *(source.theStructure);
242         return *this;
243     }

```

52. Here we see why *copyfrom* can be much faster: the assignment operator for the representation didn't involve any memory allocation. Now, the **dense** assignment operator will not allocate memory too, and will use an optimized routine for copying elements. We can not use a single *memcpy* because there is a possibility that the source storage has a different size of allocated memory (because of *max_rows* and *max_cols*).

⟨Dense storage methods 29⟩ +≡

```

243     dense⟨T⟩ &operator←(const dense⟨T⟩ &source)
244     {
245         if (source.num_rows < num_rows ∨ source.num_cols < num_cols)
246             throw error::generic("Incompatible_dimension_in_dense_assignment_operator.");
247         for (index i ← 1; i ≤ num_cols; ++i)
248             memcpy(data[i] + 1, source.data[i] + 1, num_rows * sizeof(element_type));
249         return *this;
250     }

```

53. Destroying a matrix. Destroying a matrix would be very simple if it wasn't for the fact that another matrix can be sharing the same representation. This fact makes the destruction of a matrix to be downgraded to the "simple" category. When a matrix is deleted, it can happen that there is another matrix sharing the representation. We only delete the representation if we are the sole matrix using it. If we are not, we only update the number of matrices sharing the representation.

```

⟨Matrix methods 19⟩ +≡
249   ~matrix(void)
250   {
251   #ifdef __MATH_INSANE_DEBUG__
252       cout << "[math]:_~matrix,_"this=" << this << ",_"theRepresentation=" << theRepresentation <<
           '\n';
253   #endif
254       if (theRepresentation ^ --theRepresentation->instances() ≡ 0) delete theRepresentation;
255   }

```

54. Deleting a representation requires a representation destructor. If you try to delete a representation that is still shared MATH gives you a little piece of its mind.

```

⟨Matrix representation methods 18⟩ +≡
256   ~representation(void)
257   {
258       if (num_instances) throw error::generic("Deleting_a_referenced_representation!");
259       delete theStorage;
260       delete theStructure;
261   }

```

55. Of course, deleting a representation requires structure and storage destructors. For the unstructured class we rely on the default destructor, since the class is empty anyway. For the dense storage it is convenient, like with the constructing part, to define an auxiliary method and to call it within the destructor. In deleting the *data* vector we only have to keep in mind that it was adjusted during creation to obey the (1, 1) origin default.

```

⟨Dense storage methods 29⟩ +≡
262   ~dense() { destroy(); }

```

```

56. ⟨Dense storage methods 29⟩ +≡
263   void destroy(void)
264   {
265       if (elements) delete[] elements;
266       if (data) delete[] ++data;
267       elements ← 0;
268       data ← 0;
269   }

```


57. Setting and getting elements. At this point we are able to create and destroy matrices. The next step is to allow the user to set individual elements. Since the behavior of the *set* operation depends heavily on the matrix structure, the matrix type calls the representation to perform the structure-storage communication protocol. The matrix must, however, create a new representation in case the current one is shared, otherwise we will modify other matrices too. Since this operation is useful in other situations (for example when calling LAPACK routines), we define a method that creates a new representation for the matrix if necessary.

```

270   <Matrix methods 19> +≡
271   void detach(void)
272   {
273     if (theRepresentation ^ theRepresentation-instances() > 1) {
274       --theRepresentation-instances();
275       theRepresentation ← new rep_type(*theRepresentation);
276     }
  
```

58. We are now in position to set an element.

```

277   <Matrix methods 19> +≡
278   element_type set(index row, index col, element_type value)
279   {
280     detach();
281     return theRepresentation-set(row, col, value);
  
```

59. The representation performs the handling. First it calls the structure *preprocess* method. The structure will then modify the index accordingly so that a call to the storage *set* method will modify the correct element. If the structure returns *false* it means the element is not assignable.

```

282   <Matrix representation methods 18> +≡
283   element_type set(index row, index col, element_type value)
284   {
285     if (theStructure-preprocess(&row, &col)) return theStorage-set(row, col, value);
286     return theStorage-get(row, col);
  
```

60. For **unstructured** structures the *preprocess* method is empty, and the **dense** *set* method is trivial (remember the way the data is stored).

```

287   <Unstructured structure methods 27> +≡
288   bool preprocess(index *, index *) { return true; }
  
```

```

289   <Dense storage methods 29> +≡
290   element_type set(const index &row, const index &col, const element_type &value)
291   { return data[col][row] ← value; }
  
```

62. Wow! We are now able to assign elements to matrices! Next I guess you will want to retrieve them, and the most intuitive way to do that is to use the parenthesis operator. For reasons that will become clear, we should not allow the parenthesis operator to be used for assignment, that is, we should not allow something like $A(i, j) = 0$ (if you insist then yes, we can allow it, but with a huge price in performance). For this reason we do not return a reference, but the value itself. Now to the protocol. Again, a matrix doesn't know how to retrieve the value, since the way to retrieve something depends on the structure, so it calls the representation *get* method to do the job. For consistency reasons we also define a *get* method (to pair the *set* one).

⟨Matrix methods 19⟩ +≡

```
290  const element_type operator()(const index row, const index col ← 1) const
291  { return theRepresentation→get(row, col); }
292  const element_type get(const index row, const index col) const
293  { return theRepresentation→get(row, col); }
```

63. The representation needs now to perform the communication again. First it calls the structure *preprocess* to get the right index, and then retrieve the element from the storage. Strictly speaking we would need a postprocessing on the element value for Hermitian matrices, but in this feature will wait a little more.

⟨Matrix representation methods 18⟩ +≡

```
294  const element_type get(index row, index col) const
295  {
296  theStructure→preprocess(&row, &col);
297  return theStorage→get(row, col);
298  }
```

64. A storage *get* method must return zero if either the column or the row indexes are zero. In this way it is possible for a structure to force certain entries to be zero. In a diagonal matrix, for example, the structure *preprocess* method should assign zero to the index whenever $row \neq col$ (and also return *false*).

⟨Dense storage methods 29⟩ +≡

```
299  const element_type get(const index &row, const index &col) const
300  {
301  if (¬row ∨ ¬col) return element_type(0);
302  return data[col][row];
303  }
```

65. Now to the main issue. How to deal with $A(i, j) = x$? Assigning a zero element to a sparse matrix means deleting it, so we cannot return a reference to the data. There are two possible solutions: either we force the user to use the *set* method or we return a structure. The first solution is cumbersome and hard to read, hence we choose the second, and define a method that will return an *element*. An *element* works like a pointer to a **matrix** entry. When you assign an *element* a value, it calls the matrix *set* method accordingly.

⟨Element definition 65⟩ ≡

```
304  template<class matrix_type>
305  class element {
306  ⟨Matrix element internal variables 66⟩
307  public:
308  ⟨Matrix element methods 68⟩
309  };
```

This code is used in section 3.

66. An **element** doesn't need to know about its value, since, as said before, it works like a reference to a particular index of a particular matrix. These are, then the only data an **element** must have.

⟨Matrix element internal variables 66⟩ ≡

```
310     matrix_type *theMatrix;
```

```
311     index i, j;
```

See also section 67.

This code is used in section 65.

67. The numeric type of the element is retrieved from the matrix. This type is useful (essential!) for readability in the sequel.

⟨Matrix element internal variables 66⟩ +≡

```
312     typedef typename matrix_type::element_type element_type;
```

68. An element is not intended to be used as a user variable, and is also not intended to be constructed without arguments. Hence, we define the default constructor to throw an error. An element should be constructed with all necessary arguments, which are the matrix instance and the index to the matrix entry.

⟨Matrix element methods 68⟩ ≡

```
313     element(void):theMatrix(0),i(0),j(0)
```

```
314     { throw error::generic("Default constructor of element should not be used!"); }
```

See also sections 69, 70, 95, 97, 100, 103, and 107.

This code is used in section 65.

69. ⟨Matrix element methods 68⟩ +≡

```
315     element(matrix_type *mat, const index row, const index col):theMatrix(mat),i(row),j(col) { }
```

70. As aligned before, the main purpose of an element is to call the *set* method of its matrix when assigned a value. This will ensure the correct processing of the assignment taking into account the matrix structure.

⟨Matrix element methods 68⟩ +≡

```
316     inline element_type operator←(const element_type &value) const {
```

```
        return theMatrix→set(i, j, value); }
```

71. Now we are finally able to define the matrix method to be used for readable assignments. The method returns an element which you can assign a value in a friendly way. We name the method *entry*, so that it makes logical sense. In a program, you would write *A.entry*(1,1) ← 0 or something similar. As you can see, that construction is much better, in visual terms, than the equivalent *set* method. Depending on the compiler, however, it is slower. But you can't always get what you want.

⟨Matrix methods 19⟩ +≡

```
317     element⟨matrix_type⟩ entry(const index row, const index col ← 1)
```

```
318     {
```

```
319         return element⟨matrix_type⟩(this, row, col);
```

```
320     }
```

72. By now we are able to set and get individual elements values. It is, however, useful sometimes to be able to have direct access to the storage elements (for example, when interfacing with another library such as LAPACK). In this case, we also *detach* the matrix so that there is no danger of messing around with shared representations.

⟨Matrix representation methods 18⟩ +≡

```
321     storage⟨T⟩ *storg(void)
```

```
322     { return theStorage; }
```

73. \langle Matrix methods 19 $\rangle +\equiv$

```

323 storage_type *storg(void)
324 {
325     detach();
326     return (theRepresentation ? theRepresentation->storg() : 0);
327 }
```

74. And sometimes we may want the representation too.

\langle Matrix methods 19 $\rangle +\equiv$

```

328 rep_type *rep(void)
329 {
330     detach();
331     return theRepresentation;
332 }
```

75. Now that we have a means to get the matrix storage instance, we provide a means to have access to the storage elements themselves. Since this method will be used only when we know which kind of storage we have, it will differ between every storage type, and it is not required for one.

\langle Dense storage methods 29 $\rangle +\equiv$

```

333 element_type *memory(void) { return elements; }
```

76. Resizing. We define a resize operation as one that changes the size of the matrix but leaves the matrix elements in the same place. You can't assume anything about the value of elements that were not present in the original matrix. A matrix that has no representation simply creates one with the required dimensions. Otherwise, the matrix detaches itself from others and asks for a representation resizing.

⟨Matrix methods 19⟩ +≡

```

334   void resize(const index rows, const index cols)
335   {
336     if ( $\neg$ theRepresentation) {
337       init(rows, cols);
338       return;
339     }
340     detach();
341     theRepresentation→resize(rows, cols);
342   }
```

77. The protocol is as follows: first the representation calls the structure *resize* method. The structure returns the new dimensions for the storage (or throws a **dimension** error in case the new size is not compatible). Since the storage will throw an error if the new dimensions are wrong, we can safely update the representation variables before resizing. Next, the representation calls the storage *resize* method, which will take care of everything else.

⟨Matrix representation methods 18⟩ +≡

```

343   void resize(index rows, index cols)
344   {
345     num_rows ← rows;
346     num_cols ← cols;
347     theStructure→resize(&rows, &cols);
348     theStorage→resize(rows, cols);
349   }
```

78. ⟨Unstructured structure methods 27⟩ +≡

```

350   void resize(const index *, const index *) const { }
```

79. For the dense *resize* we will use some tricks. If the matrix can be resized without any memory allocation, then that's what will be done. This means that in those cases the resizing operation is very fast and that it can happen that resizing a big matrix to a small one doesn't free any memory. If we need to allocate memory we make sure that *init* won't erase the old data. We do that by copying it to another **dense** instance and zeroing out our data pointers.

⟨Dense storage methods 29⟩ +≡

```

351   void resize(const index rows, const index cols);
```

```
80. <export-waiting big definitions 33> +≡
352  template<class T> void dense<T>::resize(const index rows, const index cols)
353  {
354      if (cols ≤ max_cols ∧ rows ≤ max_rows) {
355          num_rows ← rows;
356          num_cols ← cols;
357          return;
358      }
359      dense backup(*this);
360      elements ← 0;
361      data ← 0;
362      init(rows, cols);
363      for (index j ← 1; j ≤ cols ∧ j ≤ backup.num_cols; ++j)
364          for (index i ← 1; i ≤ rows ∧ i ≤ backup.num_rows; ++i) data[j][i] ← backup.data[j][i];
365  }
```

81. Submatrices. Our next task is to provide a way to access submatrices, an operation that is very handy in many cases. A submatrix behaves like an **element** in the sense that it does not have any storage of its own, only a pointer to a matrix. The difference is that, while an element is a pointer to only one matrix entry, a submatrix is a pointer to a matrix block. The internal variables are the matrix the block refers to and the range, in the form (using MATLAB notation) $A(i1 : i2, j1 : j2)$. The basic design decision is that a submatrix, although a pointer to a generic matrix, produces only **unstructured** matrices when algebraic operations are performed. An example will clarify this necessity: suppose you have a submatrix of a symmetric matrix, and that this submatrix is not symmetric. What happens when you multiply this submatrix by a number is that you should return a **unstructured** submatrix.

```
#define submatrix_template(M) class SUBM ##M
```

```
<Submatrix definition 81> ≡
```

```
366     template<class T>
367     class submatrix {
368         <Submatrix internal variables 83>
369     public:
370         <Submatrix methods 85>
371     };
```

This code is used in section 3.

82. In order to be able to use submatrices from matrices (recall that the **submatrix** definition comes after the **matrix** one in the header file) we need to predeclare the **submatrix** template. We could have done this the other way around, but this way seems to me to be simpler.

```
<Basic definitions 4> +≡
```

```
372     template<class T> class submatrix;
```

83. In order to be able to use submatrices and matrices with the same template functions, we need to provide a means of detecting the matrix type, element type and so on. For a submatrix, the **matrix_type** is always **unstructured**, so we have an additional definition for the matrix the submatrix is referring to.

```
<Submatrix internal variables 83> ≡
```

```
373     typedef T internal_matrix_type;
374     typedef typename T::element_type element_type;
375     typedef submatrix<T> submatrix_type;
376     typedef matrix<element_type, unstructured, dense> matrix_type;
```

See also section 84.

This code is used in section 81.

84. The only data that is needed is a pointer to the matrix and the submatrix corners coordinates.

```
<Submatrix internal variables 83> +≡
```

```
377     internal_matrix_type *theMatrix;
378     index i1, i2, j1, j2;
```

85. A submatrix is not intended to be used as a user variable, and is also not intended to be constructed without arguments. This “not intended” is strong, in the sense that no care is currently taken to insure that the pointer a **submatrix** stores is valid in any sense. A submatrix is intended to use immediately as a shorthand for assignments and one-line formula references. We will adopt, for submatrices, a (0,0) internal origin standard in order to make assignments more efficient. One of the annoying things in defining constructors is that we have to use **const_cast** in some cases (if you are passing a submatrix as a **const** argument of some function you’ll need this unless you want your screen filled with “discards const” warnings).

⟨Submatrix methods 85⟩ ≡

```

379  submatrix(void):theMatrix(0)
380  { throw error::generic("Default constructor of submatrix should not be used!"); }
381  submatrix(const internal_matrix_type *mat, const index row1, const index row2, const index
      col1 ← 1, const index col2 ← 1)
382  {
383    theMatrix ← const_cast⟨internal_matrix_type *⟩(mat);
384    i1 ← row1 - 1;
385    i2 ← row2 - 1;
386    j1 ← col1 - 1;
387    j2 ← max(col2, col1) - 1;
388  }
```

See also sections 86, 87, 90, 91, 93, 98, 101, and 104.

This code is used in section 81.

86. A submatrix can be assigned values in two ways: by a matrix or by a submatrix. We deal with the former first for simplicity.

⟨Submatrix methods 85⟩ +≡

```

389  submatrix(const internal_matrix_type &mat) { *this ← mat; }
390  submatrix_type &operator←(const internal_matrix_type &mat)
391  {
392    if (mat.rows() ≠ i2 - i1 + 1 ∨ mat.cols() ≠ j2 - j1 + 1) throw error::dimension();
393    for (index i ← 1; i ≤ mat.rows(); ++i)
394      for (index j ← 1; j ≤ mat.cols(); ++j) theMatrix→set(i1 + i, j1 + j, mat.get(i, j));
395    return *this;
396  }
```

87. For the later type we need to make sure we are protected against things like $A(1 : 2, 1 : 2) = A(2 : 3, 2 : 3)$. In order to do that we will need to create a matrix from a submatrix. The thing to be aware of is that in assigning a submatrix to a matrix we don’t handle the shared representations anymore, instead we really create a new matrix; but before doing that we need some auxiliary methods from **submatrix** (note that the parenthesis operator obeys the (1,1) origin default). These auxiliary methods, together with others that will be defined later, will make a submatrix behave like a matrix for the basic operations.

⟨Submatrix methods 85⟩ +≡

```

397  index rows(void) const { return i2 - i1 + 1; }
398  index cols(void) const { return j2 - j1 + 1; }
399  element_type operator()(const index i, const index j ← 1) const
400  { return theMatrix→get(i1 + i, j1 + j); }
```

88. ⟨Matrix internal types 14⟩ +≡

```

401  typedef submatrix⟨matrix_type⟩ submatrix_type;
```


89. The following **matrix** methods are auxiliary to **submatrix**, but useful in their own right.

```

⟨Matrix methods 19⟩ +≡
402  template⟨class SUBM⟩ matrix(const submatrix⟨SUBM⟩ &mat)
403  : theRepresentation(0) { *this ← mat; }
404  template⟨class SUBM⟩ matrix_type &operator←(const submatrix⟨SUBM⟩ &mat)
405  {
406    init(mat.rows(), mat.cols());
407    for (index i ← 1; i ≤ rows(); ++i)
408      for (index j ← 1; j ≤ cols(); ++j) set(i, j, mat(i, j));
409    return *this;
410  }

```

90. Now to the last type of assignment. We need to check if our matrix is the same as the one being assigned, and if so if there is intersection between the source and destination elements. If that's the case we create a new matrix based on the source and copy from this matrix.

```

⟨Submatrix methods 85⟩ +≡
411  submatrix(const submatrix_type &mat) { *this ← mat; }
412  submatrix_type &operator←(const submatrix_type &mat)
413  {
414    if (theMatrix ≡ mat.theMatrix ∧ j1 ≤ mat.j2 ∧ j2 ≥ mat.j1 ∧ i1 ≤ mat.i2 ∧ i2 ≥ mat.i1)
415      *this ← internal_matrix_type(mat);
416    else
417      for (index i ← 1; i ≤ rows(); ++i)
418        for (index j ← 1; j ≤ cols(); ++j) theMatrix→set(i1 + i, j1 + j, mat(i, j));
419    return *this;
420  }

```

91. In order to make the submatrix type to resemble a matrix we define *set* and *get* methods for it.

```

⟨Submatrix methods 85⟩ +≡
421  void set(index row, index col, element_type value) { theMatrix→set(row + i1, col + j1, value); }
422  const element_type get(index row, index col ← 1) { return theMatrix→get(row + i1, col + j1); }

```

92. All we need now is a means to get a submatrix from a matrix. Since we can't overload the ':' operator, there is no good way to create one except by a specialized method (instead of overloading the parenthesis operator). We maintain some similarity with MATLAB in the sense that the first two arguments define the row range (and not the upper corner). The default behavior for the last column (which is to get the value of the starting column if that is greater) is handy sometimes. Note that the method is declared **const** even if it isn't. The reason is the same as the one described when defining **submatrix** constructors.

```

⟨Matrix methods 19⟩ +≡
423  submatrix_type subm(const index row1, const index row2, const index col1 ← 1, index
      col2 ← 1) const
424  {
425    col2 ← max(col2, col1);
426    return submatrix_type(this, row1, row2, col1, col2);
427  }

```

93. Again we provide a *subm* method for submatrices in order to make them resemble an actual matrix.

(Submatrix methods 85) +≡

```
428 submatrix_type subm(const index row1, const index row2, const index col1 ← 1, index
      col2 ← 1) const
429 {
430   col2 ← max(col2, col1);
431   return theMatrix-subm(i1 + row1, i1 + row2, j1 + col1, j1 + col2);
432 }
```

94. Basic algebraic operations. We are now able to define the most basic algebraic operations, such as sum, multiplication by scalars, and so on. We begin the unary operators. The methods are straightforward, but it is not as elegant as one might want it to be because we have to use the *set* and *get* methods: it is possible to use the already defined *entry*, but it could be slower, and the parenthesis operator is very cumbersome to use with a pointer (**this**). A common characteristic of the methods that follow is that the outer loop is generally the loop on the columns. We do that because we know that the unstructured matrix stores by column, so that it is more likely that the computer will make better use of the fast cache memory if we access the matrix this way. Let's start with scalar multiplication:

⟨Matrix methods 19⟩ +=

```

433     matrix_type &operator *= (const T &value)
434     {
435         for (index j ← 1; j ≤ cols(); ++j)
436             for (index i ← 1; i ≤ rows(); ++i) set(i, j, value * get(i, j));
437         return *this;
438     }
```

95. ⟨Matrix element methods 68⟩ +=

```

439     element_type operator *= (const element_type &val)
440     {
441         *this ← theMatrix-get(i, j) * val;
442         return value();
443     }
```

96. ⟨Matrix methods 19⟩ /=

```

444     matrix_type &operator /= (const T &value)
445     {
446         for (index j ← 1; j ≤ cols(); ++j)
447             for (index i ← 1; i ≤ rows(); ++i) set(i, j, get(i, j)/value);
448         return *this;
449     }
```

97. ⟨Matrix element methods 68⟩ /=

```

450     element_type operator /= (const element_type &val)
451     {
452         *this ← theMatrix-get(i, j)/val;
453         return value();
454     }
```

98. ⟨Submatrix methods 85⟩ +=

```

455     submatrix_type &operator /= (const element_type &value) {
456         for (index j ← 1; j ≤ cols(); ++j)
457             for (index i ← 1; i ≤ rows(); ++i) theMatrix-entry(i1 + i, j1 + j) /= value;
458         return *this;
459     }
```

99. Next we define the unary addition operator. In this case we test for dimension – we assume that the computational overhead in doing this is negligible when compared to the sum itself.

⟨Matrix methods 19⟩ +≡

```

460     template⟨submatrix_template(A)⟩
461     matrix_type &operator+=(const SUBMA &value)
462     {
463         if (rows() ≠ value.rows() ∨ cols() ≠ value.cols()) throw error::dimension();
464         for (index j ← 1; j ≤ cols(); ++j)
465             for (index i ← 1; i ≤ rows(); ++i) set(i, j, get(i, j) + value(i, j));
466         return *this;
467     }

```

100. ⟨Matrix element methods 68⟩ +≡

```

468     element_type operator+=(const element_type &val)
469     {
470         *this ← theMatrix-get(i, j) + val;
471         return value();
472     }

```

101. ⟨Submatrix methods 85⟩ +≡

```

473     submatrix_type &operator+=(const matrix_type &value)
474     {
475         if (value.rows() ≠ rows() ∨ value.cols() ≠ cols()) throw error::dimension();
476         for (index j ← 1; j ≤ cols(); ++j)
477             for (index i ← 1; i ≤ rows(); ++i) theMatrix-entry(i1 + i, j1 + j) += value.get(i, j);
478         return *this;
479     }

```

102. ⟨Matrix methods 19⟩ +≡

```

480     template⟨submatrix_template(A)⟩
481     matrix_type &operator−=(const SUBMA &value)
482     {
483         if (rows() ≠ value.rows() ∨ cols() ≠ value.cols()) throw error::dimension();
484         for (index j ← 1; j ≤ cols(); ++j)
485             for (index i ← 1; i ≤ rows(); ++i) set(i, j, get(i, j) − value(i, j));
486         return *this;
487     }

```

103. ⟨Matrix element methods 68⟩ +≡

```

488     element_type operator−=(const element_type &val)
489     {
490         *this ← theMatrix-get(i, j) − val;
491         return value();
492     }

```

104. <Submatrix methods 85> +=

```

493   submatrix_type &operator-=(const matrix_type &value)
494   {
495       if (value.rows() ≠ rows() ∨ value.cols() ≠ cols()) throw error::dimension();
496       for (index j ← 1; j ≤ cols(); ++j)
497           for (index i ← 1; i ≤ rows(); ++i) theMatrix-entry(i1 + i, j1 + j) -= value.get(i, j);
498       return *this;
499   }

```

105. The transpose operator has to be defined outside, because we'll want to specialize it (for example, for a symmetric matrix it's a void method, for a upper triangular it returns another type of matrix, and so on).

<Basic algebraic operations 105> ≡

```

500   template<matrix_simple_template>
501   matrix<T, structure, storage> transpose(const matrix<T, structure, storage> &x)
502   {
503       matrix<T, structure, storage> result(x.cols(), x.rows());
504       for (index j ← 1; j ≤ x.cols(); ++j)
505           for (index i ← 1; i ≤ x.rows(); ++i) result.set(j, i, x(i, j));
506       return result;
507   }
508   template<matrix_simple_template> matrix<T, structure, storage> transpose(const
509       submatrix<matrix<T, structure, storage>> &x)
510   {
511       matrix<T, structure, storage> result(x.cols(), x.rows());
512       for (index j ← 1; j ≤ x.cols(); ++j)
513           for (index i ← 1; i ≤ x.rows(); ++i) result.set(j, i, x(i, j));
514       return result;

```

This code is used in section 3.

106. Basic specializations and utilities. We're done with the matrix type and the basic structure and storage types. You could, using what we wrote, code anything from a matrix multiplication procedure to a complete eigenvalue/eigenvector decomposition. Sometimes, however, it is good to have some utilities predefined, and that's what this section is about.

While writing code for LU decomposition I felt the need for a *swap* algorithm for matrix elements. Since we cannot use the parenthesis operator for assignment, the STL's *swap* function doesn't work. The reason is interesting: a swap function must store one value in a temporary variable. The STL's *swap* assumes that this temporary variable has the same type as the function arguments, so it does something like **element** *aux* ← *A*. But, for our purposes, this is useless – the *value* of the element was not saved, only the matrix instance and the index to the entry! Hence, the need to specialize.

⟨Specializations 106⟩ ≡

```
515     template<class T> inline void swap(element<T> x, element<T> y)
516     {
517         typename T::element_type aux ← x.value();
518         x ← y.value();
519         y ← aux;
520     }
```

See also sections 109, 110, 111, and 116.

This code is used in section 3.

107. ⟨Matrix element methods 68⟩ +≡

```
521     inline element_type value(void) const { return theMatrix-get(i, j); }
```

108. Of course, I needed a *swap* method because I wanted to swap rows. This operation is useful in many other places, so I add a method that does exactly this.

⟨Matrix methods 19⟩ +≡

```
522     void swapprows(const index i, const index k)
523     {
524         for (index j ← 1; j ≤ cols(); ++j) swap(entry(i, j), entry(k, j));
525     }
```

109. Comparisons between matrices can be made in two ways: either we compare their elements one by one or we check if they share the same representation. The next method does exactly that, and it's useful in many contexts (there's a similar function in Lisp, for example, where we can have shared objects).

⟨Specializations 106⟩ +≡

```
526     template<matrix_simple_template>
527     bool same(const matrix<T, structure, storage> &x, const matrix<T, structure, storage> &y)
528     {
529         return (x.theRepresentation ≡ y.theRepresentation);
530     }
```

110. Next we define a new type of pair, a pair of indexes. We can use this pair for using stl's **map** with a **less** comparison function.

⟨Specializations 106⟩ +≡

```
531     typedef pair<math::index, math::index> pair;
```

111. ⟨Specializations 106⟩ +≡

```
532     bool operator < (const math::pair &x, const math::pair &y);
```

112. `<Big definitions 112> ≡`

```
533   bool operator <(const math::pair &x, const math::pair &y)
534   {
535       if (x.first > y.first) return false;
536       if (x.first < y.first) return true;
537       return (x.second < y.second);
538   }
```

This code is used in section 2.

113. `<Include files math 6> +≡`

```
539   #include <utility>
```

114. We will also interface this library with LAPACK at some point. In order to define the necessary Fortran stuff we include a file automatically generated during the library build. The file contains some LAPACK function prototypes and some definitions to take care of system-dependent LAPACK features (in some systems the function names need underscore, for example). This file will also have the LAPACK definition built-in.

`<Include files math 6> +≡`

```
540   #include <math/private/fortran.h>
```

115. While we are at it, some sparse structures used in other programs (such as PCx) have a field with the number of nonzero elements in the matrix. In order to make interface easier we define a method that returns this number.

`<Matrix methods 19> +≡`

```
541   index numnonzeros(void) const
542   {
543       index result ← 0;
544       for (index j ← 1; j ≤ cols(); ++j)
545           for (index i ← 1; i ≤ rows(); ++i)
546               if (get(i, j)) result++;
547       return result;
548   }
```

116. The *reshape* function works exactly as in Matlab (that is, the new matrix has the same elements taken columnwise from the original). We define it only for **unstructured** and **dense** matrices for now. We use our knowledge about the **dense** storage to make this operation very fast. The source matrix is overwritten with the new one.

`<Specializations 106> +≡`

```
549   template<class T>
550   matrix<T, unstructured, dense> &reshape(matrix<T, unstructured, dense> *x, index rows, index
551       cols)
552   {
553       if (x->rows() * x->cols() ≠ rows * cols) throw error::dimension();
554       x->storg()->init(rows, cols);
555       x->rep()->rows() ← rows;
556       x->rep()->cols() ← cols;
557       return *x;
558   }
```

117. Algebraic operations. We now reach the point where serious optimization begins: basic algebraic operations are the heart of any matrix package, and if your basic operations suck your entire package will suck too. Since we're defining generic interfaces, we can't possibly reach state-of-the-art performance – but we can get pretty close.

```
/* Empty, waiting for export */
```

```
118. <algebra.h 118> ≡
558 #ifndef __MATH_ALGEBRA__
559 #define __MATH_ALGEBRA__ 1.0
560 #include <math.h> /* For sqrt. */
561 #include <math/math.h>
562 namespace math {
563     <Algebraic operations 119>
564 }
565 #endif
```

119. We begin our definitions with binary algebraic operators. Unary operators were defined already. These are a little trickier than unary operators because we have structures: the matrix returned can have a different type of structure than that of the arguments (for example, a symmetric matrix times a symmetric matrix is not necessarily symmetric). We therefore start with scalar operations because this problem does not arise.

```
<Algebraic operations 119> ≡
566 template<matrix_simple_template> matrix<T, structure, storage> operator*(const
      matrix<T, structure, storage> &x, const T &y)
567 {
568     matrix<T, structure, storage> result ← x;
569     return result *= y;
570 }
```

See also sections 120, 121, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, and 134.

This code is used in section 118.

```
120. <Algebraic operations 119> +≡
571 template<matrix_simple_template> matrix<T, structure, storage> operator*(const
      submatrix<matrix<T, structure, storage>> &x, const T &y)
572 {
573     matrix<T, structure, storage> result ← x;
574     return result *= y;
575 }
```

```
121. <Algebraic operations 119> +≡
576 template<matrix_simple_template> matrix<T, structure, storage> operator+(const
      matrix<T, structure, storage> &x, const matrix<T, structure, storage> &y)
577 {
578     if (x.rows() ≠ y.rows() ∨ x.cols() ≠ y.cols()) throw error::dimension();
579     matrix<T, structure, storage> result ← x;
580     return result += y;
581 }
```


122. The first function we define is one that overwrites a matrix Y with the value of $\alpha x + y$. The name *saxpy* is of common use (it's used, for example, in Golub and Van Loan and in the BLAS package), so we keep it. Note that with this method the variable y is always overwritten, although some uses can be visually misleading. Also, it is the programmer who has to be sure that the structure of y and $\alpha x + y$ is the same (or compatible), since y is overwritten.

⟨ Algebraic operations 119 ⟩ +≡

```

582  template(class T, submatrix_template(X), submatrix_template(Y))
583  SUBMY &saxpy(const T &a, const SUBMX &x, SUBMY *y)
584  {
585    if (y-rows() ≠ x.rows() ∨ y-cols() ≠ x.cols()) throw error::dimension();
586    for (index j ← 1; j ≤ y-cols(); ++j)
587      for (index i ← 1; i ≤ y-rows(); ++i) y-set(i, j, a * x(i, j) + y-get(i, j));
588    return *y;
589  }
```

123. Next we define the dot product. This method only makes sense for vectors, and we test the dimensions before proceeding.

⟨ Algebraic operations 119 ⟩ +≡

```

590  template(submatrix_template(X), submatrix_template(Y))
591  typename SUBMX::element_type dot(const SUBMX &x, const SUBMY &y)
592  {
593    if (x.cols() ≠ 1 ∨ y.cols() ≠ 1 ∨ x.rows() ≠ y.rows()) throw error::dimension();
594    typename SUBMX::element_type result ← 0;
595    for (index i ← 1; i ≤ x.rows(); ++i) result += x(i) * y(i);
596    return result;
597  }
```

124. Sometimes we want to multiply a row vector by a column vector. We call this operation the “transposed *dot*” operation. Here x must be a row vector.

⟨ Algebraic operations 119 ⟩ +≡

```

598  template(submatrix_template(X), submatrix_template(Y))
599  typename SUBMX::element_type tdot(const SUBMX &x, const SUBMY &y)
600  {
601    if (x.rows() ≠ 1 ∨ y.cols() ≠ 1 ∨ x.cols() ≠ y.rows()) throw error::dimension();
602    typename SUBMX::element_type result ← 0;
603    for (index i ← 1; i ≤ x.cols(); ++i) result += x(1, i) * y(i);
604    return result;
605  }
```

125. And, just to make you nervous, sometimes we want to take the “dot product” of two row vectors. We call this operation the “transposed-transposed *dot*” operation. Here x and y must be row vectors.

⟨Algebraic operations 119⟩ +≡

```

606     template(submatrix_template(X), submatrix_template(Y))
607     typename SUBMX::element_type tdot(const SUBMX &x, const SUBMY &y)
608     {
609         if (x.rows() ≠ 1 ∨ y.rows() ≠ 1 ∨ x.cols() ≠ y.cols()) throw error::dimension();
610         typename SUBMX::element_type result ← 0;
611         for (index i ← 1; i ≤ x.cols(); ++i) result += x(i) * y(i);
612         return result;
613     }

```

126. Instead of defining the inner product of a vector we define the 2-norm operator. Since we’re calling the *dot* function we don’t need to check dimensions here.

⟨Algebraic operations 119⟩ +≡

```

614     template(submatrix_template(X))
615     typename SUBMX::element_type norm2(const SUBMX &x)
616     {
617         return (sqrt(dot(x, x)));
618     }

```

127. The following method returns the result of the operation Ax , where x is a vector or a matrix. The third argument is the destination matrix.

⟨Algebraic operations 119⟩ +≡

```

619     template(submatrix_template(A), submatrix_template(X), matrix_simple_template)
620     matrix⟨T, structure, storage⟩ &axmul(const SUBMA &A, const SUBMX
        &x, matrix⟨T, structure, storage⟩ *dest)
621     {
622         index n ← A.rows();
623         index m ← x.cols();
624         index p ← x.rows();
625         if (A.cols() ≠ p) throw error::dimension();
626         dest-resize(n, m);
627         for (index i ← 1; i ≤ n; ++i)
628             for (index j ← 1; j ≤ m; ++j) {
629                 T aux ← 0.0;
630                 for (index k ← 1; k ≤ p; ++k) aux += A(i, k) * x(k, j);
631                 dest-entry(i, j) ← aux;
632             }
633         return *dest;
634     }

```

128. And this one returns the result of the operation $A^T x$, where x can be a matrix. The third argument is the destination matrix.

```

635  <Algebraic operations 119> +=
636  template<submatrix_template(A), submatrix_template(X), matrix_simple_template>
637  matrix<T, structure, storage> &atxmul(const SUBMA &A, const SUBMX
638  &x, matrix<T, structure, storage> *dest)
639  {
640  index n ← A.cols();
641  index m ← x.cols();
642  index p ← x.rows();
643  if (A.rows() ≠ p) throw error::dimension();
644  dest→resize(n, m);
645  for (index i ← 1; i ≤ n; ++i)
646  for (index j ← 1; j ≤ m; ++j) {
647  T aux ← 0.0;
648  for (index k ← 1; k ≤ p; ++k) aux += A(k, i) * x(k, j);
649  dest→entry(i, j) ← aux;
650  }
651  return *dest;
652  }

```

129. Next we present the generalized saxpy. In this method, the value a of *saxpy* is substituted by a matrix, that is, now we have $Y = AX + Y$. This method is normally heavily used, and it involves a matrix multiplication – it’s a target of serious optimization in some libraries. Our generic version must be, well, generic. It shouldn’t be much worse than routines for dense and unstructured matrices, however.

```

635  <Algebraic operations 119> +=
636  template<submatrix_template(A), submatrix_template(X), submatrix_template(Y)>
637  SUBMY &gaxpy(const SUBMA &A, const SUBMX &x, SUBMY *y)
638  {
639  if (y→rows() ≠ A.rows() ∨ y→cols() ≠ x.cols() ∨ x.rows() ≠ A.cols()) throw error::dimension();
640  for (index j ← 1; j ≤ y→cols(); ++j)
641  for (index i ← 1; i ≤ y→rows(); ++i)
642  y→set(i, j, y→get(i, j) + tdot(A.subm(i, i, 1, A.cols()), x.subm(1, x.rows(), j)));
643  return *y;
644  }

```

130. We also define a method in case what we want to do is to compute $Y = A^T X + Y$.

```

635  <Algebraic operations 119> +=
636  template<submatrix_template(A), submatrix_template(X), submatrix_template(Y)>
637  SUBMY &gatxpy(const SUBMA &A, const SUBMX &x, SUBMY *y)
638  {
639  if (y→rows() ≠ A.cols() ∨ y→cols() ≠ x.cols() ∨ x.rows() ≠ A.rows()) throw error::dimension();
640  for (index j ← 1; j ≤ y→cols(); ++j)
641  for (index i ← 1; i ≤ y→rows(); ++i)
642  y→set(i, j, y→get(i, j) + dot(A.subm(1, A.rows(), i), x.subm(1, A.rows(), j)));
643  return *y;
644  }

```

131. The outer product operation (returns $X \cdot X^T$) is sometimes useful.

```

669  <Algebraic operations 119> +≡
670  template<matrix_simple_template, submatrix_template(X)>
671  matrix<T, structure, storage> &outerp(const SUBMX &x, matrix<T, structure, storage> *dest)
672  {
673      dest->resize(x.rows(), x.rows());
674      T aux;
675      for (index i ← 1; i ≤ x.rows(); ++i)
676          for (index j ← i; j ≤ x.rows(); ++j) {
677              aux ← ttdot(x.subm(i, i, 1, x.cols()), x.subm(j, j, 1, x.cols()));
678              dest->entry(i, j) ← aux;
679              dest->entry(j, i) ← aux;
680          }
681      return *dest;
682  }
```

132. And just to facilitate, we enable the operation $X^T \cdot X$ too.

```

683  <Algebraic operations 119> +≡
684  template<matrix_simple_template, submatrix_template(X)>
685  matrix<T, structure, storage> &outterp(const SUBMX &x, matrix<T, structure, storage> *dest)
686  {
687      dest->resize(x.cols(), x.cols());
688      T aux;
689      for (index i ← 1; i ≤ x.cols(); ++i)
690          for (index j ← i; j ≤ x.cols(); ++j) {
691              aux ← dot(x.subm(1, x.rows(), i, i), x.subm(1, x.rows(), j, j));
692              dest->entry(i, j) ← aux;
693              dest->entry(j, i) ← aux;
694          }
695      return *dest;
696  }
```

133. And finally a function to compute $xy^T + yx^T$, where x and y are vectors.

```

697  <Algebraic operations 119> +≡
698  template<submatrix_template(Y), submatrix_template(X), matrix_simple_template>
699  matrix<T, structure, storage> &xyyx(const SUBMX &x, const SUBMY &y, matrix<T,
700  structure, storage> *dest)
701  {
702      if (y.cols() ≠ 1 ∨ x.cols() ≠ 1 ∨ x.rows() ≠ y.rows()) throw error::dimension();
703      index n ← x.rows();
704      dest->resize(n, n);
705      for (index i ← 1; i ≤ n; ++i)
706          for (index j ← 1; j ≤ n; ++j) dest->entry(i, j) ← x(i) * y(j) + x(j) * y(i);
707      return *dest;
708  }
```

134. Now to the outer product update. The function overwrites the matrix A with $A + \beta xy^T$.

(Algebraic operations 119) +≡

```

705     template<submatrix_template(A), submatrix_template(X), submatrix_template(Y)> SUBMA
        &outerp_update(SUBMA *A, typename SUBMA::element_type beta, const SUBMX
        &x, const SUBMY &y)
706     {
707         if (y.cols() ≠ 1 ∨ x.cols() ≠ 1 ∨ A→rows() ≠ x.rows() ∨ A→cols() ≠ y.rows())
            throw error::dimension();
708         index n ← A→rows();
709         index m ← A→cols();
710         for (index j ← 1; j ≤ m; ++j)
711             for (index i ← 1; i ≤ n; ++i) A→set(i, j, A→get(i, j) + beta * x(i) * y(j));
712         return *A;
713     }

```

135. The sparse storage. Now that we have the basics defined, let us exemplify how to create new types of matrices by defining a new storage type.

```
/* Empty, waiting for export */
```

136. \langle `sparse.h` 136 $\rangle \equiv$

```
714 #ifndef __MATH_SPARSE__
715 #define __MATH_SPARSE__ 1.0
716  $\langle$  Include files sparse 139  $\rangle$ 
717 namespace math {
718  $\langle$  Sparse storage definition 137  $\rangle$ 
719 }
720 #endif
```

137. \langle Sparse storage definition 137 $\rangle \equiv$

```
721 template<class T>
722 class sparse {
723  $\langle$  Sparse storage internal variables 138  $\rangle$ 
724 public:
725  $\langle$  Sparse storage methods 140  $\rangle$ 
726 };
```

This code is used in section 136.

138. A sparse matrix stores its elements in a map in which the key is the element index. To simplify notation we define a new variable for the map type.

\langle Sparse storage internal variables 138 $\rangle \equiv$

```
727 typedef map<math::pair, T, less<math::pair>> storage;
728 storage *elements;
```

This code is used in section 137.

139. \langle Include files `sparse` 139 $\rangle \equiv$

```
729 #include <math/math.h> /* For math::pair – otherwise unnecessary. */
730 #include <map>
```

This code is used in section 136.

140. Now to the interesting part, that is, to the definition of methods that a matrix storage must have. We begin by the constructors and destructor: as with the **dense** case, we define a default constructor, a “dimension” constructor and a copy constructor. The interesting part is that we don’t need to initialize any data or keep any information whatsoever in the creation phase.

\langle Sparse storage methods 140 $\rangle \equiv$

```
731 sparse(void):elements(0) {}
732 sparse(const index &rows, const index &cols):elements(0) { elements ← new storage; }
733 sparse(const sparse &source):elements(0)
734 {
735     elements ← new storage;
736     elements→insert(source.elements→begin(), source.elements→end());
737 }
```

See also sections 141, 142, 143, and 144.

This code is used in section 137.

141. \langle Sparse storage methods 140 $\rangle + \equiv$

```
738 ~sparse() { if (elements) delete elements; }
```

142. We have now to define how to set and get elements. When setting elements we have to consider setting an element to zero – if the element exists we remove it, otherwise we don't set anything.

⟨Sparse storage methods 140⟩ +=

```

739     void set(const index &row, const index &col, const T &value)
740     {
741         math::pair i(row, col);
742         if (value ≠ T(0)) (*elements)[i] ← value;
743         else if (elements→find(i) ≠ elements→end()) elements→erase(i);
744     }

```

143. Getting elements is not trivial only because if the default behavior of maps is to create non-existing entry when accessing it. Recall that if either *row* or *col* are zero we are supposed to return 0 – which will happen because there is no element with index (0, 0) stored in a sparse matrix.

⟨Sparse storage methods 140⟩ +=

```

745     const T get(const index &row, const index &col) const
746     {
747         math::pair i(row, col);
748         if (elements→find(i) ≠ elements→end()) return (*elements)[i];
749         return T(0);
750     }

```

144. When resizing we let elements outside the new matrix to stay there just to make resizing fast (very fast, just an empty method). It is assumed that you won't be resizing a sparse matrix to very different sizes.

⟨Sparse storage methods 140⟩ +=

```

751     void resize(const index &rows, const index &cols) { }

```

145. And we are done. The methods listed for this class are the only required methods for a valid matrix storage type. Now that we have the dense and sparse types, however, there is little need for other types, so I consider this work done. The only other kind of storage I can imagine is the block-diagonal storage, but then we have some difficulties (we can have dense and sparse block-diagonal matrices and so on).

146. The symmetric structure. Now that we have a new storage defined, let us end the lesson on how to define new matrices by creating a new structure.

```
/* Empty, waiting for export */
```

```
147. <symmetric.h 147> ≡
752 #ifndef __MATH_SYMMETRIC__
753 #define __MATH_SYMMETRIC__ 1.0
754 #include <math/math.h>
755 namespace math {
756     <Symmetric structure definition 148>
757 }
758 #endif
```

```
148. <Symmetric structure definition 148> ≡
759 template<class T>
760 class symmetric {
761     public:
762     <Symmetric structure methods 149>
763 };
```

This code is used in section 147.

149. We start by defining the constructor that takes the dimensions as arguments. A symmetric matrix has to be square, and it can be represented with only $n(n+1)/2$ elements. We will store them as a vector as follows: if we have a matrix

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ * & a_{22} & a_{23} \\ * & * & a_{33} \end{bmatrix},$$

then we store is as the vector

$$[a_{11} \ a_{12} \ a_{22} \ a_{13} \ a_{23} \ a_{33}].$$

This scheme is the same one used in the LAPACK package, so we will use it to facilitate interfacing with it.

```
<Symmetric structure methods 149> ≡
764 void resize(math::index *rows, math::index *cols)
765 {
766     if (*rows ≠ *cols) throw error::nonsquare();
767     *rows ← (*rows) * (*rows + 1) / 2;
768     *cols ← 1;
769 }
770 symmetric(math::index *rows, math::index *cols) { resize(rows, cols); }
```

See also sections 150 and 151.

This code is used in section 148.

150. Next we define the *preprocess* method. After some algebra, you will see that the index of the element (i, j) is given by

$$\frac{j(j-1)}{2} + i$$

if $j \geq i$, that is, if the element is on the upper triangle of the matrix. In this class these are the elements that are stored. The elements of the lower triangle are not assignable.

⟨Symmetric structure methods 149⟩ +≡

```

771   bool preprocess(math::index *row, math::index *col) const
772   {
773     if (*col ≥ *row) {
774       *row ← (*col - 1) * (*col)/2 + (*row);
775       *col ← 1;
776       return true;    /* Assignable. */
777     }
778     *row ← (*row - 1) * (*row)/2 + (*col);
779     *col ← 1;
780     return false;    /* Not assignable. */
781   }

```

151. We need a copy constructor, which is empty.

⟨Symmetric structure methods 149⟩ +≡

```

782   symmetric(const symmetric &) {}

```

152. And we are done. I cannot help feeling proud of the flexibility of this system. In less than two pages we defined a whole new matrix structure, and this is taken into account documentation!

153. File streams: input. Files are the most convenient interface with other programs. I chose the MATLAB file format for simple interface with a widely used program. Since MATLAB's version 4 file format is public we will use it against the new format, which is proprietary and secret and sucks.

```
#include "fstream.h"
```

```
154. <fstream.h 154> ≡
783 #ifndef __MATH_FSTREAM__
784 #define __MATH_FSTREAM__ 1.0
785 <Include files fstream 156>
786 #include <math/math.h>
787 namespace math {
788     <fstream structures 162>
789     <fstream declarations 155>
790 }
791 #endif
```

155. We will first define the input stream, so that you will be able to load matrices from a MATLAB file. Almost all the work is already done by the standard library **ifstream** class, the only thing we need to do is to specialize it a little bit. We will want, for example, to find a matrix by name.

```
<fstream declarations 155> ≡
792 class ifstream : public std::ifstream {
793     public:
794     <ifstream methods 157>
795 };
```

See also section 181.

This code is used in section 154.

```
156. <Include files fstream 156> ≡
796 #include <fstream>
```

See also sections 163 and 171.

This code is used in section 154.

157. The first thing to do is to overload some constructors for compatibility with the standard **ifstream** class.

```
<ifstream methods 157> ≡
797 ifstream()
798 : std::ifstream() {}
799 ifstream(int fd)
800 : std::ifstream(fd) {}
```

See also sections 158, 159, 160, 164, 166, 167, 169, 170, 172, 174, 176, and 178.

This code is used in section 155.

158. Now to the opening of files. A matrix file is binary, so we change the default **ifstream** mode on both the constructor and the *open* method accordingly.

```
<ifstream methods 157> +≡
801 ifstream(const char *name, int mode ← ios::in | ios::binary, int prot ← °664)
802 : std::ifstream(name, mode, prot) {}
```

159. `< ifstream methods 157 > +≡`

```
803 void open(const char *name, int mode ← ios::in | ios::binary, int prot ← °664)
804 {
805     std::ifstream::open(name, mode, prot);
806 }
```

160. Unfortunately, the binary format of my library does not work for **long** or **double**, that is, I cannot use the operator `>>` to get these value types from a file. But we can fix it very easily by defining new operators. The trick, in both cases, is to declare an union of **unsigned char** and **long** (or **double**) and read **unsigned chars** from the file. Since we have the desired type in the union, what we are actually doing is filling in the value. We will only overload the `>>` operator for the **long** case, because it will be convenient when reading matrices headers. For the **double** case it is simpler to do the trick when reading the matrix data itself.

`< ifstream methods 157 > +≡`

```
807     ifstream &operator>>(long &);
```

161. `math::ifstream &math::ifstream::operator>>(long &dest)`

```
808 {
809     union {
810         long Long;
811         unsigned char Char[sizeof(long)];
812     } tricky;
813     for (int i ← 0; i ≠ sizeof(long); this->get(tricky.Char[i++])) ;
814     dest ← tricky.Long;
815     return *this;
816 }
```

162. Now we define a method that will read a matrix header. In MATLAB files, a matrix header consists of a sequence of fields that we will read in a structure as described below.

`< fstream structures 162 > ≡`

```
817 struct fheader {
818     long type; /* Type of the object. */
819     long rows; /* Number of rows. */
820     long cols; /* Number of cols. */
821     bool iflag; /* true if matrix has imaginary part. */
822     string name; /* Matrix name. */
823 };
```

This code is used in section 154.

163. `< Include files fstream 156 > +≡`

```
824 #include <string>
```

164. The actual data on a MATLAB file is somewhat different. The three first fields from **fheader** are the same. The *iflag* field is a **long** on the original file, and after it there is a **long** field with the number of characters in the name (including the terminating `'\0'`) followed by the matrix name. If we encounter an error while parsing the header we return, instead of throwing an error, since this method is not intended for end users.

`< ifstream methods 157 > +≡`

```
825 void parse_header(fheader &header);
```

```

165. void math::ifstream::parse_header(math::fheader &header)
826 {
827     long name_length;
828     long file_flag;
829     *this >> header.type >> header.rows >> header.cols >> file_flag >> name_length;
830     if (rdstate() != goodbit) return;
831     header.iflag ← bool(file_flag);
832     header.name.assign(name_length, 0); /* Reserve space for matrix name. */
833     for (int i ← 0; i ≠ int(name_length); this->get(header.name[i++])) ;
834 }

```

166. We are now in position to load a matrix. In a MATLAB file the matrix data is stored by columns, in two separate blocks for real and imaginary parts. If the matrix does not have an imaginary part there is only the real block in the file. There is only one complication: if we are to define a single method, then we must call, in case of a complex matrix, the method *set* with a complex argument. Even if this method is not called it must be compiled. Now, if the matrix is of type, say, **double**, then there will be a compiler error, because there is no conversion from **complex** to **double**. The solution to this problem is to define two functions, specializing for complex matrices. First we will define the method for non-complex matrices. In this case we simply disconsider the eventual imaginary part of the file matrix.

```

⟨ifstream methods 157⟩ +≡
835     template<matrix_simple_template>
836     ifstream &operator>>(matrix⟨T, structure, storage⟩ &dest)

```

167. Until the compiler accepts **export** we are stuck with the inline method.

```

⟨ifstream methods 157⟩ +≡
837 {
838     math::fheader header;
839     parse_header(header);
840     if (rdstate() != goodbit) throw math::error::filerr();
841     dest.init(header.rows, header.cols);
842     for (int ipart ← 0; ipart ≤ header.iflag; ipart++)
843         for (int j ← 1; j ≤ header.cols; j++)
844             for (int i ← 1; i ≤ header.rows; i++) {
845                 double number;
846                 ⟨Get number from file, using the same tricky method as for longs 168⟩
847                 if (¬ipart) dest.set(i, j, number);
848             }
849     return *this;
850 }

```

```

168. ⟨Get number from file, using the same tricky method as for longs 168⟩ ≡
851     union {
852         double Double;
853         unsigned char Char[sizeof(double)];
854     } tricky;
855     for (int k ← 0; k ≠ sizeof(double); this->get(tricky.Char[k++])) ;
856     number ← tricky.Double;

```

This code is used in sections 167 and 170.

169. Next we do the same thing for **complex** matrices. In this case there is nothing to worry about: if the file matrix does not have an imaginary part, the matrix itself will have the imaginary part zeroed.

⟨**ifstream** methods 157⟩ +≡

```
857     template⟨matrix_simple_template⟩
858     ifstream &operator≫(matrix⟨complex⟨T⟩, structure, storage⟩ &dest)
```

170. Until the compiler accepts **export** we are stuck with the inline method.

⟨**ifstream** methods 157⟩ +≡

```
859     {
860         math::fheader header;
861         parse_header(header);
862         if (rdstate() ≠ goodbit) throw math::error::filerr();
863         dest.init(header.rows, header.cols);
864         for (int ipart ← 0; ipart ≤ header.iflag; ipart++)
865             for (int j ← 1; j ≤ header.cols; j++)
866                 for (int i ← 1; i ≤ header.rows; i++) {
867                     double number;
868                     ⟨Get number from file, using the same tricky method as for longs 168⟩
869                     if (−ipart) dest.set(i, j, number);
870                     else dest.set(i, j, dest(i, j) + number * std::complex⟨double⟩(0, 1));
871                 }
872         return *this;
873     }
```

171. ⟨Include files *fstream* 156⟩ +≡

```
874     #include <complex>
```

172. At this point we are able to load a matrix. Sometimes, however, we will need to skip matrices (to load the second one, for example) or to find matrices by name. We begin by defining a method to skip matrices. First we parse the current matrix header and then skip the data using a *skip_data* method.

⟨**ifstream** methods 157⟩ +≡

```
875     void skip(int num_matrices);
```

173. **void** **math**::**ifstream**::*skip*(**int** *num_matrices*)

```
876     {
877         math::fheader header;
878         while (num_matrices --) {
879             parse_header(header);
880             if (rdstate() ≠ goodbit)
881                 throw math::error::generic("Too_many_matrices_to_skip.");
882             skip_data(header);
883         }
884     }
```

174. Skipping the data is a simple matter of reading the necessary amount of bytes from the body of the matrix.

⟨**ifstream** methods 157⟩ +≡

```
885     void skip_data(const fheader &header);
```

```

175. void math::ifstream::skip_data(const math::fheader &header)
886     {
887     long data_size ← header.cols * header.rows * sizeof(double);
888     if (header.iflag) data_size *= 2;
889     for (unsigned char garbage; data_size; data_size--) this->get(garbage);
890     }

```

176. With the above definitions we can jump to an arbitrary matrix on the file. The first matrix has position 1.

```

⟨ifstream methods 157⟩ +≡
891 void skipto(const int position);

```

```

177. void math::ifstream::skipto(int position)
892     {
893     seekg(0);
894     if (--position > 0) skip(position);
895     }

```

178. Now to matrix names. What we do when we want to skip to a certain matrix name is to first go to the beginning of the file and search for the matrix with the correct name. If we succeed, then the file will be *after* the header, so that the input operator cannot be used. The solution is to rewind the file again and skip to the correct matrix.

```

⟨ifstream methods 157⟩ +≡
896 void skipto(const char *matrix_name);

```

```

179. void math::ifstream::skipto(const char *matrix_name)
897     {
898     int position ← 0;
899     seekg(0);
900     ⟨Search for correct matrix name 180⟩
901     skipto(position);
902     }

```

```

180. ⟨Search for correct matrix name 180⟩ ≡
903 math::fheader header;
904 header.name ← "";
905 while (strcmp(header.name.c_str(), matrix_name)) {
906     if (position) skip_data(header); /* Skip data of previous matrix */
907     parse_header(header);
908     if (rdstate() ≠ goodbit)
909         throw math::error::generic("No_matrix_with_supplied_name_in_the_file.");
910     position++;
911 }

```

This code is used in section 179.

181. File streams: output. The input part is complete, so let us go on the next part. Again, the biggest job is already done by the standard library **ofstream** class, we just need to specialize it a little. The problem now is that we must give matrices a name before writing them to a file. We will do that by keeping a private string for the matrix name in the class and providing a means to modify it.

```

912  <fstream declarations 155> +≡
913  class ofstream : public std::ofstream {
914      string matrix_name;
915  public:
916      <ofstream methods 182>
917  };

```

182. We overload some constructors for compatibility with the standard **ofstream** library.

```

917  <ofstream methods 182> ≡
918  ofstream()
919  : std::ofstream() {}
920  ofstream(int fd)
921  : std::ofstream(fd) {}

```

See also sections 183, 184, 185, 187, 189, 190, 193, and 194.

This code is used in section 181.

183. Now to the opening of files. A matrix file is binary, so we change the default **ofstream** mode on both the constructor and the *open* method accordingly.

```

922  <ofstream methods 182> +≡
923  ofstream(const char *name, int mode ← ios::out | ios::binary, int prot ← °664)
924  : std::ofstream(name, mode, prot) {}

```

```

925  184. <ofstream methods 182> +≡
926  void open(const char *name, int mode ← ios::in | ios::binary, int prot ← °664)
927  {
928      std::ofstream::open(name, mode, prot);
929  }

```

185. Now to the writing algorithms: since we normally write matrices in sequence, providing a method like the **ifstream**'s *skipto* is unhandy and cumbersome to use. The solution is to “write the matrix name” to the stream before writing the matrix itself. This operation will assign the matrix name field of the header to the specified name.

```

930  <ofstream methods 182> +≡
931  ofstream &operator<<(const char *name);

```

```

932  186. math::ofstream &math::ofstream::operator<<(const char *name)
933  {
934      matrix_name ← name;
935      return *this;
936  }

```

187. Like with the input stream case, we define a method to output **longs**. The trick used is exactly the same.

```

937  <ofstream methods 182> +≡
938  ofstream &operator<<(const long &);

```

```

188.  math::ofstream &math::ofstream::operator<<(const long &source)
933   {
934     union {
935       long Long;
936       unsigned char Char[sizeof(long)];
937     } tricky;
938     tricky.Long ← source;
939     for (int i ← 0; i ≠ sizeof(long); this→put(tricky.Char[i++])) ;
940     return *this;
941   }

```

189. We are now ready to write matrices. We have here the same problem with **complex** types as we had in the input stream case, and we adopt the same specialization solution.

```

⟨ofstream methods 182⟩ +≡
942   template<matrix_simple_template>
943   ofstream &operator<<(const matrix⟨T, structure, storage⟩ &source)

```

190. Waiting for **export**...

```

⟨ofstream methods 182⟩ +≡
944   {
945     long iflag ← 0;
946     ⟨Write matrix header 191⟩
947     for (math::index j ← 1; j ≤ source.cols(); ++j)
948       for (math::index i ← 1; i ≤ source.rows(); ++i) {
949         double number;
950         number ← double(source(i, j));
951         ⟨Write number to file, using the same tricky mehtod as for longs 192⟩
952       }
953     return *this;
954   }

```

191. ⟨Write matrix header 191⟩ ≡

```

955   *this << long(0) << long(source.rows()) << long(source.cols()) << iflag <<
956   long(matrix_name.size() + 1);
957   for (unsigned int i ← 0; i ≤ matrix_name.size(); this→put(matrix_name.c_str()[i++])) ;

```

This code is used in sections 190 and 194.

192. ⟨Write *number* to file, using the same tricky mehtod as for **longs** 192⟩ ≡

```

957   union {
958     double Double;
959     unsigned char Char[sizeof(double)];
960   } tricky;
961   tricky.Double ← number;
962   for (int k ← 0; k ≠ sizeof(double); this→put(tricky.Char[k++])) ;

```

This code is used in sections 190 and 194.

193. Now to the complex specialization.

```

⟨ofstream methods 182⟩ +≡
963   template<matrix_simple_template>
964   ofstream &operator<<(const matrix⟨complex⟨T⟩, structure, storage⟩ &source)

```


194. Waiting for **export...**

```

965  {
966      long iflag ← 1;
967      ⟨Write matrix header 191⟩
968      for (int ipart ← 0; ipart ≤ 1; ++ipart)
969          for (math::index j ← 1; j ≤ source.cols(); ++j)
970              for (math::index i ← 1; i ≤ source.rows(); ++i) {
971                  double number;
972                  complex<double> aux;
973                  aux ← complex<double>(source(i, j));
974                  if (-ipart) number ← aux.real();
975                  else number ← aux.imag();
976                  ⟨Write number to file, using the same tricky method as for longs 192⟩
977              }
978      return *this;
979  }

```

195. With the end of the file streams classes we finished the main body of the library. With the definitions so far provided, you are able to do everything you want, including interfacing with MATLAB. Of course, what we have is too basic for complex algorithms, so our task now is to provide common functions, such as various types of matrix decompositions, solution of systems, and functions like determinant, inverse and so on. The point is that the main job is done, that is, the definition of the structures and methods necessary to develop algorithms. From now on there is really nothing really new, no impacting decisions to be made. The fun is not gone, but instead is replaced by the fun of using the power of the **matrix** class to write beautiful code.

196. The LU decomposition. We now begin to build the backbone of any linear algebra library: the matrix decomposition functions. We begin with the most simple, the LU decomposition. The goal is to decompose a matrix A into a product of a lower-triangular matrix L and an upper-triangular matrix U , that is, we want to have $A = LU$. The main purpose of this decomposition is to solve linear systems: if we want to solve $Ax = b$, we decompose matrix A and then solve two systems, $Ly = b$ and $Ux = y$ to find the solution. The point is that solving triangular systems is very easy by back or forward substitution.

```
/* Empty, waiting for export */
```

```
197. <lu.h 197> ≡
980 #ifndef __MATH_LU__
981 #define __MATH_LU__ 1.0
982 #include <math/math.h>
983 #include <vector>
984 #include <complex> /* For templated abs. */
985 namespace math {
986 namespace lu {
987 <LU prototypes 199>
988 } }
989 #endif
```

198. We can understand the way the decomposition works by considering the first step in a 3×3 matrix. We build a matrix M_1 such that the first column of M_1A is zero, except for the first element. The construction is

$$M_1A = \begin{bmatrix} 1 & 0 & 0 \\ -\frac{a_{21}}{a_{11}} & 1 & 0 \\ -\frac{a_{31}}{a_{11}} & 0 & 1 \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ 0 & a_{22} - \frac{a_{21}a_{12}}{a_{11}} & a_{23} - \frac{a_{21}a_{13}}{a_{11}} \\ 0 & a_{32} - \frac{a_{31}a_{12}}{a_{11}} & a_{33} - \frac{a_{31}a_{13}}{a_{11}} \end{bmatrix}.$$

Now we can repeat the procedure with the remaining blocks of the matrix so that

$$M_2M_1A = U,$$

where U is upper-triangular. Since M_i is lower-triangular, so it is its inverse, and defining $L = M_1^{-1}M_2^{-1}$ we finally have $A = LU$. A matrix of the form M_i is called a *Gauss transformation*, and the i th column of M_i is represented by $m_i = (0, 0, \dots, 1, -\tau_{i+1}, \dots, -\tau_n)$. The vector $\tau = (\tau_{i+1}, \dots, \tau_n)$ is called the *Gauss vector*. In this way, a Gauss transformation M_i can be represented as $M_i = (I - \tau^T e_i)$.

There is only one possible problem with the algorithm. Looking at the first step of the 3×3 decomposition example we can see that, if a_{11} is small, we have numerical problems because the Gauss vector could have too large elements. This effect can propagate through the rest of the algorithm. The solution is to find the decomposition for a permuted version of A , such that at every step we have the largest possible denominator in the Gauss vector. The denominator is called the *pivot*, and the method of permutations is called *pivoting*. The LU decomposition has a workload of $O(n^3)$ flops (actually exactly $2n^3/3$) and, if with pivoting, $O(n^2)$ comparisons.

199. Our basic function will be one that will overwrite the input matrix with U on the upper part and L on the lower part (it can be easily shown – see Golub and Van Loan – that L is a row permutation of the Gauss vectors). A vector *pivots* will store the data about the permutations, such that at step k we multiply A by the identity matrix with rows k and $pivots(k)$ swapped. The function returns 1 if the number of permutations is even and -1 otherwise. Since we will overwrite the matrix at will, we require it to have no structure.

⟨LU prototypes 199⟩ ≡

```

990  template<class T, template<class> class storage>
991  int decompose(matrix<T, unstructured, storage> *A, vector<index> *pivots)
992  {
993      int permutations ← 1;    /* No permutations yet. */
994      index n ← A→rows();
995      if (n ≠ A→cols()) throw error::nonsquare();
996      pivots→resize(n - 1, 0); /* Reserve space for permutation data. */
997      for (index k ← 1; k ≠ n; ++k) {
998          T pivot;
999          ⟨Search for pivot and swap rows if necessary 200⟩
1000         if (pivot ≠ 0) ⟨Apply Gauss transformation 201⟩
1001         }
1002         return permutations;
1003     }

```

See also sections 202, 203, and 204.

This code is used in section 197.

200. ⟨Search for pivot and swap rows if necessary 200⟩ ≡

```

1004  pivot ← A→get(k, k);
1005  (*pivots)[k - 1] ← k;    /* STL vector has base index 0. */
1006  for (index i ← k + 1; i ≤ n; ++i)
1007      if (abs(A→get(i, k)) > abs(pivot)) {
1008          pivot ← A→get(i, k);
1009          (*pivots)[k - 1] ← i;
1010      }
1011  if ((*pivots)[k - 1] ≠ k) /* If need to swap. */
1012  {
1013      permutations *= -1;
1014      A→swaprows(k, (*pivots)[k - 1]);
1015  }

```

This code is used in section 199.

201. ⟨Apply Gauss transformation 201⟩ ≡

```

1016  for (index i ← k + 1; i ≤ n; ++i) {
1017      if (¬finite(A→entry(i, k) ← A→get(i, k)/A→get(k, k))) throw error::singular();
1018      for (index j ← k + 1; j ≤ n; ++j) A→entry(i, j) ← A→get(i, j) - A→get(i, k) * A→get(k, j);
1019  }

```

This code is used in section 199.

202. Solving a linear system. As hinted before, the principal application of the LU decomposition is the solution of linear systems. The function we will define will solve a matrix linear system $AX = B$. The decomposition just defined gives us information to compute $P_{n-1} \dots P_1 A = LU$, where P_i are permutation matrices defined by the *pivots* vector. Having the LU decomposition, we solve the system $Ax = b$ by first solving $Ly = P_{n-1} \dots P_1 b$ and then $Ux = y$. By doing this to all columns of X and B we are able to solve $AX = B$. We define now a function that, given a decomposition and a matrix B , overwrite in B the solution to $AX = B$. We define this function separately because if you want to solve $A^k X = B$, you need to perform only one decomposition and call this function k times.

⟨LU prototypes 199⟩ +=

```

1020  template<class T, template<class> class storage>
1021  void finish(const matrix<T, unstructured, storage> &A, const vector<index>
      &pivots, matrix<T, unstructured, storage> *B)
1022  {
1023      index n ← A.rows();
1024      for (index i ← 1; i < n; ++i) /* B ← Pn-1 ... P1B */
1025          B→swaprows(i, pivots[i - 1]);
1026      for (index k ← B→cols(); k ≠ 0; --k) {
1027          for (index i ← 1; i ≤ n; ++i) /* Solve Ly = B(:, k) */
1028              {
1029                  T inner ← 0;
1030                  for (index j ← 1; j ≠ i; ++j) inner += A(i, j) * B→get(j, k);
1031                  B→entry(i, k) ← B→get(i, k) - inner;
1032              }
1033          for (index i ← n; i ≥ 1; --i) /* Solve Ux = B(:, k). */
1034              {
1035                  T inner ← 0;
1036                  for (index j ← i + 1; j ≤ n; ++j) inner += A(i, j) * B→get(j, k);
1037                  if (¬finite(B→entry(i, k) ← (B→get(i, k) - inner)/A(i, i))) throw error::singular();
1038              }
1039          }
1040      }

```

203. Finally, we provide an interface (that will destroy the original matrix, by the way). Note that we can compute the inverse of a matrix by calling this function with $B = I_n$.

⟨LU prototypes 199⟩ +=

```

1041  template<class T, template<class> class storage>
1042  matrix<T, unstructured, storage> &solve(matrix<T, unstructured, storage>
      *A, matrix<T, unstructured, storage> *B)
1043  {
1044      vector<index> pivots;
1045      decompose(A, &pivots);
1046      finish(*A, pivots, &B);
1047      return *B;
1048  }

```

204. Interfacing with LAPACK. The routines just written enable you to compute the LU decomposition of any type of matrix. For certain types, however, we have extremely efficient decomposition functions already coded in the LAPACK package. It only makes sense to call these routines when possible, and the MATH library provides the ideal transparent interface: you call the exact same function, and if there is a LAPACK function to do the job, then the function will be called.

```

1049  <LU prototypes 199> +=
      #ifdef LAPACK
1050    <LU lapack interface 205>
1051  #endif

```

205. The key for the transparent operation is specialization. The trick here is that we need to make sure that the matrix does not share the representation, otherwise the call to the LAPACK routine will modify all matrices sharing it.

```

1052  <LU lapack interface 205> ≡
      template<
1053    int decompose(matrix<fortran::double_precision, unstructured, dense> *A, vector<index>
          *pivots)
1054    {
1055      <Prepare for lapack LU 206>
1056      dgetrf(&m, &n, A->storg()-memory(), &m, (int *) pivots->begin(), &status);
1057      <Check for lapack LU errors 207>
1058      <Compute number of permutations and store it in status 208>
1059      return status;
1060    }

```

See also section 209.

This code is used in section 204.

```

206.  <Prepare for lapack LU 206> ≡
1061    fortran::integer m ← fortran::integer(A->rows());
1062    fortran::integer n ← fortran::integer(A->cols());
1063    fortran::integer status ← 0;
1064    pivots->resize(min(m, n), 0);

```

This code is used in sections 205 and 209.

```

207.  <Check for lapack LU errors 207> ≡
1065    if (status > 0) throw error::singular();
1066    if (status < 0) throw error::generic();

```

This code is used in sections 205 and 209.

208. In order to be compatible with our own LU decomposition routine, we need to compute the number of permutations (1 means even number, -1 means odd).

```

<Compute number of permutations and store it in status 208> ≡
1067    status ← 1;
1068    for (index i ← 1; i < (index) min(m, n); i++)
1069      if ((*pivots)[i - 1] ≠ i) status *= -1;

```

This code is used in sections 205 and 209.

209. <LU lapack interface 205> +≡

```
1070 template<>
1071 int decompose(matrix<fortran::real, unstructured, dense> *A, vector<index> *pivots)
1072 {
1073     <Prepare for lapack LU 206>
1074     sgetrf(&m, &n, A-storg()-memory(), &m, (int *) pivots-begin(), &status);
1075     <Check for lapack LU errors 207>
1076     <Compute number of permutations and store it in status 208>
1077     return status;
1078 }
```

210. The Cholesky decomposition. The LU decomposition works for all types of square matrices. If, however, a matrix is symmetric, then we will see that we can cut the work in half: first, instead of decomposing $A = LU$ we do $A = LDU$, where D is diagonal and both L and U have unit diagonal (the matrix L is already unit diagonal, and we can make the matrix U the same by scaling it with a diagonal matrix D). Now, if A is symmetric, then we have $L^{-1}AL^{-T} = DUL^{-T}$ is also symmetric, but this can only be true if $U = L^T$. Hence, if A is symmetric we can decompose it such that $A = LL^T$, so we need to find only one matrix. If, in addition, A is positive definite, then there is no need for pivoting, and the resulting decomposition is called the “Cholesky decomposition.”

```
/* Empty, waiting for export */
```

```
211. <cholesky.h 211> ≡
1079 #ifndef __MATH_CHOLESKY__
1080 #define __MATH_CHOLESKY__ 1.0
1081 #include <math/math.h>
1082 #include <math/symmetric.h>
1083 #include <math.h> /* for sqrt. */
1084 namespace math {
1085 namespace cholesky {
1086 <Cholesky prototypes 212>
1087 } }
1088 #endif
```

212. Another way to see that we can decompose $A = LL^T$ if A is symmetric is by establishing the equality

$$\begin{bmatrix} a_{11} & \alpha^T \\ \alpha & B \end{bmatrix} = \begin{bmatrix} \sqrt{a_{11}} & 0 \\ \alpha/\sqrt{a_{11}} & I \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & B - \alpha\alpha^T/a_{11} \end{bmatrix} \begin{bmatrix} \sqrt{a_{11}} & \alpha^T/\sqrt{a_{11}} \\ 0 & I \end{bmatrix},$$

which already hints the algorithm. One characteristic of the decomposition is that it is stable even without pivoting, and that during the entire algorithm all diagonal elements remain positive if the matrix itself is positive definite. The function we will define will overwrite the upper triangular part of A with L^T . We do that because now we know that a symmetric matrix structure stores only the upper triangle elements.

```
<Cholesky prototypes 212> ≡
1089 template<matrix_simple_template>
1090 void decompose(matrix<T, structure, storage> *A)
1091 {
1092     index n ← A→rows();
1093     if (n ≠ A→cols()) throw error::nonsquare();
1094     for (index k ← 1; k ≤ n; ++k) {
1095         if (A→get(k, k) ≤ 0) throw error::nonpositivedef();
1096         <Build LT row 213>;
1097         <Apply Cholesky transformation 214>;
1098     }
1099 }
```

See also sections 215, 216, and 217.

This code is used in section 211.

```
213. <Build LT row 213> ≡
1100 A→entry(k, k) ← sqrt(A→get(k, k));
1101 for (index j ← k + 1; j ≤ n; ++j)
1102     if (¬finite(A→entry(k, j) ← A→get(k, j)/A→get(k, k))) throw error::singular();
```

This code is used in section 212.

214. \langle Apply Cholesky transformation 214 $\rangle \equiv$

```
1103   for (index  $i \leftarrow k + 1$ ;  $i \leq n$ ;  $++i$ )
1104       for (index  $j \leftarrow i$ ;  $j \leq n$ ;  $++j$ )  $A\text{-entry}(i, j) \leftarrow A\text{-get}(i, j) - A\text{-get}(k, i) * A\text{-get}(k, j)$ ;
```

This code is used in section 212.

215. Solving a linear system. The Cholesky decomposition uses half the number of flops as the LU, and in addition there is no pivoting overhead. It is only advisable to use it to solve linear equations with positive definite matrices. We solve a system $AX = B$ vector by vector of X by first solving $Ly = b$ and then $L^T x = y$. As with the LU decomposition we define a function that solves the system given a previously computed decomposition. The solution X is overwritten in matrix B .

\langle Cholesky prototypes 212 $\rangle + \equiv$

```
1105   template<class T, template<class> class structure_A, template<class> class
        structure_B, template<class> class storage>
1106   void finish(const matrix<T, structure_A, storage> &A, matrix<T, structure_B, storage> *B)
1107   {
1108       index  $n \leftarrow A\text{-rows}()$ ;
1109       for (index  $k \leftarrow B\text{-cols}()$ ;  $k \neq 0$ ;  $--k$ ) {
1110           for (index  $i \leftarrow 1$ ;  $i \leq n$ ;  $++i$ ) /* Solve  $Ly = B(:, k)$ . */
1111               {
1112                   T inner  $\leftarrow 0$ ;
1113                   for (index  $j \leftarrow 1$ ;  $j < i$ ;  $++j$ )
1114                       /* Here we have to remember that we overwrote only the upper triangular part of A, so that
1115                        now we have to get the element  $A(j, i)$  instead of  $A(i, j)$ . */
1116                       inner  $+= A(j, i) * B\text{-get}(j, k)$ ;
1117                   if ( $\neg\text{finite}(B\text{-entry}(i, k) \leftarrow (B\text{-get}(i, k) - \text{inner})/A(i, i))$ ) throw error::singular();
1118               }
1119           for (index  $i \leftarrow n$ ;  $i \geq 1$ ;  $--i$ ) /* Solve  $L^T x = y$ . */
1120               {
1121                   T inner  $\leftarrow 0$ ;
1122                   for (index  $j \leftarrow i + 1$ ;  $j \leq n$ ;  $++j$ ) inner  $+= A(i, j) * B\text{-get}(j, k)$ ;
1123                   if ( $\neg\text{finite}(B\text{-entry}(i, k) \leftarrow (B\text{-get}(i, k) - \text{inner})/A(i, i))$ ) throw error::singular();
1124               }
1125           }
1126       }
```

216. Finally, we provide an interface for the linear solver. The big warning is that the A matrix will be overwritten too, not only B !!!

\langle Cholesky prototypes 212 $\rangle + \equiv$

```
1125   template<class T, template<class> class structure_A, template<class> class
        structure_B, template<class> class storage>
1126   matrix<T, structure_B, storage> &solve(matrix<T, structure_A, storage> *A, matrix<T,
        structure_B, storage> *B)
1127   {
1128       decompose(A);
1129       finish(*A, B);
1130       return *B;
1131   }
```


217. Interfacing with LAPACK. Again, the routines we just defined can compute the Cholesky decomposition of any matrix. If you have LAPACK installed, however, you can take advantage of years and years of laborious code optimization – why not use it?

```

1132  <Cholesky prototypes 212> +≡
1133  #include <math/private/fortran.h>
1134  #ifdef HAVE_LIBLAPACK
1135  <Cholesky lapack interface 218>
1136  #endif

```

218. We begin with the double precision symmetric matrix case. The others are very similar, we just have to replace the function call.

```

1137  <Cholesky lapack interface 218> ≡
1138  template<
1139  void decompose(matrix<fortran::double_precision, symmetric, dense> *A)
1140  {
1141  <Prepare for lapack Cholesky 219>
1142  dpptrf(&mode, &n, A->storg()-memory(), &status);
1143  <Check lapack Cholesky errors 220>
1144  }

```

See also sections 221, 222, and 223.

This code is used in section 217.

```

1145  <Prepare for lapack Cholesky 219> ≡
1146  fortran::integer n ← A->rows();
1147  fortran::integer m ← A->cols();
1148  fortran::integer status ← 0;
1149  fortran::character mode ← 'U';
1150  if (n ≠ m) throw error::nonsquare();

```

This code is used in sections 218, 221, 222, and 223.

```

1151  <Check lapack Cholesky errors 220> ≡
1152  if (status > 0) throw error::nonpositivedef();
1153  if (status < 0) throw error::generic();

```

This code is used in sections 218, 221, 222, and 223.

221. As said before, the rest of the routines is essentially the same, except that for unstructured matrices we need to pass also the number of columns to the LAPACK routines.

```

1154  <Cholesky lapack interface 218> +≡
1155  template<
1156  void decompose(matrix<fortran::real, symmetric, dense> *A)
1157  {
1158  <Prepare for lapack Cholesky 219>
1159  spptrf(&mode, &n, A->storg()-memory(), &status);
1160  <Check lapack Cholesky errors 220>
1161  }

```

222. We begin with the symmetric matrix cases.

⟨ Cholesky lapack interface 218 ⟩ +≡

```
1157     template<>
1158     void decompose(matrix<fortran::double_precision, unstructured, dense> *A)
1159     {
1160         ⟨ Prepare for lapack Cholesky 219 ⟩
1161         dpotrf(&mode, &n, A->storg()-memory(), &m, &status);
1162         ⟨ Check lapack Cholesky errors 220 ⟩
1163     }
```

223. ⟨ Cholesky lapack interface 218 ⟩ +≡

```
1164     template<>
1165     void decompose(matrix<fortran::real, unstructured, dense> *A)
1166     {
1167         ⟨ Prepare for lapack Cholesky 219 ⟩
1168         spotrf(&mode, &n, A->storg()-memory(), &m, &status);
1169         ⟨ Check lapack Cholesky errors 220 ⟩
1170     }
```

224. The QR decomposition. The last two decompositions worked with square matrices, decomposing them into matrices with some *structure* that made easy the task of solving linear systems. The QR decomposition, on the other hand, produces $A = QR$, where R is upper triangular and Q is *orthogonal*, that is, $Q'Q = QQ' = I$. This decomposition provides a straightforward way to solve the problem $\min_x \|Ax - b\|_2$, and is also the basis for one of the most widely used algorithms for computing the SVD decomposition and the eigenvalues of a matrix.

```
/* Empty, waiting for export. */
```

225. `<qr.h 225> ≡`

```
1171 #ifndef __MATH_QR__
1172 #define __MATH_QR__ 1.0
1173 #include <math/math.h>
1174 #include <math/algebra.h>
1175 namespace math {
1176 namespace qr {
1177     <QR prototypes 226>
1178 } }
1179 #endif
```

226. The QR decomposition algorithm we will implement works by finding orthogonal matrices H_i such that $H_n \cdots H_1 A = R$, where R is upper triangular. The H_i matrices are called *Householder* matrices, and what they do is to selectively zero out elements of a column of A . We'll see how to compute these matrices later on. A straightforward algorithm would then be one that repeatedly computed and applied Householder transformations to the original matrix. This is exactly the algorithm we implement here until necessity arrives for pivoting.

`<QR prototypes 226> ≡`

```
1180 template<class T, template<class> class storage>
1181 matrix<T, unstructured, storage> &decompose(matrix<T, unstructured, storage> *A)
1182 {
1183     matrix<T, unstructured, dense> v(A->rows(), 1);
1184     matrix<T, unstructured, dense> w(A->cols(), 1);
1185     T beta;
1186     for (index j ← 1; j ≤ A->cols(); ++j) {
1187         <Compute Householder vector for column j 227>;
1188         <Apply Householder transformation 228>;
1189     }
1190     return *A;
1191 }
```

See also section 229.

This code is used in section 225.

227. A Householder reflection can in fact be represented by a single vector v and a scalar β . In fact, the definition of a Householder matrix is any $n \times n$ matrix P of the form

$$I - \frac{2}{v^T v} v v^T.$$

Synonyms are Householder reflection, Householder matrix. The vector v is called a *Householder vector*. Householder matrices are easily verified to be symmetric and orthogonal. If a vector x is multiplied by a Householder matrix, then it is reflected in the hyperplane defined by $\text{span}\{v\}^\perp$. In particular, suppose we have a vector x and want to zero out all but the first component, that is, we want Px to be a multiple of e_1 . After some algebra we can see that

$$v = x \pm P\beta x^0_2 e_1$$

gives the desired transformation. The following piece of code will compute a Householder vector such that $v(1) = 1$, $\beta = 2/v^T v$, and $Px = P\beta x^0_2 e_1$. The normalization of the first entry of the Householder vector is desirable because we can store it with one less component (which enables us to store them directly in the A matrix together with R). Also, we don't use the above formula because of numerical problems. We use instead

$$v_1 = x_1 - P\beta x^0_2 = \frac{-(x_2^2 + \cdots + x_n^2)}{x_1 + P\beta x^0_2}$$

which is numerically more stable when $x_1 > 0$.

⟨ Compute Householder vector for column j 227 ⟩ ≡

```

1192   v.resize(A→rows() - j + 1, 1);
1193   v.entry(1) ← T(0.0);
1194   v.subm(2, v.rows()) ← A→subm(j + 1, A.rows(), j, j);
1195   T sigma ← dot(v, v);
1196   v.entry(1) ← T(1.0);
1197   if (sigma ≡ 0) beta ← 0;
1198   else {
1199     T x ← A→get(j, j);
1200     T mu ← sqrt(sigma + x * x);
1201     v.entry(1) ← (x ≤ 0 ? x - mu : -sigma / (x + mu));
1202     T v1 ← v(1);
1203     beta ← 2 * v1 * v1 / (sigma + v1 * v1);
1204     v /= v1;
1205   }
```

This code is used in section 226.

228. In order to apply the Householder transformation we need to be careful. We actually don't need any matrix-matrix products if we realize that

$$PA = (I - \beta vv^T)A = A - \beta vv^T A,$$

and $vv^T Av^T = v(A^T v)^T$, which consists only of matrix-vector and vector-vector products. Thus, the Householder update requires only a vector-matrix multiplication and an outer product update. In our case, we can also make use of some facts: we know that $v(1) = 1$ and that a block of the A matrix is already zeroed out.

⟨Apply Householder transformation 228⟩ ≡

```
1206   submatrix⟨matrix⟨T, unstructured, storage⟩⟩ Ablock(A, j, A-rows(), j, A-cols());
1207   w ← atmul(Ablock, v, &w);
1208   (void) outerp_update(&Ablock, -beta, v, w);
1209   if (j < A-rows()) A-subm(j + 1, A-rows(), j, j) ← v.subm(2, v.rows());
```

This code is used in section 226.

229. Solving linear equations. We now are in position to solve two important problems. The first is the so-called *least-squares problem*, where we find x that solves $\min \|Ax - b\|_2$, where $A \in \mathbf{R}^{m \times n}$ and $m > n$. In this case the system is *overdetermined*, and an exact solution to $Ax = b$ may not exist.

The second problem occurs when $A \in \mathbf{R}^{m \times n}$ and $n > m$. The system $Ax = b$ either has an infinite number of solutions or none. If it does have one, we compute the *minimum norm* solution.

⟨QR prototypes 226⟩ +≡

```
1210   template⟨class T, template⟨class⟩ class storage⟩
1211   matrix⟨T, unstructured, storage⟩ &solve(matrix⟨T, unstructured, storage⟩
      *A, matrix⟨T, unstructured, dense⟩ *B)
1212   {
1213     index l ← B-cols();
1214     index m ← A-rows();
1215     index n ← A-cols();
1216     matrix⟨T, unstructured, dense⟩ w(l, 1);
1217     if (m ≥ n) {
1218       ⟨Solve least squares problem 230⟩;
1219     }
1220     else {
1221       ⟨Solve minimum norm problem 233⟩;
1222     }
1223     return *B;
1224   }
```

230. Least squares. The least squares problem is solved via a direct matrix decomposition. The solution matrix is guaranteed to be smaller than B and has the same number of columns, so we overwrite B with the solution and resize it accordingly before returning.

⟨Solve least squares problem 230⟩ ≡

```
1225   (void) decompose(A);
1226   ⟨Solve  $Qy = b$  231⟩;
1227   ⟨Solve  $Rx = y$  232⟩;
1228   B-resize(n, l);
```

This code is used in section 229.

231. To solve $Qy = b$ we have to remember that Q is orthogonal, so all we have to do is to compute $y = Q^T b$.

```

1229  <Solve  $Qy = b$  231> ≡
1230      matrix(T, unstructured, storage)  $v(m, 1)$ ;
1231       $v.entry(1) \leftarrow \mathbf{T}(1)$ ;
1232      for (index  $j \leftarrow 1$ ;  $j \leq n$ ;  $++j$ ) {
1233           $v.resize(m - j + 1, 1)$ ;
1234           $v.subm(2, v.rows()) \leftarrow A\text{-}subm(j + 1, m, j, j)$ ;
1235          submatrix(matrix(T, unstructured, dense))  $Bblock(B, j, m, 1, l)$ ;
1236           $w \leftarrow atxmul(Bblock, v, \&w)$ ;
1237           $outerp\_update(\&Bblock, -2/dot(v, v), v, w)$ ;
1238      }

```

This code is used in section 230.

232. The last part is just back substitution with an upper triangular matrix.

```

1239  <Solve  $Rx = y$  232> ≡
1240      for (index  $k \leftarrow 1$ ;  $k \leq l$ ;  $++k$ )
1241          for (index  $i \leftarrow n$ ;  $i \geq 1$ ;  $--i$ ) {
1242               $\mathbf{T} \textit{ inner} \leftarrow 0$ ;
1243              for (index  $j \leftarrow i + 1$ ;  $j \leq n$ ;  $++j$ )  $inner += A\text{-}get(i, j) * B\text{-}get(j, k)$ ;
1244              if ( $\neg finite(B\text{-}entry(i, k) \leftarrow (B\text{-}get(i, k) - inner)/A\text{-}get(i, i))$ ) throw error::rankdeficient();
1245          }

```

This code is used in section 230.

233. Minimum norm. The minimum norm problem is solved through a decomposition of the transposed matrix. If $Ax = b$, then we take the decomposition of A^T and solve the system $(QR)^T x = b$. The matrix dimensions now will be $Q \in \mathbf{R}^{n \times n}$ and $R \in \mathbf{R}^{n \times m}$. The matrix we'll work with is At , which is in $\mathbf{R}^{n \times m}$.

```

1246  <Solve minimum norm problem 233> ≡
1247      matrix(T, unstructured, storage)  $At \leftarrow transpose(*A)$ ;
1248      (void)  $decompose(\&At)$ ;
1249      <Solve  $R^T y = b$  234>;
1250      <Solve  $Q^T x = y$  235>;

```

This code is used in section 229.

234. Our system consists of $R^T Q^T x = b$. We first solve $R^T y = b$. This equation actually translates to $[R_1 \ 0][y_1; y_2] = b$, so we have no means of determining y_2 . This poses no problem since we can determine y_1 and that's all we need. We'll overwrite y_1 on B .

```

1251  <Solve  $R^T y = b$  234> ≡
1252      for (index  $k \leftarrow 1$ ;  $k \leq l$ ;  $++k$ )
1253          for (index  $i \leftarrow 1$ ;  $i \leq m$ ;  $++i$ ) {
1254               $\mathbf{T} \textit{ inner} \leftarrow 0$ ;
1255              for (index  $j \leftarrow 1$ ;  $j < i$ ;  $++j$ )  $inner += At(j, i) * B\text{-}get(j, k)$ ;
1256              if ( $\neg finite(B\text{-}entry(i, k) \leftarrow (B\text{-}get(i, k) - inner)/At(i, i))$ ) throw error::rankdeficient();
1257          }

```

This code is used in section 233.

235. Now we use the fact that Q is orthogonal, so that $Q^T x = y$ translates to $x = Qy$. The only problem is that we don't have y , just y_1 , so that we can't use the *outerp_update* function here because the matrix dimensions wouldn't match.

⟨Solve $Q^T x = y$ 235⟩ ≡

```

1254     B-resize(n, l);
1255     B-subm(m + 1, n, 1, l) ← matrix⟨T, unstructured, storage⟩(n - m, l);
1256     matrix⟨T, unstructured, storage⟩ v(n, 1);
1257     v.entry(1) ← T(1);
1258     for (index k ← m; k ≥ 1; --k) {
1259         index nrows ← n - k + 1;
1260         v.resize(nrows, 1);
1261         v.subm(2, nrows) ← At.subm(k + 1, n, k, k);
1262         submatrix⟨matrix⟨T, unstructured, dense⟩⟩ Bblock(B, k, n, 1, l);
1263         w ← atxmul(Bblock, v, &w);
1264         outerp_update(&Bblock, -2/dot(v, v), v, w);
1265     }
```

This code is used in section 233.

236. Matrix creation functions. We have now the power of the definitions, but using only the basics is unhandy. If we want to create an identity matrix, for example, we should not expect to have to set the diagonal elements to 1 manually. In this part we define functions to create matrices of common use.

237. Eye. Following MATLAB syntax, the function that creates the identity matrix is *eye*. How nice.

```
/* Empty, waiting for export */
```

238. `<eye.h 238>` ≡

```
1266 #ifndef __MATH_EYE__
1267 #define __MATH_EYE__ 1.0
1268 #include <algorithm> /* For min. */
1269 #include <math/math.h>
1270 namespace math {
1271     template<matrix_simple_template>
1272     matrix<T, structure, storage> eye(const math::index rows, const math::index cols)
1273     {
1274         matrix<T, structure, storage> dest(rows, cols);
1275         for (math::index i ← 1; i ≤ min(rows, cols); ++i) dest.entry(i, i) ← T(1);
1276         return dest;
1277     }
1278 }
1279 #endif
```


239. Ones. Following MATLAB syntax, the function that creates the matrix with ones all around is *ones*. How nice, again. Since we have a matrix constructor that sets values, the function itself is nothing more than a matrix creation. For people used with MATLAB (or Scilab, or Octave, and so on) the use of *ones* in a program may seem more intuitive.

```
/* Empty, waiting for export */
```

240. `<ones.h 240>` ≡

```
1280 #ifndef __MATH_ONES__
1281 #define __MATH_ONES__ 1.0
1282 #include <math/math.h>
1283 namespace math {
1284     template<matrix_simple_template>
1285     matrix<T, structure, storage> ones(const math::index rows, const math::index cols)
1286     {
1287         return matrix<T, structure, storage>(rows, cols, T(1));
1288     }
1289 }
1290 #endif
```

241. Matrix functions.

242. Determinant. The determinant of a matrix $A \in \mathbf{R}^{n \times n}$ is given by

$$\det(A) = \sum_{i=j}^n (-1)^{j+1} a_{1j} \det(A_{1j}),$$

where A_{1j} is an $(n-1)$ -by- $(n-1)$ matrix obtained by deleting the first row and j th column of A . Computing the determinant this way would require $O(n!)$ operations, which is unacceptable. Fortunately, we have $\det(AB) = \det(A)\det(B)$, and for an upper or lower triangular matrix $\det(A) = \prod a_{ii}$. With this in mind we are able to compute the determinant in $2n^3/3$ operations via the LU decomposition: we have $\det(A) = \det(LU) = \det(L)\det(U) = \det(U) = \prod u_{ii}$. The only trick is that we have a *permuted* version, so we have to take into account the number of permutations. The LU decomposition function returns the necessary information.

```
/* Empty, waiting for export */
```

243. `<det.h 243> ≡`

```
1291 #ifndef __MATH_DET__
1292 #define __MATH_DET__ 1.00
1293 #include <math/math.h>
1294 #include <math/lu.h>
1295 namespace math {
1296     template<matrix_simple_template>
1297     T det(const matrix<T, structure, storage> &A)
1298     {
1299         matrix<T, unstructured, storage> aux ← A;
1300         vector<math::index> pivots;
1301         T determinant ← 0;
1302         try {
1303             determinant ← math::lu::decompose(&aux, &pivots);
1304             math::index i ← 0;
1305             while (++i ≤ aux.rows()) determinant *= aux(i, i)
1306         }
1307         catch(math::error::singular e) {}
1308         return determinant;
1309     }
1310 }
1311 #endif
```

244. Functions. Functions are things that take a matrix as argument and return a matrix. We will define them before functionals (which take vectors and return scalars) because, for optimization purposes, functionals are more useful but they need functions in order to be of use. Here we define the basic interface in a base class from which specific functionals will be derived.

```
/* Not until export */
```

```

1312 245. <functionbase.h 245> ≡
1312 #ifndef __MATH_FUNCTION__
1313 #define __MATH_FUNCTION__ 1.00
1314 #include <math/math.h>
1315 #include <math/symmetric.h>
1316 #include <math/sparse.h>
1317     namespace math {
1318         namespace function
1319         {
1320             template<class T>
1321             class base {
1322             public:
1323                 <Function base class methods 246>
1324             };
1325         }
1326     }
1327 #endif

```

246. The first thing we do is to declare a virtual destructor, so that we can derive classes from the base class.

```

1328 <Function base class methods 246> ≡
1328     virtual ~base(void) {}

```

See also sections 247, 248, and 249.

This code is used in section 245.

247. Before defining the interface we need to consider some points about functions (this rant will be repeated when defining functionals): first, a function should be able to get any kind of matrix as the point, that is, the x in $f(x)$ could be sparse, dense and so on. Also, we should be able to create a list containing various types of functions. The impossibility of this ideal situation is summarized by Stroustrup: “A member template cannot be virtual.” This is a design decision, and a wrong one in my point of view. We have then only two possibilities: either we make the **matrix** class derived, so we can pass pointers, or we fix the type of the parameter. The first option results in performance degradation, so we’re stuck with the second. Fortunately most functions take vectors as arguments, so that matrix structures are not a big deal. We will assume that **dense** vectors are the best compromise. Also, we fix the return value to a **dense** and **unstructured** matrix because we consider that in most cases we’ll be returning vectors. Storage for the result must be provided by the user.

```

1329 <Function base class methods 246> +≡
1329     matrix<T, unstructured, dense> &operator()(const matrix<T, unstructured, dense>
1330         &x, matrix<T, unstructured, dense> *dest)
1330     { return eval(x, dest); }
1331     virtual matrix<T, unstructured, dense> &eval(const matrix<T, unstructured, dense>
1331         &x, matrix<T, unstructured, dense> *dest) ← 0;

```

248. If possible, a function should compute the Jacobian. If $f : \mathbf{R}^n \rightarrow \mathbf{R}^m$, then the Jacobian is a matrix $J \in \mathbf{R}^{m \times n}$ defined by $J_{ij} = df_i/dx_j$. Here we face another interface decision: functions are supposed to be used to compute lots of points, that is, when you create a function, usually you want to evaluate it in a set of points, not just only one. If a evaluation method created storage for the result each time it was called, then we would be facing a serious performance threat. For this reason we require the user to pass the result variable as an argument. This requirement, in turn, makes the type of the result matrix fixed for the same reasons discussed above. The method returns the result (that is, *dest*), which is a useful behavior in certain applications, as for example in function compositions like $f(g(x))$. The size of the *dest* matrix at input *should be checked by the method*. This will be normally performed via the `matrix::resize` method. Since this method checks for same dimensions, the performance does not suffer too much.

⟨Function base class methods 246⟩ +≡

```
1332     virtual matrix⟨T, unstructured, dense⟩ &jacobian(const matrix⟨T, unstructured, dense⟩
           &x, matrix⟨T, unstructured, dense⟩ *dest) ← 0;
```

249. The same thing is valid for the Hessian, except here we require the Hessian to be symmetric (which it is by definition). The same remarks about *dest* size made above are valid here. We provide two methods, one for **sparse** matrices (for large optimization problems it may be crucial to get a sparse version).

⟨Function base class methods 246⟩ +≡

```
1333     virtual matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
           &x, matrix⟨T, symmetric, dense⟩ *dest, const index i, const index j ← 1) ← 0;
1334     virtual matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
           &x, matrix⟨T, symmetric, sparse⟩ *dest, const index i, const index j ← 1) ← 0;
```

250. The gaxpy function. Our first functional is arguably the most simple. We define a function that computes the scalar *gaxpy* operation, that is, $f(x) = Ax + b$. This is also one of the most useful functions for optimization – think of $Ax \prec b$.

```
/* Empty, waiting for export */
```

251. ⟨function/gaxpy.h 251⟩ ≡

```
1335 #ifndef __MATH_GAXPY_FUNCTION__
1336 #define __MATH_GAXPY_FUNCTION__ 1.00
1337 #include <math/math.h>
1338 #include <math/algebra.h>
1339 #include <math/functionbase.h>
1340     namespace math {
1341         namespace function
1342         {
1343             template⟨matrix_simple_template⟩
1344             class gaxpy:public base⟨T⟩ {
1345                 ⟨Gaxpy function internal variables 252⟩
1346             public:
1347                 ⟨Gaxpy function class methods 253⟩
1348             };
1349         }
1350     }
1351 #endif
```

252. To compute a **gaxpy** operation we need two parameters, the A and b matrices. We assume b will be a vector most times.

⟨Gaxpy function internal variables 252⟩ ≡

```
1352     matrix⟨T, structure, storage⟩ ay;
1353     matrix⟨T, unstructured, dense⟩ bee;
```

This code is used in section 251.

253. We provide a means to modify these values in two ways: by returning a reference to them and at the time of construction.

⟨Gaxpy function class methods 253⟩ ≡

```
1354     matrix⟨T, structure, storage⟩ &A(void) { return ay; }
1355     matrix⟨T, unstructured, dense⟩ &b(void) { return bee; }
1356     gaxpy(const matrix⟨T, structure, storage⟩ &a, const T &B)
1357     : ay(a), bee(B) { }
```

See also sections 254, 255, and 256.

This code is used in section 251.

254. Now to the fun stuff. Evaluating is a simple task of calling some algebra functions.

⟨Gaxpy function class methods 253⟩ +≡

```
1358     virtual matrix⟨T, unstructured, dense⟩ &eval(const matrix⟨T, unstructured, dense⟩
1359           &x, matrix⟨T, unstructured, dense⟩ *dest)
           { return (gaxpy(ay, x, &(dest-copyfrom(b)))); }
```

255. The Jacobian is simply A .

⟨Gaxpy function class methods 253⟩ +≡

```
1360     virtual matrix⟨T, unstructured, dense⟩ &jacobian(const matrix⟨T, unstructured, dense⟩
1361           &x, matrix⟨T, unstructured, dense⟩ *dest)
1362     {
1363         return (*dest) ← ay;
1364     }
```

256. The hessian of an affine function is zero.

⟨Gaxpy function class methods 253⟩ +≡

```
1364     virtual matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
1365           &x, matrix⟨T, symmetric, dense⟩ *dest, const index, const index)
1366     {
1367         dest->resize(x.rows(), x.rows());
1368         dest->fillwith(0);
1369         return *dest;
1370     }
1370     virtual matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
1371           &x, matrix⟨T, symmetric, sparse⟩ *dest, const index, const index)
1372     {
1373         *dest ← matrix⟨T, symmetric, sparse⟩(x.rows(), x.rows());
1374         return *dest;
1375     }
```

257. Functionals. Functionals are the MATH library representation of functionals. In the most simple situation, functionals take arguments and return the value of a function at some point. In other situations, such as optimization algorithms, we may want the gradient and/or the Hessian of the function at some point. Here we define the basic interface in a base class from which specific functionals will be derived.

```
/* Not until export */
```

258. `<functionalbase.h 258>` \equiv

```
1375 #ifndef __MATH_FN__
1376 #define __MATH_FN__ 1.00
1377 #include <math/math.h>
1378 #include <math/symmetric.h>
1379 #include <math/sparse.h>
1380 namespace math {
1381     namespace functional
1382     {
1383         template<class T>
1384         class base {
1385         public:
1386             <Functional base class methods 259>
1387         };
1388     }
1389 }
1390 #endif
```

259. The first thing we do is to declare a virtual destructor, so that we can derive classes from the base class.

`<Functional base class methods 259>` \equiv

```
1391 virtual ~base(void) {}
```

See also sections 260, 261, and 262.

This code is used in section 258.

260. Before defining the interface we need to consider some points about functionals: first, a functional should be able to get any kind of matrix as the point, that is, the x in $f(x)$ could be sparse, dense and so on. Also, we should be able to create a list containing various types of functionals. The impossibility of this ideal situation is summarized by Stroustrup: “A member template cannot be virtual.” This is a design decision, and a wrong one in my point of view. We have then only two possibilities: either we make the **matrix** class derived, so we can pass pointers, or we fix the type of the parameter. The first option results in performance degradation, so we’re stuck with the second. Fortunately most functionals take vectors as arguments, so that matrix structures are not a big deal. We will assume that **dense** vectors are the best compromise.

`<Functional base class methods 259>` $+ \equiv$

```
1392 T operator()(const matrix<T, unstructured, dense> &x) { return eval(x); }
1393 virtual T eval(const matrix<T, unstructured, dense> &x) ← 0;
```

261. If possible, a functional should compute the gradient at a point. Here we face another interface decision: functionals are supposed to be used to compute lots of points, that is, when you create a functional, usually you want to evaluate it in a set of points, not just only one. If a evaluation method created storage for the result each time it was called, then we would be facing a serious performance threat. For this reason we require the user to pass the result variable as an argument. This requirement, in turn, makes the type of the result matrix fixed for the same reasons discussed above. The method returns the result (that is, *dest*), which is a useful behavior in certain applications, as for example in function compositions like $f(g(x))$. The size of the *dest* matrix at input *should be checked by the method*. This will be normally performed via the `matrix::init` method. Since this method checks for same dimensions, the performance does not suffer too much.

⟨Functional base class methods 259⟩ +≡

```
1394     virtual matrix⟨T, unstructured, dense⟩ &grad(const matrix⟨T, unstructured, dense⟩
        &x, matrix⟨T, unstructured, dense⟩ &dest) ← 0;
```

262. The same thing is valid for the Hessian, except here we require the Hessian to be symmetric (which it is by definition). The same remarks about *dest* size made above are valid here. We provide two methods, one for `sparse` matrices (for large optimization problems it may be crucial to get a sparse version).

⟨Functional base class methods 259⟩ +≡

```
1395     virtual matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
        &x, matrix⟨T, symmetric, dense⟩ &dest) ← 0;
```

```
1396     virtual matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
        &x, matrix⟨T, symmetric, sparse⟩ &dest) ← 0;
```

263. The gaxpy functional. Our first functional is arguably the most simple. We define a functional that computes the scalar `gaxpy` operation, that is, $f(x) = a^T x + b$.

```
/* Empty, waiting for export */
```

264. ⟨functional/gaxpy.h 264⟩ ≡

```
1397 #ifndef __MATH_GAXPY_FUNCTIONAL__
1398 #define __MATH_GAXPY_FUNCTIONAL__ 1.00
1399 #include <math/algebra.h>
1400 #include <math/functionalbase.h>
1401     namespace math {
1402         namespace functional
1403         {
1404             template⟨matrix_simple_template⟩
1405             class gaxpy :public base⟨T⟩ {
1406                 ⟨Gaxpy functional internal variables 265⟩
1407             public:
1408                 ⟨Gaxpy functional class methods 266⟩
1409             };
1410         }
1411     }
1412 #endif
```

265. To compute a `gaxpy` operation we need two parameters, the *a* vector and the scalar *b*.

⟨Gaxpy functional internal variables 265⟩ ≡

```
1413     matrix⟨T, structure, storage⟩ ay;
1414     T bee;
```

This code is used in section 264.

266. We provide a means to modify these values in two ways: by returning a reference to them and at the time of construction.

```

1415 <Gaxpy functional class methods 266> ≡
      matrix<T, structure, storage> &a(void) { return ay; }
1416   T &b(void) { return bee; }
1417   gaxpy(const matrix<T, structure, storage> &A, const T &B)
1418   : ay(A), bee(B) { }
```

See also sections 267, 268, and 269.

This code is used in section 264.

267. Now to the fun stuff. Evaluating is a simple task of calling some algebra functions.

```

1419 <Gaxpy functional class methods 266> +=
      virtual T eval(const matrix<T, unstructured, dense> &x)
1420   { return dot(ay, x) + bee; }
```

268. The gradient is simply a .

```

1421 <Gaxpy functional class methods 266> +=
      virtual matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
      &x, matrix<T, unstructured, dense> &dest)
1422   { return (dest ← ay); }
```

269. The hessian of an affine function is zero.

```

1423 <Gaxpy functional class methods 266> +=
      virtual matrix<T, symmetric, dense> &hess(const matrix<T, unstructured, dense>
      &x, matrix<T, symmetric, dense> &dest)
1424   {
1425     dest.init(x.rows(), x.rows());
1426     dest.fillwith(0);
1427     return dest;
1428   }
1429   virtual matrix<T, symmetric, sparse> &hess(const matrix<T, unstructured, dense>
      &x, matrix<T, symmetric, sparse> &dest)
1430   {
1431     dest ← matrix<T, symmetric, sparse>(x.rows(), x.rows());
1432     return dest;
1433   }
```

270. The linear combination of functionals. Our next definition is a functional that is the weighed sum of other specified functionals. It will be useful, for example, in optimization methods (when summing barrier functions).

```
/* Empty, waiting for export */
```



```

271. <functional/linear.h 271> ≡
1434 #ifndef __MATH_LINEAR_FUNCTIONAL__
1435 #define __MATH_LINEAR_FUNCTIONAL__ 1.00
1436 #include <algorithm>
1437 #include <math/functionalbase.h>
1438 namespace math {
1439     namespace functional
1440     {
1441         template<class T>
1442         class linear:public base<T> {
1443             <Linear combination of functionals internal variables 272>
1444         public:
1445             <Linear combination of functionals class methods 273>
1446         };
1447     }
1448 }
1449 #endif

```

272. The functionals to be added are stored in a vector of pairs of weights and functionals.

```

<Linear combination of functionals internal variables 272> ≡
1450 typedef std::pair<T, base<T> *> element_type;
1451 vector<element_type> elements;
1452 typedef typename vector<element_type>::iterator iterator;

```

See also sections 278 and 280.

This code is used in section 271.

273. Adding elements to the sum is a matter of passing weights and functional pointers to the *add* function. We also allow the user to specify initial values at the time of construction.

```

<Linear combination of functionals class methods 273> ≡
1453 linear(void) { }
1454 linear(base<T> *element, T weight ← 1) { add(element, weight); }
1455 linear(const vector<element_type> &elem) { elements ← elem; }
1456 void add(base<T> *element, T weight ← 1)
1457 {
1458     if (weight ≡ 0) return;
1459     elements.push_back(element_type(weight, element));
1460 }
1461 void add(const linear<T> *element, T weight ← 1)
1462 {
1463     for (int i ← 0; i ≠ element->size(); ++i) add(element->get_term(i), weight * element->get_weight(i));
1464 }

```

See also sections 274, 275, 276, 277, 279, and 281.

This code is used in section 271.

274. We can remove functionals from the end of the list.

```

<Linear combination of functionals class methods 273> +=
1465 void pop_back(void)
1466 {
1467     elements.pop_back();
1468 }

```

275. We will need to modify weights

```

1469 <Linear combination of functionals class methods 273> +=
      void set_weight(int i, T value) { elements[i].first ← value; }
1470   T get_weight(int i) const { return elements[i].first; }
1471   void set_term(int i, base<T> *value) { elements[i].second ← value; }
1472   base<T> *get_term(int i) const { return elements[i].second; }
1473   int size(void) const { return elements.size(); }

```

276. The pointers will be erased, but not the functionals they point to, when the **linear** is deleted.

```

1474 <Linear combination of functionals class methods 273> +=
      ~linear(void) { elements.erase(elements.begin(), elements.end()); }

```

277. Now to evaluation. Pretty simple, as you may expect. The default behavior is to return zero if there are no functionals in the list.

```

1475 <Linear combination of functionals class methods 273> +=
      T eval(const matrix<T, unstructured, dense> &x)
1476   {
1477     T result ← 0;
1478     for (iterator i ← elements.begin(); i ≠ elements.end(); ++i) result += i->first * i->second->eval(x);
1479     return result;
1480   }

```

278. For the gradient, the default when no elements are present is also to return zero. We set up an internal variable *aux* so that we don't need to create/resize vectors each time a new gradient is computed. Also, we test for unity weights in order to save computation.

```

1481 <Linear combination of functionals internal variables 272> +=
      matrix<T, unstructured, dense> aux;

```

```

1482 <Linear combination of functionals class methods 273> +=
      matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
1483     &x, matrix<T, unstructured, dense> &dest)
1484   {
1485     if (elements.size() ≡ 0) {
1486       dest.resize(x.rows(), 1);
1487       dest.fillwith(0.0);
1488       return dest;
1489     }
1490     iterator i ← elements.begin();
1491     i->second->grad(x, dest);
1492     if (i->first ≠ 1) dest *= i->first;
1493     for (i++; i ≠ elements.end(); ++i) {
1494       i->second->grad(x, aux);
1495       if (i->first ≠ 1) aux *= i->first;
1496       dest += aux;
1497     }
1498     return dest;

```

280. We follow the same algorithm for the Hessian, again testing for unity weights in order to save computation.

⟨Linear combination of functionals internal variables 272⟩ +≡

```
1499   matrix⟨T, symmetric, dense⟩ Hdense;
1500   matrix⟨T, symmetric, sparse⟩ Hsparse;
```

281. ⟨Linear combination of functionals class methods 273⟩ +≡

```
1501   matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
1502   {
1503     if (elements.size() ≡ 0) {
1504       dest.resize(x.rows(), x.rows());
1505       dest.fillwith(0.0);
1506       return dest;
1507     }
1508     iterator i ← elements.begin();
1509     i→second→hess(x, dest);
1510     if (i→first ≠ 1) dest *= i→first;
1511     for (i++; i ≠ elements.end(); ++i) {
1512       i→second→hess(x, Hdense);
1513       if (i→first ≠ 1) Hdense *= i→first;
1514       dest += Hdense;
1515     }
1516     return dest;
1517   }
1518   matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
1519   {
1520     if (elements.size() ≡ 0) {
1521       dest ← matrix⟨T, symmetric, sparse⟩(x.rows(), x.rows());
1522       return dest;
1523     }
1524     iterator i ← elements.begin();
1525     i→second→hess(x, dest);
1526     if (i→first ≠ 1) dest *= i→first;
1527     for (i++; i ≠ elements.end(); ++i) {
1528       i→second→hess(x, Hsparse);
1529       if (i→first ≠ 1) Hsparse *= i→first;
1530       dest += Hsparse;
1531     }
1532     return dest;
1533   }
```

282. The quadratic functional. The next functional we implement is the quadratic. It computes the value of $f(x) = x^T P x + p^T x + \pi$. We will require P to be symmetric. If it is not, it is always possible to find a new P that results in the same functional. To the code:

```
/* Empty, waiting for export */
```

```

283. <functional/quadratic.h 283> ≡
1534 #ifndef __MATH_QUADRATIC_FUNCTIONAL__
1535 #define __MATH_QUADRATIC_FUNCTIONAL__ 1.00
1536 #include <math/algebra.h>
1537 #include <math/functionalbase.h>
1538 namespace math {
1539     namespace functional
1540     {
1541         template<class T, template<class> class storage ← dense>
1542         class quadratic:public base<T> {
1543             <Quadratic functional internal variables 284>
1544         public:
1545             <Quadratic functional methods 285>
1546         };
1547     }
1548 }
1549 #endif

```

284. We first provide storage for the functional parameters. As a design decision, the vector p will be dense and unstructured.

```

<Quadratic functional internal variables 284> ≡
1550     matrix<T, symmetric, storage> Pee;
1551     matrix<T, unstructured, dense> pee;
1552     T Pi;

```

This code is used in section 283.

285. We provide means to change the parameters in the same way we did with the **gaxpy** functional.

```

<Quadratic functional methods 285> ≡
1553     matrix<T, symmetric, storage> &P(void) { return Pee; }
1554     matrix<T, unstructured, dense> &p(void) { return pee; }
1555     T &pi(void) { return Pi; }
1556     quadratic(const matrix<T, symmetric, storage> &newP, const matrix<T, unstructured, dense>
1557               &newp, const T &newpi)
: Pee(newP), pee(newp), Pi(newpi) {}

```

See also sections 286, 287, and 288.

This code is used in section 283.

286. Evaluating: we use the *axmul* function for now, but it's a good idea to change it in the future. Probably the best thing to do is to define a good $*$ operator, but I don't have one right now.

```

<Quadratic functional methods 285> +=
1558     virtual T eval(const matrix<T, unstructured, dense> &x)
1559     {
1560         T result ← dot(pee, x) + Pi;
1561         matrix<T, unstructured, dense> y(Pee.rows(), 1);
1562         y ← axmul(Pee, x, &y);
1563         result += dot(x, y);
1564         return result;
1565     }

```

287. The gradient is given by $2Px + p$.

```

1566 <Quadratic functional methods 285> +≡
      virtual matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
1567         &x, matrix<T, unstructured, dense> &dest)
1568     {
1569         return gaxpy(Pee, x * 2, &(dest.copyfrom(pee)));
1570     }

```

288. And the Hessian is simply $2P$.

```

1570 <Quadratic functional methods 285> +≡
      virtual matrix<T, symmetric, dense> &hess(const matrix<T, unstructured, dense>
1571         &x, matrix<T, symmetric, dense> &dest)
1572     {
1573         dest ← Pee;
1574         return dest *= 2;
1575     }
1576     virtual matrix<T, symmetric, sparse> &hess(const matrix<T, unstructured, dense>
1577         &x, matrix<T, symmetric, sparse> &dest)
1578     {
1579         dest ← Pee;
1580         return dest *= 2;
1581     }

```

289. The norm-2 error. This functional computes, for a given function f and a vector y , the value of

$$\|y_i - f_i(x)\|^2.$$

This is an extremely useful functional for optimization problems: suppose you trying to approximate some set of data y to a function f that depends on the *parameters* x – then minimizing this functional is what you want to do. Note the confusing naming (of which we cannot escape). Normally the set of data is (x_i, y_i) , but in our case the x_i is stored in the function f itself – they are not the argument to f .

/ Empty, waiting for export. */*

```

290. <functional/norm2err.h 290> ≡
1580 #ifndef __MATH_NORM2ERR_FUNCTIONAL__
1581 #define __MATH_NORM2ERR_FUNCTIONAL__
1582 #include <math/algebra.h>
1583 #include <math/function.h>
1584 #include <math/functionalbase.h>
1585 namespace math {
1586     namespace functional
1587     {
1588         template<class T, template<class> class storage>
1589         class norm2err:public base<T> {
1590             <Norm-2 error functional internal variables 291>
1591         public:
1592             <Norm-2 error functional class methods 292>
1593         };
1594     }
1595 }
1596 #endif

```

291. The internal matrix variables are the vector y , two auxiliary vectors for computing the value and gradient of the function, and an auxiliary matrix for computing the Hessian. The other internal variable is the approximating function.

```

1597  <Norm-2 error functional internal variables 291> ≡
1598      matrix<T, unstructured, dense>  $Y$ ,  $aux$ ,  $fval$ ;
1599      matrix<T, symmetric, storage>  $H$ ;
1599      function :: base<T> * $F$ ;

```

This code is used in section 290.

292. We can change all the values at any time or at construction.

```

1600  <Norm-2 error functional class methods 292> ≡
1601      matrix<T, unstructured, dense> & $y$ (void) { return  $Y$ ; }
1601      function :: base<T> *& $f$ (void) { return  $F$ ; }
1602      norm2err(const matrix<T, unstructured, dense> & $newy$ , function :: base<T> * $newf$  ← 0)
1603      :  $Y(newy)$ ,  $F(newf)$  {}
1604      norm2err(function :: base<T> * $newf$  ← 0)
1605      :  $F(newf)$  {}

```

See also sections 293, 294, 295, and 296.

This code is used in section 290.

293. Now to the eval function.

```

1606  <Norm-2 error functional class methods 292> +≡
1607      virtual T eval(const matrix<T, unstructured, dense> & $x$ )
1608      {
1609           $fval$  ←  $F$ -eval( $x$ , & $fval$ );
1609           $fval$  -=  $Y$ ;
1610          return dot( $fval$ ,  $fval$ );
1611      }

```

294. The gradient is given by

$$2Df(x)^T(f - y)$$

```

1612  <Norm-2 error functional class methods 292> +≡
1613      virtual matrix<T, unstructured, dense> & $grad$ (const matrix<T, unstructured, dense>
1614          & $x$ , matrix<T, unstructured, dense> & $dest$ )
1615      {
1616           $fval$  ←  $F$ -eval( $x$ , & $fval$ );
1617           $fval$  -=  $Y$ ;
1618           $F$ -jacobian( $x$ , & $aux$ );
1619           $atxmul$ ( $aux$ ,  $fval$ , & $dest$ );
1620           $dest$  *= 2;
1621          return  $dest$ ;
1622      }

```

295. The Hessian is a little bit more complicated. It is given by

$$2 \sum_i (f_i(x) - y_i) \nabla^2 f_i(x) + 2Df(x)^T Df(x).$$

We have all the necessary algebraic functions already defined, though.

⟨Norm-2 error functional class methods 292⟩ +≡

```

1621   virtual matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
1622   {
1623     outterp(F→jacobian(x, &aux), &dest);
1624     fval ← F→eval(x, &fval);
1625     fval -= Y;
1626     for (index i ← 1; i ≤ Y.rows(); ++i) {
1627       H ← F→hess(x, &H, i);
1628       H *= fval(i);
1629       dest += H;
1630     }
1631     dest *= 2;
1632     return dest;
1633   }

```

296. ⟨Norm-2 error functional class methods 292⟩ +≡

```

1634   virtual matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
1635   {
1636     outterp(F→jacobian(x, aux), &dest);
1637     fval ← F→eval(x, &fval);
1638     fval -= Y;
1639     for (index i ← 1; i ≤ Y.rows(); ++i) {
1640       H ← F→hess(x, &H, i);
1641       H *= fval(i);
1642       dest += H;
1643     }
1644     dest *= 2;
1645     return dest;
1646   }

```

297. The product of two functionals.

/* Empty, waiting for **export** */

```

298.  ⟨functional/prod.h 298⟩ ≡
1647  #ifndef __MATH_PROD_FUNCTIONAL__
1648  #define __MATH_PROD_FUNCTIONAL__ 1.00
1649  #include <math/functionalbase.h>
1650  namespace math {
1651  namespace functional
1652  {
1653  template<class T>
1654  class prod:public base<T> {
1655  ⟨Product of functionals internal variables 299⟩
1656  public:
1657  ⟨Product of functionals class methods 300⟩
1658  };
1659  }
1660 }
1661 #endif

```

```

299.  ⟨Product of functionals internal variables 299⟩ ≡
1662  typedef base<T> *element_type;
1663  element_type f, g;

```

See also sections 302 and 304.

This code is used in section 298.

```

300.  ⟨Product of functionals class methods 300⟩ ≡
1664  prod(element_type new-f, element_type new-g):f(new-f),g(new-g) { }

```

See also sections 301, 303, and 305.

This code is used in section 298.

```

301.  ⟨Product of functionals class methods 300⟩ +=
1665  T eval(const matrix<T, unstructured, dense> &x) { return f->eval(x) * g->eval(x); }

```

302. The gradient of the product is given by $f\nabla g + g\nabla f$.

```

⟨Product of functionals internal variables 299⟩ +=
1666  matrix<T, unstructured, dense> aux;

```

```

303.  ⟨Product of functionals class methods 300⟩ +=
1667  matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
    &x, matrix<T, unstructured, dense> &dest)
1668  {
1669  f->grad(x, aux);
1670  aux *= g->eval(x);
1671  g->grad(x, dest);
1672  dest *= f->eval(x);
1673  return dest += aux;
1674  }

```

304. The Hessian is given by $g\nabla^2 f + \nabla f\nabla^T g + \nabla g\nabla^T f + f\nabla^2 g$.

```

⟨Product of functionals internal variables 299⟩ +=
1675  matrix<T, unstructured, dense> xua;
1676  matrix<T, symmetric, dense> Hdense;
1677  matrix<T, symmetric, sparse> Hsparse;

```


305. \langle Product of functionals class methods 300 $\rangle + \equiv$

```

1678   matrix $\langle$ T, symmetric, dense $\rangle$  &hess(const matrix $\langle$ T, unstructured, dense $\rangle$ 
      &x, matrix $\langle$ T, symmetric, dense $\rangle$  &dest)
1679   {
1680     f-hess(x, dest);
1681     dest *= g-eval(x);
1682     g-hess(x, Hdense);
1683     Hdense *= f-eval(x);
1684     dest += Hdense;
1685     f-grad(x, aux);
1686     g-grad(x, xua);
1687     dest += xyyx(aux, xua, &Hdense);
1688     return dest;
1689   }
1690   matrix $\langle$ T, symmetric, sparse $\rangle$  &hess(const matrix $\langle$ T, unstructured, dense $\rangle$ 
      &x, matrix $\langle$ T, symmetric, sparse $\rangle$  &dest)
1691   {
1692     f-hess(x, dest);
1693     dest *= g-eval(x);
1694     g-hess(x, Hsparse);
1695     Hsparse *= f-eval(x);
1696     dest += Hsparse;
1697     f-grad(x, aux);
1698     g-grad(x, xua);
1699     dest += xyyx(aux, xua, &Hsparse);
1700     return dest;
1701   }

```

306. The ratio of two functionals.

```

/* Empty, waiting for export */

```

307. \langle functional/ratio.h 307 $\rangle \equiv$

```

1702 #ifndef __MATH_RATIO_FUNCTIONAL__
1703 #define __MATH_RATIO_FUNCTIONAL__ 1.00
1704 #include <math/functionalbase.h>
1705 namespace math {
1706     namespace functional
1707     {
1708         template<class T>
1709         class ratio:public base<T> {
1710              $\langle$ Ratio of functionals internal variables 308 $\rangle$ 
1711         public:
1712              $\langle$ Ratio of functionals class methods 309 $\rangle$ 
1713         };
1714     }
1715 }
1716 #endif

```

308. ⟨Ratio of functionals internal variables 308⟩ ≡

```
1717   typedef base<T> *element_type;
1718   element_type f, g;
```

See also sections 310 and 312.

This code is used in section 307.

309. ⟨Ratio of functionals class methods 309⟩ ≡

```
1719   ratio(element_type new_f, element_type new_g):f(new_f),g(new_g) { }
1720   T eval(const matrix<T, unstructured, dense> &x) { return f->eval(x)/g->eval(x); }
```

See also sections 311 and 313.

This code is used in section 307.

310. The gradient of the ratio is given by $\nabla f/g - f\nabla g/g^2$.

⟨Ratio of functionals internal variables 308⟩ +=

```
1721   matrix<T, unstructured, dense> aux;
```

311. ⟨Ratio of functionals class methods 309⟩ +=

```
1722   matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
      &x, matrix<T, unstructured, dense> &dest)
1723   {
1724     T result ← g->eval(x);
1725     g->grad(x, aux);
1726     aux *= -f->eval(x)/result;
1727     f->grad(x, dest);
1728     dest += aux;
1729     return dest /= result;
1730   }
```

312. The Hessian is given by $\nabla^2 f/g - \nabla g \nabla^T f/g^2 - \nabla f \nabla^T g/g^2 + 2f \nabla g \nabla^T g/g^3 - f \nabla^2 g/g^2$.

⟨Ratio of functionals internal variables 308⟩ +=

```
1731   matrix<T, unstructured, dense> xua;
1732   matrix<T, symmetric, dense> Hdense;
1733   matrix<T, symmetric, sparse> Hsparse;
```

313. ⟨Ratio of functionals class methods 309⟩ +=

```

1734   matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
1735   {
1736     T result ← g-eval(x);
1737     T tuser ← f-eval(x);
1738     f-grad(x, aux);
1739     g-grad(x, xua);
1740     outerp(xua, &dest);
1741     dest *= 2 * tuser / result;
1742     dest -= xyyx(aux, xua, &Hdense);
1743     g-hess(x, Hdense);
1744     Hdense *= tuser;
1745     dest -= Hdense;
1746     dest /= result;
1747     dest += f-hess(x, Hdense);
1748     return dest /= result;
1749   }
1750   matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
1751   {
1752     T result ← g-eval(x);
1753     T tuser ← f-eval(x);
1754     f-grad(x, aux);
1755     g-grad(x, xua);
1756     outerp(xua, &dest);
1757     dest *= 2 * tuser / result;
1758     dest -= xyyx(aux, xua, &Hsparse);
1759     g-hess(x, Hsparse);
1760     Hsparse *= tuser;
1761     dest -= Hsparse;
1762     dest /= result;
1763     dest += f-hess(x, Hsparse);
1764     return dest /= result;
1765   }

```

314. The power functional.

```

/* Empty, waiting for export */

```

```

315. <functional/power.h 315> ≡
1766 #ifndef __MATH_POWER_FUNCTIONAL__
1767 #define __MATH_POWER_FUNCTIONAL__ 1.00
1768 #include <math.h>
1769 #include <math/functionalbase.h>
1770 namespace math {
1771     namespace functional
1772     {
1773         template<class T>
1774         class power:public base<T> {
1775             <Power of a functional internal variables 316>
1776         public:
1777             <Power of a functional class methods 317>
1778         };
1779     }
1780 }
1781 #endif

```

316.

<Power of a functional internal variables 316> ≡

```

1782     typedef base<T> *element_type;
1783     element_type f;
1784     double exponent;

```

See also section 321.

This code is used in section 315.

317.

<Power of a functional class methods 317> ≡

```

1785     power(element_type new_f, double new_e):f(new_f), exponent(new_e) { }

```

See also sections 318, 320, and 322.

This code is used in section 315.

318.

<Power of a functional class methods 317> +≡

```

1786     T eval(const matrix<T, unstructured, dense> &x)
1787     {
1788         T result ← f->eval(x);
1789         return pow(result, exponent);
1790     }

```

319. The gradient of the power is given by $nf^{n-1}\nabla f$.

320. <Power of a functional class methods 317> +≡

```

1791     matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
1792         &x, matrix<T, unstructured, dense> &dest)
1793     {
1794         T result ← f->eval(x);
1795         result ← exponent * pow(result, exponent - 1);
1796         f->grad(x, dest);
1797         return dest *= result;
1798     }

```

321. The Hessian is given by $nf^{n-1}\nabla^2 f + n(n-1)f^{n-2}\nabla f\nabla f^T$.

⟨Power of a functional internal variables 316⟩ +≡

```
1798   matrix⟨T, symmetric, dense⟩ Hdense;
1799   matrix⟨T, symmetric, sparse⟩ Hsparse;
1800   matrix⟨T, unstructured, dense⟩ aux;
```

322. ⟨Power of a functional class methods 317⟩ +≡

```
1801   matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
1802   {
1803     T result ← f-eval(x);
1804     if (result ≡ 0 ∧ exponent ≥ 2) {
1805       dest.resize(x.rows(), x.rows());
1806       dest.fillwith(0.0);
1807       return dest;
1808     }
1809     if (result ≡ 0 ∧ exponent < 2) throw error::domain();
1810     f-grad(x, aux);
1811     outerp(aux, &Hdense);
1812     Hdense *= (exponent - 1)/result;
1813     f-hess(x, dest);
1814     dest += Hdense;
1815     return dest *= pow(result, exponent - 1) * exponent;
1816   }
1817   matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
1818   {
1819     T result ← f-eval(x);
1820     if (result ≡ 0 ∧ exponent ≥ 2) {
1821       dest.resize(x.rows(), x.rows());
1822       dest.fillwith(0.0);
1823       return dest;
1824     }
1825     if (result ≡ 0 ∧ exponent < 2) throw error::domain();
1826     f-grad(x, aux);
1827     outerp(aux, &Hsparse);
1828     Hsparse *= (exponent - 1)/result;
1829     f-hess(x, dest);
1830     dest += Hsparse;
1831     return dest *= pow(result, exponent - 1) * exponent;
1832   }
```

323. The entropy of a functional.

/* Empty, waiting for export */

```

324.  ⟨functional/entr.h 324⟩ ≡
1833  #ifndef __MATH_ENTROPY_FUNCTIONAL__
1834  #define __MATH_ENTROPY_FUNCTIONAL__ 1.00
1835  #include <math.h>
1836  #include <math/functionalbase.h>
1837      namespace math {
1838          namespace functional
1839          {
1840              template<class T>
1841              class entr:public base<T> {
1842                  ⟨Entropy of a functional internal variables 325⟩
1843              public:
1844                  ⟨Entropy of a functional class methods 326⟩
1845              };
1846          }
1847      }
1848  #endif

```

```

325.  ⟨Entropy of a functional internal variables 325⟩ ≡
1849      typedef base<T> *element_type;
1850      element_type f;

```

See also section 330.

This code is used in section 324.

```

326.  ⟨Entropy of a functional class methods 326⟩ ≡
1851      entr(element_type new_f):f(new_f) { }

```

See also sections 327, 329, and 331.

This code is used in section 324.

```

327.  ⟨Entropy of a functional class methods 326⟩ +≡
1852      T eval(const matrix<T, unstructured, dense> &x)
1853      {
1854          T result ← f->eval(x);
1855          if (result ≤ 0) throw error::domain();
1856          return -result * ::log(result);
1857      }

```

328. The gradient of the entropy is given by $-(1 + \log f)\nabla f$.

```

329.  ⟨Entropy of a functional class methods 326⟩ +≡
1858      matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
1859          &x, matrix<T, unstructured, dense> &dest)
1860      {
1861          T result ← f->eval(x);
1862          if (result ≤ 0) throw error::domain();
1863          f->grad(x, dest);
1864          return dest * = - (::log(result) + 1);
1865      }

```

330. The Hessian is given by $-(1 + \log f)\nabla^2 f - (1/f)\nabla f\nabla^T f$.

⟨Entropy of a functional internal variables 325⟩ +≡

```
1865   matrix⟨T, symmetric, dense⟩ Hdense;
1866   matrix⟨T, symmetric, sparse⟩ Hsparse;
1867   matrix⟨T, unstructured, dense⟩ aux;
```

331. ⟨Entropy of a functional class methods 326⟩ +≡

```
1868   matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
1869   {
1870     T result ← f-eval(x);
1871     if (result ≤ 0) throw error::domain();
1872     f-grad(x, aux);
1873     outerp(aux, &Hdense);
1874     Hdense /= (-result);
1875     f-hess(x, dest);
1876     dest *= - (::log(result) + 1);
1877     return dest += Hdense;
1878   }
1879   matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
1880   {
1881     T result ← f-eval(x);
1882     if (result ≤ 0) throw error::domain();
1883     f-grad(x, aux);
1884     outerp(aux, &Hsparse);
1885     Hsparse /= (-result);
1886     f-hess(x, dest);
1887     dest *= - (::log(result) + 1);
1888     return dest += Hsparse;
1889   }
```

332. Relative entropy. The relative entropy (also known as Kullback Leibler distance) function is given by $x \log x/y$. The relative entropy is a measure of how much do we loose for assuming that the distribution of some random variable is y when the true distribution is x . If we knew x , then we could describe the random variable with a code with average description length $\mathbf{entr}(x)$. If we use y , however, we would need $\mathbf{entr}(x) + \mathbf{relentr}(x, y)$ to describe the variable.

/* Empty, waiting for **export** */

```

333.  ⟨functional/releintr.h 333⟩ ≡
1890  #ifndef __MATH_RELATIVE_ENTROPY_FUNCTIONAL__
1891  #define __MATH_RELATIVE_ENTROPY_FUNCTIONAL__ 1.00
1892  #include <math.h>
1893  #include <math/functionalbase.h>
1894      namespace math { namespace functional {
1895          template<class T>
1896          class releintr : public base<T> {
1897              ⟨Relative entropy functional internal variables 334⟩
1898          public:
1899              ⟨Relative entropy functional class methods 335⟩
1900          };
1901      } }
1902  #endif

```

```

334.  ⟨Relative entropy functional internal variables 334⟩ ≡
1903      typedef base<T> *element_type;
1904      element_type f, g;

```

See also sections 337 and 339.

This code is used in section 333.

```

335.  ⟨Relative entropy functional class methods 335⟩ ≡
1905      releintr(element_type new_f, element_type new_g)
1906      : f(new_f), g(new_g) { }

```

See also sections 336, 338, and 340.

This code is used in section 333.

```

336.  ⟨Relative entropy functional class methods 335⟩ +≡
1907      T eval(const matrix<T, unstructured, dense> &x)
1908      {
1909          T fval ← f→eval(x);
1910          T gval ← g→eval(x);
1911          if (gval ≤ 0 ∨ fval ≤ 0) throw error::domain();
1912          return fval * ::log(fval/gval);
1913      }

```

337. The gradient of the relative entropy is given by $(1 + \log(f/g))\nabla f - (f/g)\nabla g$.

```

⟨Relative entropy functional internal variables 334⟩ +≡
1914      matrix<T, unstructured, dense> auxvec;

```



```

338.  ⟨Relative entropy functional class methods 335⟩ +≡
1915  matrix⟨T, unstructured, dense⟩ &grad(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, unstructured, dense⟩ &dest)
1916  {
1917  T fval ← f-eval(x);
1918  T gval ← g-eval(x);
1919  if (gval ≤ 0 ∨ fval ≤ 0) throw error::domain();
1920  g-grad(x, auxvec);
1921  auxvec *= fval/gval;
1922  f-grad(x, dest);
1923  dest *= (::log(fval/gval) + 1);
1924  return dest -= auxvec;
1925  }

```

339. The Hessian is given by $(1 + \log(f/g))\nabla^2 f - (f/g)\nabla^2 g + f(\nabla g/g - \nabla f/f)(\nabla g/g - \nabla f/f)^T$.

```

⟨Relative entropy functional internal variables 334⟩ +≡
1926  matrix⟨T, unstructured, dense⟩ auxgrad;
1927  matrix⟨T, symmetric, dense⟩ auxdense;
1928  matrix⟨T, symmetric, sparse⟩ auxsparse;

```

```

340.  ⟨Relative entropy functional class methods 335⟩ +≡
1929  matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
1930  {
1931  #define auxhess auxdense
1932  ⟨relentr Hessian 341⟩;
1933  #undef auxhess
1934  }
1935  matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
1936  {
1937  #define auxhess auxsparse
1938  ⟨relentr Hessian 341⟩;
1939  #undef auxhess
1940  }

```

```

341.  ⟨relentr Hessian 341⟩ ≡
1941  T fval ← f-eval(x);
1942  T gval ← g-eval(x);
1943  if (gval ≤ 0 ∨ fval ≤ 0) throw error::domain();
1944  g-grad(x, auxvec);
1945  f-grad(x, auxgrad);
1946  auxvec /= gval;
1947  auxgrad /= fval;
1948  outerp(auxvec -= auxgrad, &dest);
1949  dest *= fval;
1950  g-hess(x, auxhess);
1951  auxhess *= fval / gval;
1952  dest -= auxhess;
1953  f-hess(x, auxhess);
1954  auxhess *= (::log(fval / gval) + 1);
1955  return dest += auxhess;

```

This code is used in section 340.

342. The Error Function. The error function (erf) is defined as

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-u^2} du$$

/* Empty, waiting for **export** */

```

343.  ⟨functional/erf.h 343⟩ ≡
1956  #ifndef __MATH_ERF_FUNCTIONAL__
1957  #define __MATH_ERF_FUNCTIONAL__ 1.00
1958  #include <math.h>
1959  #include <math/functionalbase.h>
1960  namespace math { namespace functional {
1961      template<class T>
1962      class erf:public base<T> {
1963          ⟨Error function of a functional internal variables 344⟩
1964      public:
1965          ⟨Error function of a functional class methods 345⟩
1966      };
1967  } }
1968  #endif

```

```

344.  ⟨Error function of a functional internal variables 344⟩ ≡
1969  typedef base<T> *element_type;
1970  element_type f;

```

See also section 349.

This code is used in section 343.

```

345.  ⟨Error function of a functional class methods 345⟩ ≡
1971  erf(element_type new_f):f(new_f) {}

```

See also sections 346, 348, and 350.

This code is used in section 343.

```

346.  ⟨Error function of a functional class methods 345⟩ +≡
1972  T eval(const matrix⟨T, unstructured, dense⟩ &x)
1973  {
1974    return (::erf(f→eval(x)));
1975  }

347.  The gradient of the error function is given by  $\frac{2}{\sqrt{\pi}}e^{-f^2}\nabla f$ .

348.  ⟨Error function of a functional class methods 345⟩ +≡
1976  matrix⟨T, unstructured, dense⟩ &grad(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, unstructured, dense⟩ &dest)
1977  {
1978    T result ← f→eval(x);
1979    f→grad(x, dest);
1980    return dest *= 2/::sqrt(M_PI) * ::exp(-result * result);
1981  }

349.  The Hessian is given by  $\frac{2}{\sqrt{\pi}}e^{-f^2}\nabla^2 f - \frac{4}{\sqrt{\pi}}fe^{-f^2}\nabla\nabla^T f$ .
⟨Error function of a functional internal variables 344⟩ +≡
1982  matrix⟨T, symmetric, dense⟩ dense_aux;
1983  matrix⟨T, symmetric, sparse⟩ sparse_aux;
1984  matrix⟨T, unstructured, dense⟩ aux;

350.  ⟨Error function of a functional class methods 345⟩ +≡
1985  matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
1986  {
1987  #define hess_aux dense_aux
1988    ⟨Compute Hessian for erf functional 351⟩;
1989  #undef hess_aux
1990  }
1991  matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
1992  {
1993  #define hess_aux sparse_aux
1994    ⟨Compute Hessian for erf functional 351⟩;
1995  #undef hess_aux
1996  }

351.  ⟨Compute Hessian for erf functional 351⟩ ≡
1997  T result ← f→eval(x);
1998  f→grad(x, aux);
1999  f→hess(x, dest);
2000  outerp(aux, &hess_aux);
2001  hess_aux *= -2 * result;
2002  dest += hess_aux;
2003  return dest *= 2 * ::exp(-result * result)/::sqrt(M_PI);

```

This code is used in section 350.

352. The exponential of a functional.

```
/* Empty, waiting for export */
```

```
353. <functional/exp.h 353> ≡
2004 #ifndef __MATH_EXP_FUNCTIONAL__
2005 #define __MATH_EXP_FUNCTIONAL__ 1.00
2006 #include <math.h>
2007 #include <math/functionalbase.h>
2008 namespace math {
2009     namespace functional
2010     {
2011         template<class T>
2012         class exp:public base<T> {
2013             <Exponential of a functional internal variables 354>
2014         public:
2015             <Exponential of a functional class methods 355>
2016         };
2017     }
2018 }
2019 #endif
```

```
354. <Exponential of a functional internal variables 354> ≡
2020     typedef base<T> *element_type;
2021     element_type f;
```

See also section 359.

This code is used in section 353.

```
355. <Exponential of a functional class methods 355> ≡
2022     exp(element_type new_f):f(new_f) { }
```

See also sections 356, 358, and 360.

This code is used in section 353.

```
356. <Exponential of a functional class methods 355> +≡
2023     T eval(const matrix<T, unstructured, dense> &x)
2024     {
2025         return ::exp(f->eval(x));
2026     }
```

357. The gradient of the exponential is given by $e^f \nabla f$.

```
358. <Exponential of a functional class methods 355> +≡
2027     matrix<T, unstructured, dense> &grad(const matrix<T, unstructured, dense>
2028         &x, matrix<T, unstructured, dense> &dest)
2029     {
2030         f->grad(x, dest);
2031         return dest *= eval(x);
2032     }
```

359. The Hessian is given by $e^f \nabla^2 f + e^f \nabla f \nabla^T f$.

⟨Exponential of a functional internal variables 354⟩ +≡

```
2032  matrix⟨T, symmetric, dense⟩ Hdense;
2033  matrix⟨T, symmetric, sparse⟩ Hsparse;
2034  matrix⟨T, unstructured, dense⟩ aux;
```

360. ⟨Exponential of a functional class methods 355⟩ +≡

```
2035  matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
2036  {
2037      f->grad(x, aux);
2038      outerp(aux, &Hdense);
2039      f->hess(x, dest);
2040      dest += Hdense;
2041      return dest *= eval(x);
2042  }
2043  matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
2044  {
2045      f->grad(x, aux);
2046      outerp(aux, &Hsparse);
2047      f->hess(x, dest);
2048      dest += Hsparse;
2049      return dest *= eval(x);
2050  }
```

361. The logarithm of a functional.

/* Empty, waiting for **export** */

362. ⟨functional/log.h 362⟩ ≡

```
2051  #ifndef __MATH_LOG_FUNCTIONAL__
2052  #define __MATH_LOG_FUNCTIONAL__ 1.00
2053  #include <math.h>
2054  #include <math/functionalbase.h>
2055  namespace math {
2056      namespace functional
2057      {
2058          template<class T>
2059          class log:public base⟨T⟩ {
2060              ⟨Logarithm of a functional internal variables 363⟩
2061              public:
2062                  ⟨Logarithm of a functional class methods 364⟩
2063              };
2064      }
2065  }
2066  #endif
```

363. The functionals to be added are stored in a vector of pairs of weights and functionals.

⟨Logarithm of a functional internal variables 363⟩ ≡

```
2067     typedef base⟨T⟩ *element_type;
2068     element_type f;
```

See also section 368.

This code is used in section 362.

364.

⟨Logarithm of a functional class methods 364⟩ ≡

```
2069     log(element_type new_f):f(new_f) { }
```

See also sections 365, 367, and 369.

This code is used in section 362.

365.

⟨Logarithm of a functional class methods 364⟩ +≡

```
2070     T eval(const matrix⟨T, unstructured, dense⟩ &x)
2071     {
2072         T result ← ::log(f→eval(x));
2073         if (¬finite(result)) throw error::domain();
2074         return result;
2075     }
```

366. The gradient of the logarithm is given by $\nabla f/f$.

367. ⟨Logarithm of a functional class methods 364⟩ +≡

```
2076     matrix⟨T, unstructured, dense⟩ &grad(const matrix⟨T, unstructured, dense⟩
2077         &x, matrix⟨T, unstructured, dense⟩ &dest)
2078     {
2079         T result ← f→eval(x);
2080         f→grad(x, dest);
2081         return dest /= result;
2082     }
```

368. The Hessian is given by $\nabla^2 f/f - \nabla f \nabla^T f / f^2$.

⟨Logarithm of a functional internal variables 363⟩ +≡

```
2082     matrix⟨T, symmetric, dense⟩ Hdense;
2083     matrix⟨T, symmetric, sparse⟩ Hsparse;
2084     matrix⟨T, unstructured, dense⟩ aux;
```

369. \langle Logarithm of a functional class methods 364 $\rangle + \equiv$

```

2085   matrix $\langle$ T, symmetric, dense $\rangle$  &hess(const matrix $\langle$ T, unstructured, dense $\rangle$ 
      &x, matrix $\langle$ T, symmetric, dense $\rangle$  &dest)
2086   {
2087     T result  $\leftarrow$  f-eval(x);
2088     f-grad(x, aux);
2089     outerp(aux, &Hdense);
2090     Hdense /= (-result);
2091     f-hess(x, dest);
2092     dest += Hdense;
2093     return dest /= result;
2094   }
2095   matrix $\langle$ T, symmetric, sparse $\rangle$  &hess(const matrix $\langle$ T, unstructured, dense $\rangle$ 
      &x, matrix $\langle$ T, symmetric, sparse $\rangle$  &dest)
2096   {
2097     T result  $\leftarrow$  f-eval(x);
2098     f-grad(x, aux);
2099     outerp(aux, &Hsparse);
2100     Hdense /= (-result);
2101     f-hess(x, dest);
2102     dest += Hsparse;
2103     return dest /= result;
2104   }

```

370. Line Searching. We begin now the definitions that will enable us to perform various types of numerical optimization. The first task we face is the one of minimizing a functional along a line. There are many algorithms to perform this task, and which to use will depend heavily on the specific problem (for example, is it cheap to compute the gradient, to evaluate the function at a point and so on).

```
/* Empty, waiting for export. */
```

```
371. <linesearchbase.h 371> ≡
2105 #ifndef __MATH_LINESEARCH__
2106 #define __MATH_LINESEARCH__ 1.0
2107 #include <math/functionalbase.h>
2108 namespace math {
2109     namespace linesearch
2110     {
2111         template<class T>
2112         class base {
2113         public:
2114             virtual ~base(void) { }
2115             <Line search base class methods 372>
2116         };
2117     }
2118 }
2119 #endif
```

372. The interface is very simple: there is one method that applies the minimization algorithm to a given functional, starting from a given point and searching in a given direction. The syntax is *minimize*(**functional**, *x0*, *dir*). The direction does not always need to have unity norm, but you better check out the specific algorithm to be sure that this is the case. The *minimize* method returns the minimizing point. As a design decision, it would be beneficial in some cases (for example in the backtracking algorithm) to include the gradient in the list of arguments, since the gradient would probably be available from the computation of the search direction. Since we cannot require this (for example, in some cases the gradient would have to be computed via some expensive simulation), we provide an alternative method that takes the gradient at the starting point as the third argument (the search direction becomes the fourth).

```
<Line search base class methods 372> ≡
2120     virtual matrix<T, unstructured, dense> minimize(functional::base<T> *, const
2121             matrix<T, unstructured, dense> &, const matrix<T, unstructured, dense> &) ← 0;
2121     virtual matrix<T, unstructured, dense> minimize(functional::base<T> *, const
2122             matrix<T, unstructured, dense> &, const matrix<T, unstructured, dense> &, const
2123             matrix<T, unstructured, dense> &) ← 0;
```

This code is used in section 371.

373. The bisection algorithm. Although not really a true bisection algorithm, it works by halving the step size when appropriate as we will see. This is only an example of line searching, but this algorithm can fail miserably in practice: depending on the function and on the disposition of local minima, you can end up in a local minima that is bigger than the closest one. You're advised do use other methods if possible. That said, we move on: As with many algorithms, we must provide a stopping criterion. In this case, the algorithm stops whenever $\|x - x^*\| \leq tol$, where *tol* is a given tolerance. Also, some people claim that halving the step size is not always the best thing to do. We then provide a means to alter this behavior via a **ratio** parameter, so that when the step size is changed, it changes via $step \leftarrow step * ratio$. Of course, we need $0 < ratio < 1$ for the algorithm to work.

```
/* Empty, waiting for export. */
```



```

374. <linsearch/bisection.h 374> ≡
2122 #ifndef __MATH_BISECTION__
2123 #define __MATH_BISECTION__ 1.0
2124 #include <math.h>
2125     const double bisection_infty ← HUGE_VAL;
        /* Some compilers have problems inlining this constant. */
2126 #include <math/algebra.h>
2127 #include <math/linsearchbase.h>
2128     namespace math {
2129         namespace linsearch
2130         {
2131             template<class T>
2132             class bisection:public base<T> {
2133                 double tol,
2134                     ratio;
2135                 unsigned long maxiter;    /* Stop in case its unbounded below or something. */
2136                 public: ~bisection(void)
2137                 { }
2138                 <Bisection line search methods 375>
2139             };
2140             <bisection big definitions 378>;
2141         }
2142     }
2143 #endif

```

375. First we provide a constructor that enable us to set all parameters at creation time, with some defaults in case we don't care about it.

```

<Bisection line search methods 375> ≡
2144     bisection(double t ← 1 · 10-3, double r ← 0.5, unsigned long i ← 1000)
2145     : tol(t), ratio(r), maxiter(i) { }

```

See also sections 376 and 377.

This code is used in section 374.

376. The first thing we do is to define the method that takes the gradient as an argument. We don't use the gradient for this algorithm, so the only thing we do is to call the other method.

```

<Bisection line search methods 375> +≡
2146     matrix<T, unstructured, dense> minimize(functional::base<T> *f,
        const matrix<T, unstructured, dense> &x0, const matrix<T, unstructured, dense>
        &g0, const matrix<T, unstructured, dense> &dir)
2147     {
2148         return minimize(f, x0, dir);
2149     }

```

377. The algorithm works as follows: we make $x \leftarrow x_0$. At each step, we start from x and go in the descending direction until we pass the minimum, as indicated by an increase in the function. At this point, we know the descent direction will be to the other side, so we turn around and proceed slower, that is, with a smaller stepsize. The stopping criterion is independent of x^* and x since

$$\|x - x^*\| \leq \|x - (x + \delta x)\| = \|\delta x\|,$$

where $\|\delta x\|$ is a positive multiple of $\|dir\|$. We consider that reaching the maximum number of iterations is an error. Note that the maximum number of iterations is with respect to the *inner* loop, since the outer loop always finishes. The algorithm follows:

⟨ Bisection line search methods 375 ⟩ +≡

```
2150   matrix⟨T, unstructured, dense⟩ minimize(functional::base⟨T⟩ *func, const matrix⟨T,
      unstructured, dense⟩ &x0, const matrix⟨T, unstructured, dense⟩ &dir);
```

378. ⟨ bisection big definitions 378 ⟩ ≡

```
2151   template⟨class T⟩ matrix⟨T, unstructured, dense⟩ bisection⟨T⟩::minimize(functional::base⟨T⟩
      *func, const matrix⟨T, unstructured, dense⟩ &x0, const matrix⟨T, unstructured, dense⟩
      &dir){ double stepsize ← 1, dirnorm ← norm2(dir);
2152   matrix⟨T, unstructured, dense⟩ x(x0);
2153   unsigned long outer_iter ← 0; while ( fabs ( stepsize / ratio )
      *dirnorm > tol ∧ ++outer_iter ≤ maxiter ) { double f ← func-eval(x);
2154   double newf;
2155   ⟨ Compute next function value for bisection 379 ⟩;
2156   unsigned long inner_iter ← 0;
2157   while ( newf < f ) {
2158     f ← newf;
2159     ⟨ Compute next function value for bisection 379 ⟩;
2160     if ( ++inner_iter > maxiter ) throw error::unboundedbelow();
2161   }
2162   stepsize *= -ratio; } return x; }
```

This code is used in section 374.

```

379.  ⟨ Compute next function value for bisection 379 ⟩ ≡
2163    saxpy(stepsize, dir, &x);
2164    try {
2165        newf ← func-eval(x);
2166    }
2167    catch(error :: domain e)
2168    {
2169        newf ← bisection_infty;
2170    }
2171    while (¬finite(newf)) { saxpy(-stepsize, dir, &x); stepsize *= ratio;
2172    saxpy(stepsize, dir, &x);
2173    try {
2174        newf ← func-eval(x);
2175    }
2176    catch(error :: domain e)
2177    {
2178        newf ← bisection_infty;
2179    }
2180    }
2181    if (stepsize ≡ 0) return x;

```

This code is used in section 378.

380. Backtracking. This line search algorithm uses the gradient information. We begin with a unit step λ , and until we have

$$f(x_0 + \lambda \cdot dir) \leq f(x_0) + \lambda \alpha \nabla f(x_0)^T dir,$$

we update the step with $\beta\lambda$. The algorithm parameters are α and β , where $0 < \alpha < 0.5$ and $0 < \beta < 1$.

```
/* Empty, waiting for export. */
```

```

381.  ⟨ linesearch/backtracking.h 381 ⟩ ≡
2182  #ifndef __MATH_BACKTRACKING_LINESEARCH__
2183  #define __MATH_BACKTRACKING_LINESEARCH__ 1.0
2184    const double backtracking_infty ← HUGE_VAL;
2185    /* Some compilers have problems using HUGE_VAL. */
2186  #include <math/algebra.h>
2187  #include <math/linesearchbase.h>
2188    namespace math {
2189        namespace linesearch
2190        {
2191            template<class T>
2192            class backtracking:public base<T> {
2193                double alpha, beta;
2194                unsigned long maxiter;
2195            public: ~backtracking(void)
2196                {}
2197                ⟨ Backtracking line search methods 382 ⟩
2198            };
2199        }
2200    #endif

```

382. First we provide a constructor that enable us to set all parameters at creation time, with some defaults in case we don't care about it.

⟨Backtracking line search methods 382⟩ ≡

```
2201   backtracking(double a ← 0.3, double b ← 0.8, unsigned long m ← 1000)
2202   : alpha(a), beta(b), maxiter(m) { }
```

See also sections 383 and 384.

This code is used in section 381.

383. The first thing we do is to define the method that doesn't take the gradient as an argument. Since we need the gradient, what this method does is to compute it and call the correct method.

⟨Backtracking line search methods 382⟩ +≡

```
2203   matrix⟨T, unstructured, dense⟩ minimize(functional::base⟨T⟩ *f, const matrix⟨T,
      unstructured, dense⟩ &x0, const matrix⟨T, unstructured, dense⟩ &dir)
2204   {
2205     matrix⟨T, unstructured, dense⟩ g0;
2206     g0 ← f-grad(x0, g0);
2207     return minimize(f, x0, g0, dir);
2208   }
```

384.

⟨Backtracking line search methods 382⟩ +≡

```
2209   matrix⟨T, unstructured, dense⟩ minimize(functional::base⟨T⟩ *func,
      const matrix⟨T, unstructured, dense⟩ &x0, const matrix⟨T, unstructured, dense⟩
      &g0, const matrix⟨T, unstructured, dense⟩ &dir)
2210   {
2211     matrix⟨T, unstructured, dense⟩ x;
2212     double ftreshold ← func-eval(x0);
2213     double gtreshold ← alpha * dot(g0, dir);
2214     double stepsize ← 1/beta;
2215     unsigned long iter ← 0;
2216     double fval;
2217     #ifdef __MATH_DEBUG__
2218       cout << "[math]:_backtracking_line_search_begin.\n";
2219       cout << "[math]:_x0=_";
2220       for (index i ← 1; i ≤ x0.rows(); ++i) cout << x0(i) << '_';
2221       cout << '\n';
2222       cout << "[math]:_dir=_";
2223       for (index i ← 1; i ≤ dir.rows(); ++i) cout << dir(i) << '_';
2224       cout << '\n';
2225     #endif
2226     do {
2227       ⟨Compute next function value for backtracking 385⟩;
2228       if (++iter ≡ maxiter) throw error::maxiterations();
2229     } while (fval > ftreshold + stepsize * gtreshold);
2230     #ifdef __MATH_DEBUG__
2231       cout << "[math]:_backtracking_line_search_end.\n";
2232     #endif
2233     return x;
2234   }
```

```

385.  ⟨ Compute next function value for backtracking 385 ⟩ ≡
2235      fval ← backtracking_infty;
2236      while (¬finite(fval)) {
2237          x ← dir;
2238          x *= (stepsize *= beta);
2239          x += x0;
2240      #ifdef __MATH_DEBUG__
2241          cout << "[math]:_x=_";
2242          for (index i ← 1; i ≤ x.rows(); ++i) cout << x(i) << ' ';
2243          cout << '\n';
2244      #endif
2245          try {
2246              fval ← func-eval(x);
2247          }
2248          catch(error::domain e)
2249          {
2250              fval ← backtracking_infty;
2251          }
2252      #ifdef __MATH_DEBUG__
2253          cout << "[math]:_f(x)=" << fval << ",_stepsize=" << stepsize << ",_treshold=" <<
                ftreshold + stepsize * gtreshold << '\n';
2254      #endif
2255          if (stepsize ≡ 0.0) fval ← ftreshold + stepsize * gtreshold;
2256      }

```

This code is used in section 384.

386. Computing a search direction. Once you have a line minimization algorithm the only thing that you still need in order to define a local minimization algorithm is a method for computing the search direction.

```
/* Empty, waiting for export. */
```

```

387. <searchdirbase.h 387> ≡
2257 #ifndef __MATH_SEARCHDIRBASE__
2258 #define __MATH_SEARCHDIRBASE__ 1.0
2259 #include <math/math.h>
2260 #include <math/functionalbase.h>
2261 namespace math {
2262     namespace searchdir
2263     {
2264         template<class T>
2265         class base {
2266         public:
2267             virtual ~base(void) { }
2268             <Search direction base class methods 388>
2269         };
2270     }
2271 }
2272 #endif

```

388. Now to the interface: The only thing a search direction method has to do is to compute a descent direction based on a functional and a starting point. Our method will return the descent direction.

<Search direction base class methods 388> ≡

```

2273     virtual matrix<T, unstructured, dense> &dir(functional::base<T> *f, const
        matrix<T, unstructured, dense> &xi, matrix<T, unstructured, dense> *dest) ← 0;

```

This code is used in section 387.

389. The gradient direction. If you're able to compute the gradient easily (that is, not requiring simulation or something) *and* the Hessian is too costly for you, then the gradient search direction can be a good choice.

```
/* Empty, waiting for export. */
```

```

390. <searchdir/gradient.h 390> ≡
2274 #ifndef __MATH_SEARCHDIR_GRADIENT__
2275 #define __MATH_SEARCHDIR_GRADIENT__ 1.0
2276 #include <math/searchdirbase.h>
2277 namespace math {
2278     namespace searchdir
2279     {
2280         template<class T>
2281         class gradient:public base<T> {
2282         public:
2283             ~gradient(void)
2284             {}
2285             < Gradient search direction methods 391 >
2286         };
2287     }
2288 }
2289 #endif

```

391. The search direction is basically $-\nabla f$.

```

< Gradient search direction methods 391 > ≡
2290     matrix<T, unstructured, dense> &dir(functional::base<T> *f, const matrix<T, unstructured,
2291         dense> &x, matrix<T, unstructured, dense> *dest)
2292     {
2293         f->grad(x, *dest);
2294         return (*dest) *= -1;
2295     }

```

This code is used in section 390.

392. The Newton direction. If you are lucky enough that the Hessian computation is not a big deal and you can afford solving a linear system each time the search direction is needed, then the Newton direction is a no-brainer. The only thing that needs to be defined is what type of storage you want to use for the Hessian. For large, sparse problems, you probably want a sparse storage.

```

/* Empty, waiting for export. */

```

```

393. <searchdir/newton.h 393> ≡
2295 #ifndef __MATH_SEARCHDIR_NEWTON__
2296 #define __MATH_SEARCHDIR_NEWTON__ 1.0
2297 #include <math/searchdirbase.h>
2298 #include <math/algebra.h>
2299 #include <math/lu.h>
2300 #include <math/cholesky.h>
2301 namespace math {
2302     namespace searchdir
2303     {
2304         template<class T, template<class> class storage>
2305         class newton:public base<T> {
2306             <Newton search direction internal variables 394>
2307         public:
2308             <Newton search direction methods 395>
2309         };
2310     }
2311 }
2312 #endif

```

394. We assume that you'll need to compute the search direction many times, so we reserve space for the Hessian.

```

<Newton search direction internal variables 394> ≡
2313     matrix<T, symmetric, storage> hess;

```

This code is used in section 393.

395. The search direction is now $-(\nabla^2 f)^{-1}\nabla f$. We first compute the gradient and store it in *aux*, following by a Hessian computation and Cholesky linear system solver. In case the Hessian is singular or close, we return the gradient direction (which is still a descent direction).

```

<Newton search direction methods 395> ≡
2314     matrix<T, unstructured, dense> &dir(functional::base<T> *f, const matrix<T, unstructured,
           dense> &x, matrix<T, unstructured, dense> *dest)
2315     {
2316         try {
2317             try {
2318                 f->grad(x, *dest);
2319                 f->hess(x, hess);
2320                 cholesky::solve(&hess, dest);
2321             }
2322             catch(error::nonpositivedef e)
2323             {
2324                 f->grad(x, *dest);
2325                 f->hess(x, hess);
2326                 lu::solve(&hess, dest);
2327             }
2328         }
2329         catch(error::singular e) { f->grad(x, *dest); }
2330         return (*dest) *= -1;
2331     }

```

This code is used in section 393.

396. Enforcing equality constraints. All the search direction classes defined above are not able to take into account the search direction must, in some cases, have zero component in some directions. The most obvious and common case is when we want to minimize something while enforcing an equality constraint $Ax = b$. We will define a class that, based on a search direction d computed in any way, transforms it so that $A(x + tv) = b$ for any real t .

```
/* Empty, waiting for export. */
```

```

397. <searchdir/equality.h 397> ≡
2332 #ifndef __MATH_SEARCHDIR_EQUALITY__
2333 #define __MATH_SEARCHDIR_EQUALITY__ 1.0
2334 #include <math/searchdirbase.h>
2335 #include <math/symmetric.h>
2336 #include <math/lu.h>
2337 namespace math {
2338     namespace searchdir
2339     {
2340         template<matrix_simple_template>
2341         class equality:public base<T> {
2342             <Equality search direction internal variables 398>
2343         public:
2344             <Equality search direction methods 399>
2345         };
2346     }
2347 }
2348 #endif

```

398. First we need a way to let the user define which basic search direction he wants to use.

```

<Equality search direction internal variables 398> ≡
2349     base<T> *dirf;

```

See also section 400.

This code is used in section 397.

```

399. <Equality search direction methods 399> ≡
2350     equality(void):dirf(0) { }
2351     equality(base<T> *newf):dirf(newf) { }
2352     base<T> *&f(void) { return dirf; }

```

See also sections 401 and 403.

This code is used in section 397.

400. Now to the A matrix. Suppose that our basic direction class provides us with a preferred direction d_0 . What we do is to compute the projection of d_0 into the kernel of A . This task can be accomplished by solving the linear system

$$\begin{bmatrix} I & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} d \\ w \end{bmatrix} = \begin{bmatrix} d_0 \\ 0 \end{bmatrix},$$

which in the end will give

$$d = (I - A^T(AA^T)^{-1}A)d_0$$

. That's the only time where we use the A matrix, so we can see that we don't need to store it: what we store is the LU decomposition (the matrix is not necessarily positive-definite) of the linear system matrix, so that no matter how many directions we compute, solving the linear system is a very fast operation (no decompositions needed!). The only catch here is that we may receive (and we will when barrier functions are defined) a preferred direction with the wrong dimension. We store the correct dimension in $xrows$ and deal with it later on.

(Equality search direction internal variables 398) +=

```
2353   matrix<T, unstructured, storage> decomp;
2354   vector<index> pivots;
2355   index xrows;
```

401. (Equality search direction methods 399) +=

```
2356   equality(const matrix<T, structure, storage> &A):dirf(0) { set_A(newa); }
2357   void set_A(const matrix<T, structure, storage> &A)
2358   {
2359   #ifdef __MATH_DEBUG__
2360       cout << "[math]:_equality_search_constructor.\n";
2361       for (index i ← 1; i ≤ A.rows(); ++i) {
2362           for (index j ← 1; j ≤ A.cols(); ++j) cout << A(i,j) << ' ';
2363           cout << '\n';
2364       }
2365   #endif
2366       xrows ← A.cols();
2367       decomp.resize(A.cols() + A.rows(), A.cols() + A.rows());
2368       decomp.fillwith(0.0);
2369       decomp.subm(1, A.cols(), 1 + A.cols(), decomp.cols()) ← transpose(A);
2370       decomp.subm(1 + A.cols(), decomp.rows(), 1, A.cols()) ← A;
2371       for (index i ← 1; i ≤ A.cols(); ++i) decomp.entry(i, i) ← 1;
2372       lu::decompose(&decomp, &pivots);
2373   }
```

402. Now we're ready to compute the direction. As said before, we may receive a preferred direction with the wrong dimensions. This happens when barrier functions are in use and a new slack variable is introduced in the problem after the A matrix is defined. Therefore, when x has the wrong dimension, we know it has exactly one more component than the original number, and what we do is to project only the first components.

403. \langle Equality search direction methods 399 $\rangle + \equiv$

```

2374     matrix<T, unstructured, dense> &dir(functional::base<T> *f, const matrix<T, unstructured,
           dense> &x, matrix<T, unstructured, dense> *dest)
2375     {
2376         double extra_component ← 0;
2377         dest→resize(decomp.rows(), 1);
2378         dest→fillwith(0.0);
2379         dirf→dir(f, x, dest);
2380     #ifdef __MATH_DEBUG__
2381         cout << "[math]:_equality_search_dir_begin.\n[math]:_orig_dir=";
2382         for (index i ← 1; i ≤ x.rows(); ++i) cout << (*dest)(i) << ' ';
2383         cout << '\n';
2384     #endif
2385         if (x.rows() ≠ xrows) extra_component ← (*dest)(xrows + 1);
2386         dest→entry(xrows + 1) ← 0.0;
2387         lu::finish(decomp, pivots, dest);
2388         dest→resize(x.rows(), 1);
2389         if (x.rows() ≠ xrows) dest→entry(xrows + 1) ← extra_component;
2390     #ifdef __MATH_DEBUG__
2391         cout << "[math]:_new_dir=";
2392         for (index i ← 1; i ≤ result.rows(); ++i) cout << (*dest)(i) << ' ';
2393         cout << "\n[math]:_equality_search_dir_end.\n";
2394     #endif
2395         return *dest;
2396     }

```

404. Newton direction with equality constraints.

```
/* Empty, waiting for export */
```

405. \langle searchdir/equality/newton.h 405 $\rangle \equiv$

```

2397     #ifndef __MATH_SEARCHDIR_EQUALITY_NEWTON__
2398     #define __MATH_SEARCHDIR_EQUALITY_NEWTON__
2399     #include <math/math.h>
2400     #include <math/searchdirbase.h>
2401     #include <math/lu.h>
2402     namespace math { namespace searchdir { namespace equality { template <class T, template
           < class > class storage > class newton : public base<T> {
2403         < Newton with equality internal variables 406 >;
2404         public: < Newton with equality functions 408 >;
2405         }; } } }
2406     #endif

```

406. Let us be wise memorywise. Unlike with the pure equality search direction, we will have to restore the system matrix every time we compute a new direction. So that's what we do: we store the A matrix literally. This matrix will not change once its initialized. Next, we keep the system matrix in a matrix M .

\langle Newton with equality internal variables 406 $\rangle \equiv$

```

2407     matrix<T, unstructured, storage> A;
2408     matrix<T, symmetric, storage> H;
2409     matrix<T, unstructured, storage> M;

```

See also section 407.

This code is used in section 405.

407. Of course, we also face the same auxiliary variable problem.

⟨Newton with equality internal variables 406⟩ +≡

```
2410     bool has_t;  
2411     index xrows;    /* Original number of variables. */
```

```

408.  ⟨Newton with equality functions 408⟩ ≡
2412  matrix⟨T, unstructured, dense⟩ &dir (functional::base⟨T⟩ *f, const matrix⟨T, unstructured,
      dense⟩ &x, matrix⟨T, unstructured, dense⟩ *dest)
2413  {
2414    if (x.rows() ≠ xrows ∧ has_t ≡ false) {
2415      has_t ← true;
2416      A.resize(A.rows(), xrows + 1);
2417      for (index i ← 1; i ≤ A.rows(); ++i) A.entry(i, xrows + 1) ← 0;
2418    }
2419    if (x.rows() ≡ xrows ∧ has_t ≡ true) {
2420      has_t ← false;
2421      A.resize(A.rows(), xrows);
2422    }
2423    try {
2424      f-grad(x, *dest);
2425      f-hess(x, H);
2426    #ifdef __MATH_DEBUG__
2427      cout << "[math]:_got_Hessian.\n";
2428      cout << "H=\n";
2429      for (index i ← 1; i ≤ H.rows(); ++i) {
2430        cout << "row_" << i << ":\n";
2431        for (index j ← 1; j ≤ H.cols(); ++j) cout << H(i, j) << ' ';
2432        cout << '\n';
2433      }
2434    #endif
2435      dest-resize(x.rows() + A.rows(), 1);
2436      M.resize(x.rows() + A.rows(), x.rows() + A.rows());
2437      for (index i ← 1; i ≤ A.rows(); ++i) {
2438        dest-entry(i + x.rows(), 1) ← 0;
2439        for (index j ← 1; j ≤ A.rows(); ++j) M.entry(i + x.rows(), j + x.rows()) ← 0;
2440        for (index j ← 1; j ≤ A.cols(); ++j) {
2441          M.entry(i + x.rows(), j) ← A(i, j);
2442          M.entry(j, i + x.rows()) ← A(i, j);
2443        }
2444      }
2445      for (index i ← 1; i ≤ x.rows(); ++i)
2446        for (index j ← i; j ≤ x.rows(); ++j) {
2447          M.entry(i, j) ← H(i, j);
2448          M.entry(j, i) ← H(j, i);
2449        }
2450    #ifdef __MATH_DEBUG__
2451      cout << "[math]:_about_to_get_search_direction.\n";
2452      cout << "M=\n";
2453      for (index i ← 1; i ≤ M.rows(); ++i) {
2454        for (index j ← 1; j ≤ M.cols(); ++j) cout << M(i, j) << ' ';
2455        cout << '\n';
2456      }
2457      cout << "dest=\n";
2458      for (index i ← 1; i ≤ dest-rows(); ++i) cout << dest-get(i, 1) << ' ';
2459      cout << '\n';
2460    #endif
2461      lu::solve(M, *dest);

```

```

2462     dest->resize(x.rows(),1);
2463 #ifdef __MATH_DEBUG__
2464     cout << "dest=\n";
2465     for (index i ← 1; i ≤ dest->rows(); ++i) cout << dest->get(i,1) << '␣';
2466     cout << '\n';
2467 #endif
2468     }
2469     catch(error::singular e)
2470     {
2471         f->grad(x,*dest);
2472     }
2473     return (*dest) *= -1;
2474     }

```

See also section 409.

This code is used in section 405.

409. Now we provide a function to set the A matrix.

⟨Newton with equality functions 408⟩ +≡

```

2475     void set_A(const matrix<T,unstructured,storage> &a)
2476     {
2477         has_t ← false;
2478         xrows ← a.cols();
2479         A ← a;
2480     }

```

410. Optimization algorithms. We have a generic matrix class, we are able to perform some useful decompositions, we have a functional class, a line search class and a descent direction class. We are in position to build up a generic optimization package. The ultimate goal would be to efficiently solve the problem

$$\begin{aligned} & \text{minimize} && f_0(x) \\ & \text{subject to} && f_i(x) < 0, && i = 1, \dots, p \\ & && g_i(x) = 0, && i = 1, \dots, q \end{aligned}$$

for any type of convex constraints. The best way to accomplish this goal is by using barrier-function based optimization algorithms. These are functions that are infinite outside the feasible set of the constraint and finite inside, that is, $\phi_i(x) = \infty$ if and only if $f_i(x) \geq 0$. With this property, we see that the minimum value of $f_0(x) + \Sigma\phi_i(x)$ is bounded above if and only if the optimization problem is feasible. In what follows we build the base that will enable us to come up with a generic sequential unconstrained optimization routine. The basic idea is as follows, we compute the minimum of $tf_0(x) + \Sigma\phi_i(x)$ for increasing values of t . As this value goes to infinity we will approach the optimal solution of the original problem. The equality constraints will be satisfied if we use a search direction class that ensures this property – so we are able to ensure $Ax = b$, for example.

411. Functional minimization. The first thing we need is to be able to minimize a functional without any constraints. Strictly speaking the correct thing to do would be to define a base class for functional minimization and specialize it depending on the stop criteria, but in the real world we face only two of them:

$$\begin{aligned} |x_i - x_i^*| &\leq |x_i^*|reltol + abstol \\ \|\nabla f(x)\| &\leq abstol. \end{aligned}$$

Of course, the first condition only makes sense when $abstol \cdot reltol = 0$. From that observation we decided to define only one functional minimization function that can take into account all of the above stopping criteria.

The function we'll define will take various arguments: the first are the functional to be minimized, the starting point, the line search algorithm and the search direction algorithm. The starting point will be overwritten with the optimal point approximation and returned on exit. The next arguments define the stopping criteria: *abstol* and *reltol* have the meaning of that first stopping condition described above. They can be both positive, in which case the optimization stops when *both* the absolute error and the relative error conditions are met. If you don't want one of them to have an effect, simply assign it zero. The *gtol* argument specifies, if nonzero, the stop condition for the gradient norm. At last, the user can input the maximum allowed number of iterations and a pointer to a function that gets the current point at each iteration (possibly for displaying interactively).

/* Empty, waiting for **export**. */

```

412.  <fmin.h 412> ≡
2481  #ifndef __MATH_FMIN__
2482  #define __MATH_FMIN__ 1.0
2483  #include <math/functionalbase.h>
2484  #include <math/lineearchbase.h>
2485  #include <math/searchdirbase.h>
2486  namespace math {
2487  template<class T>
2488  matrix<T,
          unstructured, dense> &fmin(functional::base<T> *f, matrix<T, unstructured, dense>
          &x, lineearch::base<T> *lsearch, searchdir::base<T> *sdir, T abstol ← 1 · 10-4, T
          reltol ← 1 · 10-3, T gtol ← 1 · 10-3, unsigned long maxiter ← 1000, void(*disp)(const
          matrix<T, unstructured, dense> &) ← 0)
2489  {
2490  <Functional minimization algorithm 413>
2491  }
2492  }
2493  #endif

```

413. The algorithm is pretty simple: we keep calling the line minimization function until the stop criteria is met or the maximum number of iterations is reached (in which case we return the last line minimizer). The only trick in this function is that we only use the gradient for line minimization if $gtol > 0$. In this case, our function keeps the last function gradient in $grad$ (which is otherwise not used).

<Functional minimization algorithm 413> ≡

```

2494  {
2495  matrix<T, unstructured, dense> x0, grad, dir;
2496  if (gtol) f→grad(x, grad);
2497  bool stop ← false;
2498  for (unsigned long iter ← 0; ¬stop ∧ iter ≠ maxiter; ++iter) {
2499  x0 ← x;
2500  if (gtol) x ← lsearch→minimize(f, x, grad, sdir→dir(f, x, &dir));
2501  else x ← lsearch→minimize(f, x, sdir→dir(f, x, &dir));
2502  <Update functional minimization stop criteria 414>
2503  if (disp) disp(x);
2504  }
2505  return x;
2506  }

```

This code is used in section 412.

414. We test if all stopping criteria are satisfied. For the relative tolerance we have to check if it makes sense: if $x_i = 0$ no relative tolerance can be reached (in practice).

<Update functional minimization stop criteria 414> ≡

```

2507  stop ← true;
2508  if (abstol) for (index i ← 1; i ≤ x.rows() ∧ stop; ++i) stop ← (fabs(x(i) - x0(i)) ≤ abstol);
2509  if (reltol) for (index i ← 1; i ≤ x.rows() ∧ stop; ++i)
2510  if (x(i)) stop ← (fabs(x(i) - x0(i)) ≤ fabs(x(i)) * reltol);
2511  if (gtol) stop &= (norm2(f→grad(x, grad)) ≤ gtol);

```

This code is used in section 413.

415. Barrier functions. The next step necessary in order to achieve our goal is to define barrier functions. Barrier functions, as aligned before, are functionals that are bounded above if and only if the constraint to which they relate is feasible. In our convention, a constraint is feasible when its defining function is negative.

```
/* Empty, waiting for export */
```

```

416. <barrierbase.h 416> ≡
2512 #ifndef __MATH_BARRIER__
2513 #define __MATH_BARRIER__ 1.0
2514 #include <math/functionalbase.h>
2515 #include <math/functionbase.h>
2516 namespace math {
2517     namespace barrier
2518     {
2519         template<class T> class base:public functional::base<T>
2520         {
2521             <Barrier function internal variables 417>
2522             public:
2523                 <Barrier function methods 418>
2524             }
2525         };
2526     }
2527 }
2528 #endif

```

417. In order to make things easier for the programmer, we allow a barrier function to be related to a *vector-valued function*. We will store pointers to the function or functional, and the convention is that only one of them can be nonzero at any time.

```

<Barrier function internal variables 417> ≡
2529 protected: functional::base<T> *fctnal;
2530 function::base<T> *fct;

```

See also section 420.

This code is used in section 416.

```

418.  ⟨Barrier function methods 418⟩ ≡
2531  base(functional::base⟨T⟩ *newf ← 0)
2532  : fctnal(newf), fct(0), has_t(false) { }
2533  base(function::base⟨T⟩ *newf)
2534  : fctnal(0), fct(newf), has_t(false) { }
2535  void f(functional::base⟨T⟩ *newf)
2536  {
2537      fctnal ← newf;
2538      fct ← 0;
2539  }
2540  void f(function::base⟨T⟩ *newf)
2541  {
2542      fctnal ← 0;
2543      fct ← newf;
2544  }
2545  virtual ~base(void) { }

```

See also sections 421, 422, 423, and 424.

This code is used in section 416.

419. We don't need to define the gradient and Hessian methods because we're already derived from a class that derive them. The only thing we have to define is the behavior when dealing with vector-valued functions: in this case, the gradient is the sum of the gradient of the vector elements, and similarly for the Hessian and the evaluation functions. There's another important thing that must be taken into account by evaluation methods.

420. We use barrier function in constrained optimization problems. Some algorithms don't require an initial feasible point, that is, we don't need an x such that $f(x) < 0$. One standard way to find a feasible point is to add a variable to the constraint so that $f(x) < t$ is feasible, and then minimize t . These methods will then call the *addt* function of the barrier functions to signal that we should behave as if this variable existed. From then on, until a call for *delt*, the last component of the argument for *eval* and so on is considered to be t , and the evaluation methods must deal with it accordingly.

⟨Barrier function internal variables 417⟩ +=

```

2546  bool has_t;

```

421. The *addt* function returns a value of t that satisfies $f(x) \leq t$.

⟨Barrier function methods 418⟩ +=

```

2547  T addt(const matrix⟨T, unstructured, dense⟩ &x)
2548  {
2549      T result ← T(0);
2550      if (fctnal) result ← fctnal-eval(x);
2551      else {
2552          matrix⟨T, unstructured, dense⟩ aux;
2553          fct-eval(x, aux);
2554          for (index i ← 1; i ≤ aux.rows(); ++i) result += aux(i);
2555      }
2556      has_t ← true;
2557      return result;
2558  }

```

422. `<Barrier function methods 418> +≡`

2559 `void delt(void) { has_t ← false; }`

423. Barrier functions are used in optimization problems in order to deal with constraints. To each constraint we have an associated dual variable. This is what happens: we will want to solve the problem

$$\begin{aligned} &\text{minimize} && f_0(x) \\ &\text{subject to} && f_i(x) < 0 \end{aligned}$$

and we will sequentially solve

$$\text{minimize} \quad \mu f_0(x) + \sum \phi(f_i(x))$$

for various values of μ . For each of these values, the solution of the second problem will also satisfy $\nabla f_0(x) + \sum \lambda_i \nabla f_i(x) = 0$ for some λ_i that depends on the particular barrier function used. But if x^* satisfies that equality, it is a minimizer for the Lagrangian with the same λ_i , that is, λ_i are feasible dual variables and can provide a lower bound on the optimal value. It is clear that the λ_i will be dependent on μ , but the barrier function doesn't know its value. What we do then is to return $\mu\lambda_i$, which is independent of μ .

`<Barrier function methods 418> +≡`

2560 `virtual T dual(const matrix<T, unstructured, dense> &) ← 0;`

424. On other occasions it will be useful to get the value of the dual variable times the function itself, and it is sometimes much cheaper to compute this value (see, for example, the log barrier function). The same above remarks with respect to μ are valid here, that is, the function should return $\mu\lambda_i f(x)$.

`<Barrier function methods 418> +≡`

2561 `virtual T dual_times_f(const matrix<T, unstructured, dense> &) ← 0;`

425. **The log barrier function.** Perhaps the most common, its value is $-\log(-f(x))$.

`/* Empty, waiting for export */`

426. `<barrier/log.h 426> ≡`

2562 `#ifndef __MATH_LOG_BARRIER__`

2563 `#define __MATH_LOG_BARRIER__`

2564 `#include <math/barrierbase.h>`

2565 `namespace math {`

2566 `namespace barrier`

2567 `{`

2568 `template<class T> classlog:public base<T>`

2569 `{`

2570 `matrix<T, unstructured, dense> aux, jac;`

2571 `matrix<T, symmetric, dense> dense_aux;`

2572 `matrix<T, symmetric, sparse> sparse_aux;`

2573 `public:`

2574 `<Log barrier methods 427>`

2575 `}`

2576 `;`

2577 `}`

2578 `}`

2579 `#endif`

427. \langle Log barrier methods 427 $\rangle + \equiv$

2580 **log**(**functional**::**base** \langle **T** \rangle **newf*) { *f*(*newf*); }

2581 **log**(**function**::**base** \langle **T** \rangle **newf*) { *f*(*newf*); }

See also sections 428, 429, 430, and 431.

This code is used in section 426.

428. Remember that we have to take into account that an auxiliary variable may be present. In that case, we first get the value of t , resize x and get the value and the gradient of the function. Resizing is not a big burden on performance because its optimized so that normally no memory allocation is necessary.

\langle Log barrier methods 427 $\rangle + \equiv$

```

2582 T eval(const matrix $\langle$ T, unstructured, dense $\rangle$  &x)
2583 {
2584   T result  $\leftarrow$  0, t  $\leftarrow$  0;
2585   matrix $\langle$ T, unstructured, dense $\rangle$  X  $\leftarrow$  x;    /* We may resize x. */
2586   if (has_t) {
2587     t  $\leftarrow$  X.rows();
2588     X.resize(X.rows() - 1, 1);
2589   }
2590   if (fctnal) {
2591     result  $\leftarrow$  fctnal-eval(X) - t;
2592     if (result  $\geq$  0) throw error::domain();
2593     return -::log(-result);
2594   }
2595   fct-eval(X, aux);
2596   for (index i  $\leftarrow$  1; i  $\leq$  aux.rows(); ++i) {
2597     if (aux(i) - t  $\geq$  0) throw error::domain();
2598     result -= ::log(t - aux(i));
2599   }
2600   return result;
2601 }
```

429. For the logarithm, the dual variable is $-1/(\mu f(x))$, so we return $-1/f(x)$.

\langle Log barrier methods 427 $\rangle + \equiv$

2602 **T** *dual*(**const matrix** \langle **T**, **unstructured**, **dense** \rangle &*x*) { **return** $-1/eval(x)$; }

2603 **T** *dual_times_f*(**const matrix** \langle **T**, **unstructured**, **dense** \rangle &*x*) { **return** **T**(-1.0); }

430. Now to the gradient: for a functional f we have $\nabla -\log(-f) = -(\nabla f)/f$. By our definition for a function we have the sum $-\sum_i(\nabla f_i)/f_i$.

⟨Log barrier methods 427⟩ +≡

```

2604   matrix⟨T, unstructured, dense⟩ &grad(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, unstructured, dense⟩ &dest)
2605   {
2606     T t ← 0;
2607     matrix⟨T, unstructured, dense⟩ X ← x;    /* We may resize x. */
2608     if (has_t) {
2609       t ← X(X.rows());
2610       X.resize(X.rows() - 1, 1);
2611     }
2612     if (fctnal) {
2613       fctnal→grad(X, dest);
2614       if (has_t) {
2615         dest.resize(X.rows() + 1, 1);
2616         dest.entry(dest.rows()) ← T(-1);
2617       }
2618       dest /= t - fctnal→eval(X);
2619       return dest;
2620     }
2621     fct→eval(X, aux);
2622     fct→jacobian(X, dest);
2623     if (has_t) {
2624       dest.resize(dest.rows(), dest.cols() + 1);
2625       dest.subm(1, dest.rows(), dest.cols()) ← matrix⟨T, unstructured, dense⟩(dest.rows(), 1, T(-1));
2626     }
2627     for (index i ← 1; i ≠ aux.rows(); ++i) {
2628       dest.subm(i, i, 1, dest.cols()) /= t - aux(i);
2629       if (i ≠ 1) dest.subm(1, 1, 1, dest.cols()) += dest.subm(i, i, 1, dest.cols());
2630     }
2631     dest ← transpose(dest.subm(1, 1, 1, dest.cols()));
2632     return dest;
2633   }

```

431. Finally the Hessian: for a functional we have $\nabla^2 -\log(-f) = (\nabla f \nabla^T f)/f^2 - (\nabla f)^2/f$.

⟨Log barrier methods 427⟩ +≡

```

2634   matrix⟨T, symmetric, dense⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, dense⟩ &dest)
2635   {
2636   #define hess_aux dense_aux
2637     ⟨Compute Hessian for log barrier 432⟩;
2638   #undef hess_aux
2639   }
2640   matrix⟨T, symmetric, sparse⟩ &hess(const matrix⟨T, unstructured, dense⟩
      &x, matrix⟨T, symmetric, sparse⟩ &dest)
2641   {
2642   #define hess_aux sparse_aux
2643     ⟨Compute Hessian for log barrier 432⟩;
2644   #undef hess_aux
2645   }

```

```

432.  ⟨ Compute Hessian for log barrier 432 ⟩ ≡
2646  T t ← 0;
2647  matrix⟨T, unstructured, dense⟩ X ← x;    /* We may resize x. */
2648  if (has_t) {
2649      t ← X(X.rows());
2650      X.resize(X.rows() - 1, 1);
2651  }
2652  if (fctnal) {
2653      T fval ← fctnal-eval(X) - t;
2654      fctnal-grad(X, aux);
2655      fctnal-hess(X, dest);
2656      if (has_t) {
2657          aux.resize(aux.rows() + 1, 1);
2658          dest.resize(dest.rows() + 1, dest.rows() + 1);
2659          aux.entry(aux.rows()) ← T(-1);
2660          for (index i ← 1; i ≤ dest.cols(); ++i) {
2661              dest.entry(dest.rows(), i) ← 0;
2662              dest.entry(i, dest.cols()) ← 0;
2663          }
2664      }
2665      aux /= fval;
2666      dest /= -fval;
2667      dest += outerp(aux, &hess_aux);
2668      return dest;
2669  }
2670  dest.init(X.rows() + (has_t ? 1 : 0), X.rows() + (has_t ? 1 : 0));
2671  dest.fillwith(0);
2672  fct-eval(X, aux);
2673  fct-jacobian(X, jac);
2674  if (has_t) {
2675      jac.resize(jac.rows(), jac.cols() + 1);
2676      jac.subm(1, jac.rows(), jac.cols()) ← matrix⟨T, unstructured, dense⟩(jac.rows(), 1, T(-1));
2677  }
2678  for (index i ← 1; i ≤ aux.rows(); ++i) {
2679      jac.subm(i, i, 1, jac.cols()) /= aux(i) - t;
2680      fct-hess(X, hess_aux, i);
2681      if (has_t) dest.subm(1, dest.rows() - 1, 1, dest.cols() - 1) -= (hess_aux /= t - aux(i));
2682      else dest -= (hess_aux /= aux(i));
2683      dest += outerp(jac.subm(i, i, 1, jac.cols()), &hess_aux);
2684  }
2685  return dest;

```

This code is used in section 431.

433. Sequential unconstrained minimization. function defined, we are ready to tackle our main problem, namely

$$\begin{aligned} &\text{minimize} && f_0(x) \\ &\text{subject to} && f_i(x) < 0, \quad i = 1, \dots, p \\ & && g_i(x) = 0, \quad i = 1, \dots, q. \end{aligned}$$

The method used to solve this problem is to minimize, for increasing values of t , the functional

$$\mu f_0(x) + \sum \phi_i(x),$$

where ϕ_i are barrier functions related to the inequality constraints. The equality constraints are taken care by the search direction function — the method doesn't do anything to enforce them. The algorithms that work this way are called *Sequential Unconstrained Minimization Techniques*, or SUMT for brevity.

/* Empty, waiting for **export** */

434. `<sumt.h 434>` \equiv

```
2686 #ifndef __MATH_SUMT__
2687 #define __MATH_SUMT__
2688 #include <vector>
2689 #include <math/fmin.h>
2690 #include <math/functional/linear.h>
2691 #include <math/functional/gaxpy.h>
2692 #include <math/barrier/log.h>
2693 #include <math/linesearchbase.h>
2694 namespace math {
2695     <SUMT functions 435>;
2696 }
2697 #endif
```

435. The arguments for the function are as follows: first we have the objective functional and a list of barrier functions. If the objective functional is zero, a feasibility problem will be solved. We have to enforce that the list is actually of barrier functions because we'll need information about dual variables in order to determine that our solution is precise enough, to detect infeasibility and so on. Next, the user provides an initial point, and line minimization and search direction algorithms. The last parameters are the absolute and relative precisions, followed by a pointer to a function that gets as arguments the current phase, the value of the objective function (primal) and the dual slack.

```

2698   <SUMT functions 435> ≡
2699   template<class T>
2699   void sumt (functional::base<T> *obj, vector<barrier::base<T> *> &barrier_functions, matrix<T,
2700   unstructured, dense> &x, linesearch::base<T> * line , searchdir::base<T> *dir, T
2701   abstol, T reltol, void(*disp)(int, double, double) ← 0 )
2702   {
2703   #ifdef __MATH_DEBUG__
2704       cout << "[math]:_sumt_begins.\n";
2705   #endif
2706
2707       T mu;
2708       functional::linear<T> f;
2709       typedef typename vector<barrier::base<T> *>::iterator sumt_iterator;
2710       <Build unconstrained objective function 436>;
2711       bool initial_point_is_feasible ← true;
2712       for (sumt_iterator i ← barrier_functions.begin();
2713           i ≠ barrier_functions.end() ∧ initial_point_is_feasible; ++i)
2714           try {
2715               (*i)→eval(x);
2716           }
2717       catch(error::domain)
2718       {
2719           initial_point_is_feasible ← false;
2720       }
2721       if (¬initial_point_is_feasible) {
2722           <SUMT Phase one 438>;
2723       }
2724       <SUMT Phase two 441>;
2725   #ifdef __MATH_DEBUG__
2726       cout << "[math]:_sumt_ends.\n";
2727   #endif
2728   }
```

This code is used in section 434.

436. Before doing anything we build the functional we'll minimize. We use the **linear** functional, which has provision for weighting.

```

2725   <Build unconstrained objective function 436> ≡
2726   f.add(obj);
2727   for (sumt_iterator i ← barrier_functions.begin(); i ≠ barrier_functions.end(); ++i) f.add(*i);
```

This code is used in section 435.

437. At this point we know something is wrong with the initial point. The problem is obviously not feasible here, but there can be two reasons: one is that one of the constraints is positive, so the barrier function throws a **domain** error. But it can happen that the functional called by the barrier function throws a **domain** error too. Consider that the constraint is $\sqrt{x} > 1$ and that the initial point is -1 . A log barrier function would call $\sqrt{-1}$, which would throw a **domain** error. The first type of infeasibility is easy to deal with through a Phase I approach. The other is hard in this phase, so we assume the user will provide good enough a initial point. If not, the next piece of code will end up throwing **domain**, which will flag the problem.

438. Ok. Now that we don't have a feasible point, we modify the problem so it becomes feasible. The standard way to do that is to modify all constraints of the form $f(x) < 0$ to $f(x) < t$, where t is an auxiliary variable. The *addt* function of barrier functions update the constraint and returns the minimum value of t necessary in order for the constraint to be feasible with the x provided. At the end of the loop we know that *aux_var* will be positive (the problem was infeasible), so we can multiply it by two in order to obtain an interior point.

⟨SUMT Phase one 438⟩ ≡

```
2727   T aux_var ← 0;
2728   for (sumt_iterator  $i \leftarrow$  barrier_functions.begin();  $i \neq$  barrier_functions.end(); ++ $i$ )
      aux_var ← max(aux_var, (*i)-addt(x));
2729   aux_var ← aux_var + aux_var;
```

See also sections 439 and 440.

This code is used in section 435.

439. By now we have an interior feasible point for the augmented system. What we have to do is to build the initial point vector accordingly and to build the unconstrained function to be minimized.

⟨SUMT Phase one 438⟩ +≡

```
2730    $x.resize(x.rows() + 1, 1)$ ;
2731    $x.entry(x.rows()) \leftarrow aux\_var$ ;
2732   matrix⟨T, unstructured, dense⟩ phasei_cost( $x.rows()$ , 1);
2733    $phasei\_cost.entry(x.rows()) \leftarrow 1.0$ ;
2734   functional::gaxpy⟨T, unstructured, dense⟩ phasei( $phasei\_cost$ , 0.0);
2735    $f.set\_term(0, \&phasei)$ ;
```

440. We now try to find a feasible point for the original problem. We don't have to go all the way, just enough so that the final value of t is not too close to zero, otherwise the initial point for phase II will be barely feasible and numerical problems could arise. If we find the global optimum and $t \geq 0$, then the problem is infeasible. We also compute the Lagrange dual function in order to test for infeasibility. The value of the Lagrange dual function depends on the barrier functions being used.

⟨SUMT Phase one 438⟩ +≡

```

2736     mu ← 1;
2737     T cost ← 2 * abstol;
2738     T lagrange ← 0; while (cost - lagrange ≥ abstol ∧ x(x.rows()) > -0.1) { f.set_weight(0, mu); try {
           x ← fmin (&f, x, line , dir, abstol, reltol, 0.0) ; } catch(error::maxiterations e)
2739     { }
2740     cost ← x(x.rows());
2741     lagrange ← cost;
2742     for (int i ← 1; i ≠ f.size(); ++i)
           lagrange += static_cast<barrier::base<T> *>(f.get_term(i))-dual_times_f(x)/mu;
2743     if (disp) disp(1, cost, cost - lagrange);
2744     if (lagrange ≥ 0) throw error::infeasible();
2745     mu ← mu * 50; }
2746     if (x(x.rows()) ≥ 0) throw error::infeasible();
2747     x.resize(x.rows() - 1, 1);
2748     for (sumt_iterator i ← barrier_functions.begin(); i ≠ barrier_functions.end(); ++i) (*i)-delt();

```

441. Ready we are for phase II. We have to restore the original cost function that was overwritten in phase I. After that, we proceed almost exactly as in phase I.

⟨SUMT Phase two 441⟩ ≡

```

2749     if (¬obj) return; /* It was a feasibility problem. */
2750     f.set_term(0, obj);
2751     bool stop ← false;
2752     mu ← 1; while (¬stop) { f.set_weight(0, mu); x ← fmin (&f, x, line , dir, abstol, reltol, 0.0) ;
2753     T cost ← obj-eval(x);
2754     T slack ← 0;
2755     for (int i ← 1; i ≠ f.size(); ++i)
           slack -= static_cast<barrier::base<T> *>(f.get_term(i))-dual_times_f(x)/mu;
2756     if (disp) disp(2, cost, slack);
2757     mu ← mu * 50;
2758     stop ← true;
2759     #ifdef __MATH_DEBUG
2760     cout << "[math]:";
2761     cout << "slack=" << slack << ", cost=" << cost << ", abstol=" << abstol << ", reltol=" << reltol <<
           '\n';
2762     #endif
2763     if (abstol) stop ← (slack ≤ abstol);
2764     if (reltol ∧ cost ∧ stop) stop ← (fabs(slack/cost) ≤ reltol);
2765     }

```

This code is used in section 435.

442. Index. Here is a cross-reference table for MATH. Underlined entries correspond to where the identifier was declared.

__MATH__: 3.
 __MATH_ALGEBRA__: 118.
 __MATH_BACKTRACKING_LINESEARCH__: 381.
 __MATH_BARRIER__: 416.
 __MATH_BISECTION__: 374.
 __MATH_CHOLESKY__: 211.
 __MATH_DEBUG: 441.
 __MATH_DEBUG__: 384, 385, 401, 403, 408, 435.
 __MATH_DET__: 243.
 __MATH_ENTROPY_FUNCTIONAL__: 324.
 __MATH_ERF_FUNCTIONAL__: 343.
 __MATH_EXP_FUNCTIONAL__: 353.
 __MATH_EYE__: 238.
 __MATH_FMIN__: 412.
 __MATH_FN__: 258.
 __MATH_FSTREAM__: 154.
 __MATH_FUNCTION__: 245.
 __MATH_GAXPY_FUNCTION__: 251.
 __MATH_GAXPY_FUNCTIONAL__: 264.
 __MATH_INSANE_DEBUG__: 24, 36, 53.
 __MATH_LINEAR_FUNCTIONAL__: 271.
 __MATH_LINESEARCH__: 371.
 __MATH_LOG_BARRIER__: 426.
 __MATH_LOG_FUNCTIONAL__: 362.
 __MATH_LU__: 197.
 __MATH_NORM2ERR_FUNCTIONAL__: 290.
 __MATH_ONES__: 240.
 __MATH_POWER_FUNCTIONAL__: 315.
 __MATH_PROD_FUNCTIONAL__: 298.
 __MATH_QR__: 225.
 __MATH_QUADRATIC_FUNCTIONAL__: 283.
 __MATH_RATIO_FUNCTIONAL__: 307.
 __MATH_RELATIVE_ENTROPY_FUNCTIONAL__: 333.
 __MATH_SEARCHDIR_EQUALITY__: 397.
 __MATH_SEARCHDIR_EQUALITY_NEWTON__: 405.
 __MATH_SEARCHDIR_GRADIENT__: 390.
 __MATH_SEARCHDIR_NEWTON__: 393.
 __MATH_SEARCHDIRBASE__: 387.
 __MATH_SPARSE__: 136.
 __MATH_SUMT__: 434.
 __MATH_SYMMETRIC__: 147.
 A: 127, 128, 129, 130, 134, 199, 202, 203, 205,
 209, 212, 215, 216, 218, 221, 222, 223, 226,
 229, 243, 253, 266, 401, 406.
 a: 122, 253, 266, 382, 409.
 Ablock: 228.
 abs: 197, 200.
 abstol: 412, 414, 435, 440, 441.
 add: 273, 436.
 addt: 420, 421, 438.
 alpha: 381, 382, 384.
 assign: 165.
 At: 233, 234, 235.
 atxmul: 128, 228, 231, 235, 294.
 aux: 106, 127, 128, 131, 132, 194, 243, 278, 279,
 291, 294, 295, 296, 302, 303, 305, 310, 311,
 313, 321, 322, 330, 331, 349, 351, 359, 360,
 368, 369, 421, 426, 428, 430, 432.
 aux.var: 438, 439.
 auxdense: 339, 340.
 auxgrad: 339, 341.
 auxhess: 340, 341.
 auxsparse: 339, 340.
 auxvec: 337, 338, 341.
 axmul: 127, 286.
 ay: 252, 253, 254, 255, 265, 266, 267, 268.
 B: 202, 203, 215, 216, 229, 253, 266.
 b: 253, 266, 382.
 backtracking: 381, 382.
 backtracking::minimize: 383, 384.
 backtracking_infty: 381, 385.
 backup: 80.
 barrier: 416, 426, 435, 440, 441.
 barrier::log::grad: 430.
 barrier::log::hess: 431.
 barrier::addt: 438.
 barrier::delt: 440.
 barrier_functions: 435, 436, 438, 440.
 base: 245, 246, 251, 258, 259, 264, 271, 272, 273,
 275, 283, 290, 291, 292, 298, 299, 307, 308,
 315, 316, 324, 325, 333, 334, 343, 344, 353,
 354, 362, 363, 371, 372, 374, 376, 377, 378,
 381, 383, 384, 387, 388, 390, 391, 393, 395,
 397, 398, 399, 403, 405, 408, 412, 416, 417,
 418, 426, 427, 435, 440, 441.
 Bblock: 231, 235.
 bee: 252, 253, 265, 266, 267.
 begin: 140, 205, 209, 276, 277, 279, 281, 435,
 436, 438, 440.
 beta: 134, 226, 227, 228, 381, 382, 384, 385.
 binary: 158, 159, 183, 184.
 bisection: 374, 375, 378.
 bisection::minimize: 376, 377.
 bisection_infty: 374, 379.
 c_str: 180, 191.
 Char: 161, 168, 188, 192.
 character: 219.
 cholesky: 211, 395.
 col: 58, 59, 61, 62, 63, 64, 69, 71, 91, 142, 143, 150.

- cols*: [18](#), [19](#), [24](#), [26](#), [30](#), [32](#), [33](#), 35, [36](#), 39, 44, 46, 47, 51, [76](#), [77](#), [79](#), [80](#), 86, [87](#), 89, 90, 94, 96, 98, 99, 101, 102, 104, 105, 108, 115, [116](#), 121, 122, 123, 124, 125, 127, 128, 129, 130, 131, 132, 133, 134, [140](#), [144](#), [149](#), [162](#), 165, 167, 170, 175, 190, 191, 194, 199, 202, 206, 212, 215, 219, 226, 228, 229, [238](#), [240](#), 401, 408, 409, 430, 432.
- col1*: [85](#), [92](#), [93](#).
- col2*: [85](#), [92](#), [93](#).
- complex**: 166, 169, 170, 189, 193, 194.
- copyfrom*: 40, 43, [46](#), 47, 50, 52, 254, 287.
- cost*: [440](#), [441](#).
- cout*: 24, 36, 53, 384, 385, 401, 403, 408, 435, 441.
- data*: [28](#), 29, 30, 32, 33, 48, 52, 55, 56, 61, 64, 80.
- data_size*: [175](#).
- decomp*: [400](#), 401, 403.
- decompose*: [199](#), 203, [205](#), [209](#), [212](#), 216, [218](#), [221](#), [222](#), [223](#), [226](#), 230, 233, 243, 401.
- delt*: 420, [422](#), 440.
- dense**: [10](#), 12, 26, [29](#), [30](#), 33, [48](#), 52, [55](#), 60, 79, 80, 83, 116, 140, 205, 209, 218, 221, 222, 223, 226, 229, 231, 235, 247, 248, 249, 252, 253, 254, 255, 256, 260, 261, 262, 267, 268, 269, 277, 278, 279, 280, 281, 283, 284, 285, 286, 287, 288, 291, 292, 293, 294, 295, 296, 301, 302, 303, 304, 305, 309, 310, 311, 312, 313, 318, 320, 321, 322, 327, 329, 330, 331, 336, 337, 338, 339, 340, 346, 348, 349, 350, 356, 358, 359, 360, 365, 367, 368, 369, 372, 376, 377, 378, 383, 384, 388, 391, 395, 403, 408, 412, 413, 421, 423, 424, 426, 428, 429, 430, 431, 432, 435, 439.
- dense::destroy**: 55, [56](#).
- dense::get**: [64](#).
- dense::init**: 30, [32](#), 116.
- dense::memory**: [75](#).
- dense::resize**: [79](#).
- dense::set**: [61](#).
- dense_aux*: [349](#), 350, [426](#), 431.
- dest*: [127](#), [128](#), [131](#), [132](#), [133](#), [161](#), [166](#), 167, [169](#), 170, [238](#), [247](#), [248](#), [249](#), [254](#), [255](#), [256](#), [261](#), [262](#), [268](#), [269](#), [279](#), [281](#), [287](#), [288](#), [294](#), [295](#), [296](#), [303](#), [305](#), [311](#), [313](#), [320](#), [322](#), [329](#), [331](#), [338](#), [340](#), 341, [348](#), [350](#), 351, [358](#), [360](#), [367](#), [369](#), [388](#), [391](#), [395](#), [403](#), [408](#), [430](#), [431](#), 432.
- destroy*: 55, [56](#).
- det*: 8, [243](#).
- detach*: [57](#), 58, 72, 73, 74, 76.
- determinant*: [243](#).
- dgetrf*: 205.
- dimension**: [9](#), 77, 86, 99, 101, 102, 104, 116, 121, 122, 123, 124, 125, 127, 128, 129, 130, 133, 134.
- dir*: 372, [376](#), [377](#), [378](#), 379, [383](#), [384](#), 385, [388](#), [391](#), [395](#), [403](#), [408](#), [413](#), [435](#), 440, 441.
- dirf*: [398](#), 399, 401, 403.
- dirnorm*: [378](#).
- disp*: 412, 413, 435, 440, 441.
- domain**: [9](#), 322, 327, 329, 331, 336, 338, 341, 365, 379, 385, 428, 435, 437.
- dot*: [123](#), 124, 125, 126, 130, 132, 227, 231, 235, 267, 286, 293, 384.
- Double*: [168](#), [192](#).
- double_precision**: 205, 218, 222.
- dpotrf*: 222.
- dpptrf*: 218.
- dual*: [423](#), [429](#).
- dual_times_f*: [424](#), [429](#), 440, 441.
- e*: [243](#), [379](#), [385](#), [395](#), [408](#), [440](#).
- elem*: [273](#).
- element**: [65](#), 66, [68](#), [69](#), 71, 81, 106.
- element::←**: [70](#).
- element::value**: [107](#).
- element_type**: [10](#), [11](#), [13](#), [16](#), 26, 28, 30, 33, 34, 35, 36, 48, 52, 58, 59, 61, 62, 63, 64, [67](#), 70, 75, [83](#), 87, 91, [95](#), 97, 98, 100, 103, 106, 107, 123, 124, 125, 126, 134, [272](#), 273, [299](#), 300, [308](#), 309, [316](#), 317, [325](#), 326, [334](#), 335, [344](#), 345, [354](#), 355, [363](#), 364.
- elements*: [28](#), 29, 30, 33, 48, 56, 75, 80, [138](#), 140, 141, 142, 143, [272](#), 273, 274, 275, 276, 277, 279, 281.
- elment*: [273](#).
- end*: 140, 142, 143, 276, 277, 279, 281, 435, 436, 438, 440.
- entr**: [324](#), [326](#).
- entry*: [71](#), 94, 98, 101, 104, 108, 127, 128, 131, 132, 133, 201, 202, 213, 214, 215, 227, 231, 232, 234, 235, 238, 401, 403, 408, 430, 432, 439.
- equality**: [397](#), [399](#), [401](#).
- equality**: [405](#).
- erase*: 142, 276.
- erf**: [343](#), [345](#), 346.
- error**: [5](#), 7, 25, 51, 52, 54, 68, 85, 86, 99, 101, 102, 104, 116, 121, 122, 123, 124, 125, 127, 128, 129, 130, 133, 134, 149, 167, 170, 173, 180, 199, 201, 202, 207, 212, 213, 215, 219, 220, 232, 234, 243, 322, 327, 329, 331, 336, 338, 341, 365, 378, 379, 384, 385, 395, 408, 428, 435, 440.
- eval*: [247](#), [254](#), [260](#), [267](#), [277](#), [286](#), [293](#), 294, 295, 296, [301](#), 303, 305, 309, 311, 313, [318](#), 320, 322, [327](#), 329, 331, [336](#), 338, 341, [346](#), 348, 351, [356](#), 358, 360, [365](#), [367](#), 369, 378, 379, 384, 385, 420, 421, [428](#), 429, 430, 432, 435, 441.
- exp**: 348, 351, [353](#), [355](#), 356.

- exponent*: [316](#), [317](#), [318](#), [320](#), [322](#).
extra_component: [403](#).
eye: [237](#), [238](#).
F: [291](#).
f: [292](#), [299](#), [308](#), [316](#), [325](#), [334](#), [344](#), [354](#), [363](#),
[376](#), [378](#), [383](#), [388](#), [391](#), [395](#), [399](#), [403](#), [408](#),
[412](#), [418](#), [435](#).
fabs: [378](#), [414](#), [441](#).
false: [41](#), [59](#), [64](#), [112](#), [150](#), [408](#), [409](#), [413](#), [418](#),
[422](#), [435](#), [441](#).
fast_assignment: [40](#), [41](#), [42](#), [43](#).
fct: [417](#), [418](#), [421](#), [428](#), [430](#), [432](#).
fctnal: [417](#), [418](#), [421](#), [428](#), [430](#), [432](#).
fd: [157](#), [182](#).
fheader: [162](#), [164](#), [165](#), [167](#), [170](#), [173](#), [174](#),
[175](#), [180](#).
file_flag: [165](#).
filerr: [9](#), [167](#), [170](#).
fillwith: [34](#), [35](#), [36](#), [256](#), [269](#), [279](#), [281](#), [322](#),
[401](#), [403](#), [432](#).
find: [142](#), [143](#).
finish: [202](#), [203](#), [215](#), [216](#), [403](#).
finite: [201](#), [202](#), [213](#), [215](#), [232](#), [234](#), [365](#), [379](#), [385](#).
first: [112](#), [275](#), [277](#), [279](#), [281](#).
fmin: [412](#), [440](#), [441](#).
fortran: [205](#), [206](#), [209](#), [218](#), [219](#), [221](#), [222](#), [223](#).
ftreshold: [384](#), [385](#).
func: [377](#), [378](#), [379](#), [384](#), [385](#).
function: [245](#), [251](#), [291](#), [292](#), [417](#), [418](#), [427](#).
function::base::grad: [248](#).
function::base::hess: [249](#).
function::gaxpy::A: [253](#).
function::gaxpy::b: [253](#).
function::gaxpy::eval: [254](#).
function::gaxpy::hess: [256](#).
function::gaxpy::jacobian: [255](#).
function::eval: [293](#), [294](#), [295](#), [296](#), [428](#), [430](#), [432](#).
function::hess: [295](#), [296](#), [432](#).
function::jacobian: [294](#), [295](#), [296](#), [430](#), [432](#).
functional: [258](#), [264](#), [271](#), [283](#), [290](#), [298](#), [307](#),
[315](#), [324](#), [333](#), [343](#), [353](#), [362](#), [372](#), [376](#), [377](#),
[378](#), [383](#), [384](#), [388](#), [391](#), [395](#), [403](#), [408](#), [412](#),
[416](#), [417](#), [418](#), [427](#), [435](#), [439](#).
functional::base::grad: [261](#).
functional::base::hess: [262](#).
functional::gaxpy::A: [266](#).
functional::gaxpy::b: [266](#).
functional::gaxpy::grad: [268](#).
functional::gaxpy::hess: [269](#).
functional::norm2err::eval: [293](#).
functional::norm2err::grad: [294](#).
functional::norm2err::hess: [295](#), [296](#).
functional::quadratic::eval: [286](#).
functional::quadratic::grad: [287](#).
functional::quadratic::hess: [288](#).
functional::eval: [277](#), [384](#), [385](#), [428](#), [430](#), [432](#),
[435](#).
functional::grad: [279](#), [383](#), [413](#), [414](#), [430](#), [432](#).
functional::hess: [281](#), [432](#).
fval: [291](#), [293](#), [294](#), [295](#), [296](#), [336](#), [338](#), [341](#),
[384](#), [385](#), [432](#).
g: [299](#), [308](#), [334](#).
garbage: [175](#).
gatxpy: [130](#).
gaxpy: [129](#), [250](#), [251](#), [252](#), [253](#), [254](#), [263](#), [264](#),
[265](#), [266](#), [285](#), [287](#), [439](#).
generic: [5](#), [7](#), [9](#), [25](#), [51](#), [52](#), [54](#), [68](#), [85](#), [173](#),
[180](#), [207](#), [220](#).
generic::message: [8](#).
get: [59](#), [62](#), [63](#), [64](#), [86](#), [87](#), [91](#), [94](#), [95](#), [96](#), [97](#), [99](#),
[100](#), [101](#), [102](#), [103](#), [104](#), [107](#), [115](#), [122](#), [129](#), [130](#),
[134](#), [143](#), [161](#), [165](#), [168](#), [175](#), [200](#), [201](#), [202](#), [212](#),
[213](#), [214](#), [215](#), [227](#), [232](#), [234](#), [408](#).
get_term: [273](#), [275](#), [440](#), [441](#).
get_weight: [273](#), [275](#).
goodbit: [165](#), [167](#), [170](#), [173](#), [180](#).
grad: [261](#), [268](#), [279](#), [287](#), [294](#), [303](#), [305](#), [311](#), [313](#),
[320](#), [322](#), [329](#), [331](#), [338](#), [341](#), [348](#), [351](#), [358](#), [360](#),
[367](#), [369](#), [383](#), [391](#), [395](#), [408](#), [413](#), [414](#), [430](#), [432](#).
gradient: [390](#).
gtol: [412](#), [413](#), [414](#).
gtreshold: [384](#), [385](#).
gval: [336](#), [338](#), [341](#).
g0: [376](#), [383](#), [384](#).
H: [291](#), [406](#).
has_t: [407](#), [408](#), [409](#), [418](#), [420](#), [421](#), [422](#), [428](#),
[430](#), [432](#).
HAVE_LIBLAPACK: [217](#).
Hdense: [280](#), [281](#), [304](#), [305](#), [312](#), [313](#), [321](#), [322](#),
[330](#), [331](#), [359](#), [360](#), [368](#), [369](#).
header: [164](#), [165](#), [167](#), [170](#), [173](#), [174](#), [175](#), [180](#).
hess: [249](#), [256](#), [262](#), [269](#), [281](#), [288](#), [295](#), [296](#), [305](#),
[313](#), [322](#), [331](#), [340](#), [341](#), [350](#), [351](#), [360](#), [369](#),
[394](#), [395](#), [408](#), [431](#), [432](#).
hess_aux: [350](#), [351](#), [431](#), [432](#).
Hsparse: [280](#), [281](#), [304](#), [305](#), [312](#), [313](#), [321](#), [322](#),
[330](#), [331](#), [359](#), [360](#), [368](#), [369](#).
HUGE_VAL: [374](#), [381](#).
i: [33](#), [35](#), [44](#), [52](#), [66](#), [80](#), [86](#), [87](#), [89](#), [90](#), [94](#), [96](#), [98](#),
[99](#), [101](#), [102](#), [104](#), [105](#), [108](#), [115](#), [122](#), [123](#), [124](#),
[125](#), [127](#), [128](#), [129](#), [130](#), [131](#), [132](#), [133](#), [134](#), [142](#),
[143](#), [161](#), [165](#), [167](#), [170](#), [188](#), [190](#), [191](#), [194](#),
[200](#), [201](#), [202](#), [208](#), [214](#), [215](#), [232](#), [234](#), [238](#),
[243](#), [249](#), [273](#), [275](#), [277](#), [279](#), [281](#), [295](#), [296](#),

- 375, 384, 385, 401, 403, 408, 414, 421, 428,
 430, 432, 435, 436, 438, 440, 441.
iflag: 162, 164, 165, 167, 170, 175, 190, 191, 194.
ifstream: 155, 157, 158, 159, 160, 161, 165, 166,
 169, 173, 175, 177, 179, 185.
ifstream::open: 159.
ifstream::skip: 172.
ifstream::skip_data: 174.
ifstream::skipto: 176, 178.
imag: 194.
in: 158, 159, 184.
index: 4, 17, 18, 19, 24, 26, 27, 28, 30, 32, 33, 35,
 36, 38, 39, 44, 52, 58, 59, 60, 61, 62, 63, 64, 66,
 69, 71, 76, 77, 78, 79, 80, 84, 85, 86, 87, 89, 90,
 91, 92, 93, 94, 96, 98, 99, 101, 102, 104, 105,
 108, 110, 115, 116, 122, 123, 124, 125, 127,
 128, 129, 130, 131, 132, 133, 134, 140, 142,
 143, 144, 149, 150, 190, 194, 199, 200, 201,
 202, 203, 205, 208, 209, 212, 213, 214, 215,
 226, 229, 231, 232, 234, 235, 238, 240, 243,
 249, 256, 295, 296, 384, 385, 400, 401, 403,
 407, 408, 414, 421, 428, 430, 432.
infeasible: 9, 440.
init: 30, 32, 33, 38, 39, 44, 48, 76, 79, 80, 89,
 116, 167, 170, 261, 269, 432.
initial_point_is_feasible: 435.
inner: 202, 215, 232, 234.
inner_iter: 378.
insert: 140.
instances: 22, 37, 39, 43, 46, 53, 57.
integer: 206, 219.
internal_matrix_type: 83, 84, 85, 86, 90.
ios: 158, 159, 183, 184.
ipart: 167, 170, 194.
iter: 384, 413.
iterator: 272, 277, 279, 281, 435.
i1: 84, 85, 86, 87, 90, 91, 93, 98, 101, 104.
i2: 84, 85, 86, 87, 90.
j: 35, 44, 66, 80, 86, 87, 89, 90, 94, 96, 98, 99, 101,
 102, 104, 105, 108, 115, 122, 127, 128, 129, 130,
 131, 132, 133, 134, 167, 170, 190, 194, 201, 202,
 213, 214, 215, 226, 231, 232, 234, 249, 401, 408.
jac: 426, 432.
jacobian: 248, 255, 294, 295, 296, 430, 432.
j1: 84, 85, 86, 87, 90, 91, 93, 98, 101, 104.
j2: 84, 85, 86, 87, 90.
k: 108, 127, 128, 168, 192, 199, 202, 212, 215,
 232, 234, 235.
l: 229.
lagrange: 440.
LAPACK: 204.
less: 110, 138.
linear: 271, 273, 276, 435, 436.
linesearch: 371, 374, 381, 412, 435.
linesearch::bisection::minimize: 376.
linesearch::minimize: 413.
log: 327, 329, 331, 336, 338, 341, 362, 364,
 365, 426, 427, 428.
Long: 161, 188.
lsearch: 412, 413.
lu: 197, 243, 395, 401, 403, 408.
lu::decompose: 401.
lu::finish: 403.
M: 44, 81, 406.
m: 127, 128, 134, 206, 219, 229, 382.
M_PI: 348, 351.
map: 110, 138.
mat: 69, 85, 86, 89, 90.
math: 3, 4, 41, 110, 111, 112, 118, 136, 138, 139,
 142, 143, 147, 149, 150, 154, 161, 165, 167,
 170, 173, 175, 177, 179, 180, 186, 188, 190,
 194, 197, 211, 225, 238, 240, 243, 245, 251,
 258, 264, 271, 283, 290, 298, 307, 315, 324,
 333, 343, 353, 362, 371, 374, 381, 387, 390,
 393, 397, 405, 412, 416, 426, 434.
matrix: 12, 16, 17, 18, 23, 24, 35, 36, 37, 39, 43,
 44, 45, 53, 65, 82, 83, 89, 105, 109, 116, 119,
 120, 121, 127, 128, 131, 132, 133, 166, 169, 189,
 193, 199, 202, 203, 205, 209, 212, 215, 216, 218,
 221, 222, 223, 226, 228, 229, 231, 233, 235, 238,
 240, 243, 247, 248, 249, 252, 253, 254, 255, 256,
 260, 261, 262, 265, 266, 267, 268, 269, 277, 278,
 279, 280, 281, 284, 285, 286, 287, 288, 291, 292,
 293, 294, 295, 296, 301, 302, 303, 304, 305, 309,
 310, 311, 312, 313, 318, 320, 321, 322, 327, 329,
 330, 331, 336, 337, 338, 339, 340, 346, 348, 349,
 350, 356, 358, 359, 360, 365, 367, 368, 369, 372,
 376, 377, 378, 383, 384, 388, 391, 394, 395, 400,
 401, 403, 406, 408, 409, 412, 413, 421, 423, 424,
 426, 428, 429, 430, 431, 432, 435, 439.
matrix::(): 62.
matrix::←: 42, 44.
matrix::*=: 279, 281, 385.
matrix::+=: 385.
matrix::/=: 430.
matrix::=: 385.
matrix::cols: 19, 38, 44, 46, 86, 89, 94, 96, 115,
 116, 122, 123, 124, 125, 127, 128, 129, 130,
 131, 132, 226, 255, 401.
matrix::copyfrom: 46.
matrix::detach: 57, 58, 73, 74, 76.
matrix::entry: 71, 255, 401, 430.
matrix::fillwith: 34, 36, 256, 279, 281, 401,
 403, 432.

- matrix::get:** [62](#), [86](#), [91](#), [94](#), [96](#), [115](#).
matrix::init: [38](#), [44](#), [76](#), [89](#), [256](#), [432](#).
matrix::numnonzeros: [115](#).
matrix::rep: [74](#), [116](#).
matrix::resize: [76](#), [256](#), [279](#), [281](#), [401](#), [403](#),
[428](#), [430](#), [432](#).
matrix::rows: [19](#), [38](#), [44](#), [46](#), [86](#), [89](#), [94](#), [96](#), [115](#),
[116](#), [122](#), [123](#), [124](#), [125](#), [127](#), [128](#), [129](#), [130](#),
[131](#), [132](#), [226](#), [256](#), [279](#), [281](#), [295](#), [296](#), [401](#),
[403](#), [414](#), [428](#), [430](#), [432](#).
matrix::set: [44](#), [58](#), [70](#), [86](#), [89](#), [90](#), [91](#), [94](#), [96](#).
matrix::storg: [73](#), [116](#).
matrix::subm: [92](#), [93](#), [401](#), [403](#), [430](#), [432](#).
matrix::swaprows: [108](#).
matrix_name: [178](#), [179](#), [180](#), [181](#), [186](#), [191](#).
matrix_simple_template: [12](#), [13](#), [35](#), [39](#), [43](#), [105](#),
[109](#), [119](#), [120](#), [121](#), [127](#), [128](#), [131](#), [132](#), [133](#), [166](#),
[169](#), [189](#), [193](#), [212](#), [238](#), [240](#), [243](#), [251](#), [264](#), [397](#).
matrix_template: [44](#), [45](#).
matrix_type: [16](#), [37](#), [42](#), [43](#), [44](#), [46](#), [65](#), [66](#), [67](#),
[69](#), [71](#), [83](#), [88](#), [89](#), [94](#), [96](#), [99](#), [101](#), [102](#), [104](#).
max: [85](#), [92](#), [93](#), [438](#).
max_cols: [28](#), [29](#), [33](#), [52](#), [80](#).
max_rows: [28](#), [29](#), [33](#), [52](#), [80](#).
maxiter: [374](#), [375](#), [378](#), [381](#), [382](#), [384](#), [412](#), [413](#).
maxiterations: [9](#), [384](#), [440](#).
memcpy: [31](#), [48](#), [52](#).
memory: [75](#), [205](#), [209](#), [218](#), [221](#), [222](#), [223](#).
memset: [30](#), [31](#).
message: [8](#).
min: [206](#), [208](#), [238](#).
minimize: [372](#), [376](#), [377](#), [378](#), [383](#), [384](#), [413](#).
mode: [158](#), [159](#), [183](#), [184](#), [218](#), [219](#), [221](#), [222](#), [223](#).
msg: [7](#).
mu: [227](#), [435](#), [440](#), [441](#).
n: [127](#), [128](#), [133](#), [134](#), [199](#), [202](#), [206](#), [212](#), [215](#),
[219](#), [229](#).
name: [158](#), [159](#), [162](#), [165](#), [180](#), [183](#), [184](#), [185](#), [186](#).
name_length: [165](#).
new_e: [317](#).
new_f: [300](#), [309](#), [317](#), [326](#), [335](#), [345](#), [355](#), [364](#).
new_g: [300](#), [309](#), [335](#).
newa: [401](#).
newf: [292](#), [378](#), [379](#), [399](#), [418](#), [427](#).
newp: [285](#).
newP: [285](#).
newpi: [285](#).
newton: [393](#), [405](#).
newy: [292](#).
nonpositivedef: [9](#), [212](#), [220](#), [395](#).
nonsquare: [9](#), [149](#), [199](#), [212](#), [219](#).
norm2: [126](#), [378](#), [414](#).
norm2err: [290](#), [292](#).
notimplemented: [9](#).
nrows: [235](#).
num_cols: [17](#), [18](#), [26](#), [28](#), [29](#), [30](#), [33](#), [38](#), [39](#),
[47](#), [48](#), [52](#), [77](#), [80](#).
num_instances: [21](#), [22](#), [26](#), [47](#), [54](#).
num_matrices: [172](#), [173](#).
num_rows: [17](#), [18](#), [26](#), [28](#), [29](#), [30](#), [33](#), [38](#), [39](#),
[47](#), [48](#), [52](#), [77](#), [80](#).
number: [167](#), [168](#), [170](#), [190](#), [192](#), [194](#).
numnonzeros: [115](#).
obj: [435](#), [436](#), [441](#).
ofstream: [181](#), [182](#), [183](#), [184](#), [185](#), [186](#), [187](#),
[188](#), [189](#), [193](#).
ofstream::open: [184](#).
ones: [239](#), [240](#).
open: [158](#), [159](#), [183](#), [184](#).
out: [183](#).
outer_iter: [378](#).
outerp: [131](#), [313](#), [322](#), [331](#), [341](#), [351](#), [360](#), [369](#), [432](#).
outerp_update: [134](#), [228](#), [231](#), [235](#).
outterp: [132](#), [295](#), [296](#).
P: [285](#).
p: [127](#), [128](#), [285](#).
pair: [110](#), [111](#), [112](#), [138](#), [139](#), [142](#), [143](#), [272](#).
parse_header: [164](#), [165](#), [167](#), [170](#), [173](#), [180](#).
pee: [284](#), [285](#), [286](#), [287](#).
Pee: [284](#), [285](#), [286](#), [287](#), [288](#).
permutations: [199](#), [200](#).
phasei: [439](#).
phasei_cost: [439](#).
pi: [285](#).
Pi: [284](#), [285](#), [286](#).
pivot: [199](#), [200](#).
pivots: [199](#), [200](#), [202](#), [203](#), [205](#), [206](#), [208](#), [209](#),
[243](#), [400](#), [401](#), [403](#).
pop_back: [274](#).
position: [176](#), [177](#), [179](#), [180](#).
pow: [318](#), [320](#), [322](#).
power: [315](#), [317](#).
preprocess: [59](#), [60](#), [63](#), [64](#), [150](#).
prod: [298](#), [300](#).
prot: [158](#), [159](#), [183](#), [184](#).
push_back: [273](#).
put: [188](#), [191](#), [192](#).
qr: [225](#).
quadratic: [283](#), [285](#).
r: [375](#).
rankdeficient: [9](#), [232](#), [234](#).
ratio: [307](#), [309](#), [373](#), [374](#), [375](#), [378](#), [379](#).
rdstate: [165](#), [167](#), [170](#), [173](#), [180](#).
real: [194](#), [209](#), [221](#), [223](#).

- relentr:** [333](#), [335](#).
reitol: [412](#), [414](#), [435](#), [440](#), [441](#).
rep: [74](#), [116](#).
rep_type: [14](#), [15](#), [24](#), [36](#), [39](#), [46](#), [50](#), [51](#), [57](#), [74](#).
representation: [13](#), [14](#), [25](#), [26](#), [47](#), [50](#), [54](#).
representation::cols: [18](#), [19](#), [47](#), [116](#).
representation::get: [62](#), [63](#).
representation::instances: [22](#), [37](#), [38](#), [43](#), [46](#),
[53](#), [57](#).
representation::resize: [76](#), [77](#).
representation::rows: [18](#), [19](#), [47](#), [116](#).
representation::set: [35](#), [58](#), [59](#).
representation::storg: [72](#), [73](#).
reshape: [18](#), [116](#).
resize: [76](#), [77](#), [78](#), [79](#), [80](#), [127](#), [128](#), [131](#), [132](#),
[133](#), [144](#), [149](#), [199](#), [206](#), [227](#), [230](#), [231](#), [235](#),
[248](#), [256](#), [279](#), [281](#), [322](#), [401](#), [403](#), [408](#), [428](#),
[430](#), [432](#), [439](#), [440](#).
result: [105](#), [115](#), [119](#), [120](#), [121](#), [123](#), [124](#), [125](#), [277](#),
[286](#), [311](#), [313](#), [318](#), [320](#), [322](#), [327](#), [329](#), [331](#), [348](#),
[351](#), [365](#), [367](#), [369](#), [403](#), [421](#), [428](#).
row: [58](#), [59](#), [61](#), [62](#), [63](#), [64](#), [69](#), [71](#), [91](#), [142](#),
[143](#), [150](#).
rows: [18](#), [19](#), [24](#), [26](#), [30](#), [32](#), [33](#), [35](#), [36](#), [39](#), [44](#), [46](#),
[47](#), [51](#), [76](#), [77](#), [79](#), [80](#), [86](#), [87](#), [89](#), [90](#), [94](#), [96](#),
[98](#), [99](#), [101](#), [102](#), [104](#), [105](#), [115](#), [116](#), [121](#), [122](#),
[123](#), [124](#), [125](#), [127](#), [128](#), [129](#), [130](#), [131](#), [132](#), [133](#),
[134](#), [140](#), [144](#), [149](#), [162](#), [165](#), [167](#), [170](#), [175](#), [190](#),
[191](#), [194](#), [199](#), [202](#), [206](#), [212](#), [215](#), [219](#), [226](#), [227](#),
[228](#), [229](#), [231](#), [238](#), [240](#), [243](#), [256](#), [269](#), [279](#), [281](#),
[286](#), [295](#), [296](#), [322](#), [384](#), [385](#), [401](#), [403](#), [408](#),
[414](#), [421](#), [428](#), [430](#), [432](#), [439](#), [440](#).
row1: [85](#), [92](#), [93](#).
row2: [85](#), [92](#), [93](#).
same: [109](#).
saxpy: [122](#), [129](#), [379](#).
sdir: [412](#), [413](#).
searchdir: [387](#), [390](#), [393](#), [397](#), [405](#), [412](#), [435](#).
searchdir::base::dir: [388](#).
searchdir::gradient::dir: [391](#).
searchdir::newton::dir: [395](#).
searchdir::dir: [403](#), [413](#).
second: [112](#), [275](#), [277](#), [279](#), [281](#).
seekg: [177](#), [179](#).
set: [35](#), [44](#), [57](#), [58](#), [59](#), [60](#), [61](#), [62](#), [65](#), [70](#), [71](#),
[86](#), [89](#), [90](#), [91](#), [94](#), [96](#), [99](#), [102](#), [105](#), [122](#), [129](#),
[130](#), [134](#), [142](#), [166](#), [167](#), [170](#).
set_A: [401](#), [409](#).
set_term: [275](#), [439](#), [441](#).
set_weight: [275](#), [440](#), [441](#).
sgetrf: [209](#).
sigma: [227](#).
singular: [9](#), [201](#), [202](#), [207](#), [213](#), [215](#), [243](#), [395](#), [408](#).
size: [191](#), [273](#), [275](#), [279](#), [281](#), [440](#), [441](#).
skip: [172](#), [173](#), [177](#).
skip_data: [172](#), [173](#), [174](#), [175](#), [180](#).
skipto: [176](#), [177](#), [178](#), [179](#), [185](#).
slack: [441](#).
solve: [203](#), [216](#), [229](#), [395](#), [408](#).
source: [37](#), [42](#), [43](#), [44](#), [45](#), [46](#), [47](#), [48](#), [51](#), [52](#), [140](#),
[188](#), [189](#), [190](#), [191](#), [193](#), [194](#).
sparse: [137](#), [140](#), [141](#), [249](#), [256](#), [262](#), [269](#), [280](#), [281](#),
[288](#), [296](#), [304](#), [305](#), [312](#), [313](#), [321](#), [322](#), [330](#), [331](#),
[339](#), [340](#), [349](#), [350](#), [359](#), [360](#), [368](#), [369](#), [426](#), [431](#).
sparse::get: [143](#).
sparse::resize: [144](#).
sparse::set: [142](#).
sparse_aux: [349](#), [350](#), [426](#), [431](#).
spotrf: [223](#).
spptrf: [221](#).
sqrt: [118](#), [126](#), [211](#), [213](#), [227](#), [348](#), [351](#).
status: [205](#), [206](#), [207](#), [208](#), [209](#), [218](#), [219](#), [220](#),
[221](#), [222](#), [223](#).
std: [155](#), [157](#), [158](#), [159](#), [170](#), [181](#), [182](#), [183](#),
[184](#), [272](#).
step: [373](#).
stepsize: [378](#), [379](#), [384](#), [385](#).
STO: [44](#).
STOA: [44](#), [45](#).
stop: [413](#), [414](#), [441](#).
storage: [12](#), [14](#), [16](#), [20](#), [26](#), [35](#), [39](#), [43](#), [47](#), [50](#),
[72](#), [105](#), [109](#), [119](#), [120](#), [121](#), [127](#), [128](#), [131](#), [132](#),
[133](#), [138](#), [140](#), [166](#), [169](#), [189](#), [193](#), [199](#), [202](#), [203](#),
[212](#), [215](#), [216](#), [226](#), [228](#), [229](#), [231](#), [233](#), [235](#), [238](#),
[240](#), [243](#), [252](#), [253](#), [265](#), [266](#), [283](#), [284](#), [285](#), [290](#),
[291](#), [393](#), [394](#), [400](#), [401](#), [405](#), [406](#), [409](#).
storage::get: [59](#), [63](#).
storage::resize: [77](#).
storage_type: [16](#), [73](#).
storg: [72](#), [73](#), [116](#), [205](#), [209](#), [218](#), [221](#), [222](#), [223](#).
STR: [44](#).
STRA: [44](#), [45](#).
strcmp: [180](#).
string: [4](#), [5](#), [8](#), [162](#), [181](#).
structure: [12](#), [14](#), [16](#), [20](#), [26](#), [35](#), [39](#), [43](#), [47](#), [50](#),
[105](#), [109](#), [119](#), [120](#), [121](#), [127](#), [128](#), [131](#), [132](#),
[133](#), [166](#), [169](#), [189](#), [193](#), [212](#), [238](#), [240](#), [243](#),
[252](#), [253](#), [265](#), [266](#), [401](#).
structure::preprocess: [59](#), [63](#).
structure::resize: [77](#).
structure_A: [215](#), [216](#).
structure_B: [215](#), [216](#).
structure_type: [16](#).

- subm*: [92](#), [93](#), 129, 130, 131, 132, 227, 228, 231, 235, 401, 430, 432.
- SUBM**: 81, 89.
- SUBMA**: 99, 102, 127, 128, 129, 130, 134.
- submatrix**: [81](#), [82](#), 83, [85](#), [86](#), 87, 88, 89, [90](#), 92, 105, 120, 228, 231, 235.
- submatrix::()**: 87, 89, 90.
- submatrix::cols**: [87](#), 89, 90, 122, 123, 124, 125, 127, 128, 129, 130, 131, 132.
- submatrix::get**: [91](#).
- submatrix::rows**: [87](#), 89, 90, 122, 123, 124, 125, 127, 128, 129, 130, 131, 132.
- submatrix::set**: [91](#).
- submatrix::subm**: [93](#).
- submatrix_template*: [81](#), 99, 102, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, 134.
- submatrix_type**: [83](#), 86, [88](#), 90, 92, 93, 98, 101, 104.
- SUBMX**: 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, 134.
- SUBMY**: 122, 123, 124, 125, 129, 130, 133, 134.
- sumt*: 435.
- sumt_iterator**: [435](#), 436, 438, 440.
- swap*: [106](#), 108.
- swaprows*: [108](#), 200, 202.
- symmetric**: [148](#), [149](#), [151](#), 218, 221, 249, 256, 262, 269, 280, 281, 284, 285, 288, 291, 295, 296, 304, 305, 312, 313, 321, 322, 330, 331, 339, 340, 349, 350, 359, 360, 368, 369, 394, 406, 426, 431.
- symmetric::preprocess**: [150](#).
- T**: [94](#), [96](#).
- t*: [375](#), [428](#), [430](#), [432](#).
- TA**: 44, 45.
- tdot*: [124](#), 129.
- theMatrix*: [66](#), 68, 69, 70, [84](#), 85, 86, 87, 90, 91, 93, 95, 97, 98, 100, 101, 103, 104, 107.
- theMessage*: [5](#), 7, 8.
- theRepresentation*: [15](#), 19, 23, 24, 35, 36, 37, 39, 43, 45, 46, 53, 57, 58, 62, 73, 74, 76, 89, 109.
- theStorage*: [20](#), 26, 47, 51, 54, 59, 63, 72, 77.
- theStructure*: [20](#), 26, 47, 51, 54, 59, 63, 77.
- tluser*: [313](#).
- tol*: [374](#), 375, 378.
- transpose*: [105](#), 233, 401, 430.
- tricky*: [161](#), [168](#), [188](#), [192](#).
- true*: 40, 60, 112, 150, 162, 408, 414, 421, 435, 441.
- tdot*: [125](#), 131.
- type*: [162](#), 165.
- unboundedbelow**: [9](#), 378.
- unstructured**: [11](#), 12, 26, [27](#), [49](#), 60, 81, 83, 116, 199, 202, 203, 205, 209, 222, 223, 226, 228, 229, 231, 233, 235, 243, 247, 248, 249, 252, 253, 254, 255, 256, 260, 261, 262, 267, 268, 269, 277, 278, 279, 281, 284, 285, 286, 287, 288, 291, 292, 293, 294, 295, 296, 301, 302, 303, 304, 305, 309, 310, 311, 312, 313, 318, 320, 321, 322, 327, 329, 330, 331, 336, 337, 338, 339, 340, 346, 348, 349, 350, 356, 358, 359, 360, 365, 367, 368, 369, 372, 376, 377, 378, 383, 384, 388, 391, 395, 400, 403, 406, 408, 409, 412, 413, 421, 423, 424, 426, 428, 429, 430, 431, 432, 435, 439.
- unstructured::preprocess**: [60](#).
- unstructured::resize**: [78](#).
- v*: [226](#), [231](#), [235](#).
- val*: [95](#), [97](#), [100](#), [103](#).
- value*: [34](#), [35](#), [36](#), [58](#), [59](#), [61](#), [70](#), [91](#), [94](#), 95, [96](#), 97, [98](#), [99](#), 100, [101](#), [102](#), 103, [104](#), 106, [107](#), [142](#), [275](#).
- vector**: 199, 202, 203, 205, 209, 243, 272, 273, 400, 435.
- void**: [435](#).
- v1*: [227](#).
- w*: [226](#), [229](#).
- weight*: [273](#).
- X*: [428](#), [430](#), [432](#).
- x*: [105](#), [106](#), [109](#), [111](#), [112](#), [116](#), [119](#), [120](#), [121](#), [122](#), [123](#), [124](#), [125](#), [126](#), [127](#), [128](#), [129](#), [130](#), [131](#), [132](#), [133](#), [134](#), [227](#), [247](#), [248](#), [249](#), [254](#), [255](#), [256](#), [260](#), [261](#), [262](#), [267](#), [268](#), [269](#), [277](#), [279](#), [281](#), [286](#), [287](#), [288](#), [293](#), [294](#), [295](#), [296](#), [301](#), [303](#), [305](#), [309](#), [311](#), [313](#), [318](#), [320](#), [322](#), [327](#), [329](#), [331](#), [336](#), [338](#), [340](#), [346](#), [348](#), [350](#), [356](#), [358](#), [360](#), [365](#), [367](#), [369](#), [378](#), [384](#), [391](#), [395](#), [403](#), [408](#), [412](#), [421](#), [428](#), [429](#), [430](#), [431](#), [435](#).
- xi*: [388](#).
- xrows*: [400](#), 401, 403, [407](#), 408, 409.
- xua*: [304](#), 305, [312](#), 313.
- xyyx*: [133](#), 305, 313.
- x0*: 372, [376](#), [377](#), [378](#), [383](#), [384](#), 385, [413](#), 414.
- Y*: [291](#).
- y*: [106](#), [109](#), [111](#), [112](#), [119](#), [120](#), [121](#), [122](#), [123](#), [124](#), [125](#), [129](#), [130](#), [133](#), [134](#), [286](#), [292](#).

- < Algebraic operations 119, 120, 121, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, 134 > Used in section 118.
- < Apply Cholesky transformation 214 > Used in section 212.
- < Apply Gauss transformation 201 > Used in section 199.
- < Apply Householder transformation 228 > Used in section 226.
- < Backtracking line search methods 382, 383, 384 > Used in section 381.
- < Barrier function internal variables 417, 420 > Used in section 416.
- < Barrier function methods 418, 421, 422, 423, 424 > Used in section 416.
- < Basic algebraic operations 105 > Used in section 3.
- < Basic definitions 4, 5, 40, 82 > Used in section 3.
- < Big definitions 112 > Used in section 2.
- < Bisection line search methods 375, 376, 377 > Used in section 374.
- < Build L^T row 213 > Used in section 212.
- < Build unconstrained objective function 436 > Used in section 435.
- < Check for lapack LU errors 207 > Used in sections 205 and 209.
- < Check lapack Cholesky errors 220 > Used in sections 218, 221, 222, and 223.
- < Cholesky lapack interface 218, 221, 222, 223 > Used in section 217.
- < Cholesky prototypes 212, 215, 216, 217 > Used in section 211.
- < Compute Hessian for log barrier 432 > Used in section 431.
- < Compute Hessian for **erf** functional 351 > Used in section 350.
- < Compute Householder vector for column j 227 > Used in section 226.
- < Compute next function value for **backtracking** 385 > Used in section 384.
- < Compute next function value for **bisection** 379 > Used in section 378.
- < Compute number of permutations and store it in *status* 208 > Used in sections 205 and 209.
- < Dense storage internal variables 28 > Used in section 10.
- < Dense storage methods 29, 30, 32, 48, 52, 55, 56, 61, 64, 75, 79 > Used in section 10.
- < Element definition 65 > Used in section 3.
- < Entropy of a functional class methods 326, 327, 329, 331 > Used in section 324.
- < Entropy of a functional internal variables 325, 330 > Used in section 324.
- < Equality search direction internal variables 398, 400 > Used in section 397.
- < Equality search direction methods 399, 401, 403 > Used in section 397.
- < Error function of a functional class methods 345, 346, 348, 350 > Used in section 343.
- < Error function of a functional internal variables 344, 349 > Used in section 343.
- < Exponential of a functional class methods 355, 356, 358, 360 > Used in section 353.
- < Exponential of a functional internal variables 354, 359 > Used in section 353.
- < Function base class methods 246, 247, 248, 249 > Used in section 245.
- < Functional base class methods 259, 260, 261, 262 > Used in section 258.
- < Functional minimization algorithm 413 > Used in section 412.
- < Gaxpy function class methods 253, 254, 255, 256 > Used in section 251.
- < Gaxpy function internal variables 252 > Used in section 251.
- < Gaxpy functional class methods 266, 267, 268, 269 > Used in section 264.
- < Gaxpy functional internal variables 265 > Used in section 264.
- < Generic error class methods 7, 8 > Used in section 5.
- < Get *number* from file, using the same tricky method as for **longs** 168 > Used in sections 167 and 170.
- < Gradient search direction methods 391 > Used in section 390.
- < Include files *fstream* 156, 163, 171 > Used in section 154.
- < Include files **math** 6, 31, 113, 114 > Used in section 3.
- < Include files **sparse** 139 > Used in section 136.
- < LU lapack interface 205, 209 > Used in section 204.
- < LU prototypes 199, 202, 203, 204 > Used in section 197.
- < Line search base class methods 372 > Used in section 371.
- < Linear combination of functionals class methods 273, 274, 275, 276, 277, 279, 281 > Used in section 271.
- < Linear combination of functionals internal variables 272, 278, 280 > Used in section 271.

- ⟨ Log barrier methods 427, 428, 429, 430, 431 ⟩ Used in section 426.
- ⟨ Logarithm of a functional class methods 364, 365, 367, 369 ⟩ Used in section 362.
- ⟨ Logarithm of a functional internal variables 363, 368 ⟩ Used in section 362.
- ⟨ Matrix definition 12 ⟩ Used in section 3.
- ⟨ Matrix element internal variables 66, 67 ⟩ Used in section 65.
- ⟨ Matrix element methods 68, 69, 70, 95, 97, 100, 103, 107 ⟩ Used in section 65.
- ⟨ Matrix internal types 14, 16, 88 ⟩ Used in section 12.
- ⟨ Matrix internal variables 15 ⟩ Used in section 12.
- ⟨ Matrix methods 19, 23, 24, 34, 36, 37, 38, 42, 44, 45, 46, 53, 57, 58, 62, 71, 73, 74, 76, 89, 92, 94, 96, 99, 102, 108, 115 ⟩
Used in section 12.
- ⟨ Matrix representation definition 13 ⟩ Used in section 12.
- ⟨ Matrix representation internal variables 17, 20, 21, 50 ⟩ Used in section 13.
- ⟨ Matrix representation methods 18, 22, 25, 26, 47, 51, 54, 59, 63, 72, 77 ⟩ Used in section 13.
- ⟨ Newton search direction internal variables 394 ⟩ Used in section 393.
- ⟨ Newton search direction methods 395 ⟩ Used in section 393.
- ⟨ Newton with equality functions 408, 409 ⟩ Used in section 405.
- ⟨ Newton with equality internal variables 406, 407 ⟩ Used in section 405.
- ⟨ Norm-2 error functional class methods 292, 293, 294, 295, 296 ⟩ Used in section 290.
- ⟨ Norm-2 error functional internal variables 291 ⟩ Used in section 290.
- ⟨ Power of a functional class methods 317, 318, 320, 322 ⟩ Used in section 315.
- ⟨ Power of a functional internal variables 316, 321 ⟩ Used in section 315.
- ⟨ Predefined error types 9 ⟩ Used in section 5.
- ⟨ Prepare for lapack Cholesky 219 ⟩ Used in sections 218, 221, 222, and 223.
- ⟨ Prepare for lapack LU 206 ⟩ Used in sections 205 and 209.
- ⟨ Product of functionals class methods 300, 301, 303, 305 ⟩ Used in section 298.
- ⟨ Product of functionals internal variables 299, 302, 304 ⟩ Used in section 298.
- ⟨ QR prototypes 226, 229 ⟩ Used in section 225.
- ⟨ Quadratic functional internal variables 284 ⟩ Used in section 283.
- ⟨ Quadratic functional methods 285, 286, 287, 288 ⟩ Used in section 283.
- ⟨ Ratio of functionals class methods 309, 311, 313 ⟩ Used in section 307.
- ⟨ Ratio of functionals internal variables 308, 310, 312 ⟩ Used in section 307.
- ⟨ Relative entropy functional class methods 335, 336, 338, 340 ⟩ Used in section 333.
- ⟨ Relative entropy functional internal variables 334, 337, 339 ⟩ Used in section 333.
- ⟨ SUMT Phase one 438, 439, 440 ⟩ Used in section 435.
- ⟨ SUMT Phase two 441 ⟩ Used in section 435.
- ⟨ SUMT functions 435 ⟩ Used in section 434.
- ⟨ Search direction base class methods 388 ⟩ Used in section 387.
- ⟨ Search for correct matrix name 180 ⟩ Used in section 179.
- ⟨ Search for pivot and swap rows if necessary 200 ⟩ Used in section 199.
- ⟨ Solve $Q^T x = y$ 235 ⟩ Used in section 233.
- ⟨ Solve $Qy = b$ 231 ⟩ Used in section 230.
- ⟨ Solve $R^T y = b$ 234 ⟩ Used in section 233.
- ⟨ Solve $Rx = y$ 232 ⟩ Used in section 230.
- ⟨ Solve least squares problem 230 ⟩ Used in section 229.
- ⟨ Solve minimum norm problem 233 ⟩ Used in section 229.
- ⟨ Sparse storage definition 137 ⟩ Used in section 136.
- ⟨ Sparse storage internal variables 138 ⟩ Used in section 137.
- ⟨ Sparse storage methods 140, 141, 142, 143, 144 ⟩ Used in section 137.
- ⟨ Specializations 106, 109, 110, 111, 116 ⟩ Used in section 3.
- ⟨ Storage definition 10 ⟩ Used in section 3.
- ⟨ Structure definition 11 ⟩ Used in section 3.
- ⟨ Submatrix definition 81 ⟩ Used in section 3.

<Submatrix internal variables 83, 84> Used in section 81.
 <Submatrix methods 85, 86, 87, 90, 91, 93, 98, 101, 104> Used in section 81.
 <Symmetric structure definition 148> Used in section 147.
 <Symmetric structure methods 149, 150, 151> Used in section 148.
 <Unstructured structure methods 27, 49, 60, 78> Used in section 11.
 <Update functional minimization stop criteria 414> Used in section 413.
 <Write matrix header 191> Used in sections 190 and 194.
 <Write *number* to file, using the same tricky method as for **longs** 192> Used in sections 190 and 194.
 <algebra.h 118>
 <barrier/log.h 426>
 <barrierbase.h 416>
 <cholesky.h 211>
 <det.h 243>
 <eye.h 238>
 <fmin.h 412>
 <fstream.h 154>
 <function/gaxpy.h 251>
 <functional/entr.h 324>
 <functional/erf.h 343>
 <functional/exp.h 353>
 <functional/gaxpy.h 264>
 <functional/linear.h 271>
 <functional/log.h 362>
 <functional/norm2err.h 290>
 <functional/power.h 315>
 <functional/prod.h 298>
 <functional/quadratic.h 283>
 <functional/ratio.h 307>
 <functional/releintr.h 333>
 <functionalbase.h 258>
 <functionbase.h 245>
 <linesearch/backtracking.h 381>
 <linesearch/bisection.h 374>
 <linesearchbase.h 371>
 <lu.h 197>
 <math.h 3>
 <ones.h 240>
 <qr.h 225>
 <searchdir/equality.h 397>
 <searchdir/equality/newton.h 405>
 <searchdir/gradient.h 390>
 <searchdir/newton.h 393>
 <searchdirbase.h 387>
 <sparse.h 136>
 <sumt.h 434>
 <symmetric.h 147>
 <**bisection** big definitions 378> Used in section 374.
 <**export**-waiting big definitions 33, 35, 39, 43, 80> Used in section 3.
 <fstream declarations 155, 181> Used in section 154.
 <fstream structures 162> Used in section 154.
 <**ifstream** methods 157, 158, 159, 160, 164, 166, 167, 169, 170, 172, 174, 176, 178> Used in section 155.
 <**ofstream** methods 182, 183, 184, 185, 187, 189, 190, 193, 194> Used in section 181.

⟨ **reletr** Hessian 341 ⟩ Used in section 340.

The MATH Library

(Version alpha 3-22-2002)

	Section	Page
Introduction	1	1
Basic definitions	2	2
Matrix basics	10	5
Creating and copying a matrix	21	8
Destroying a matrix	53	14
Setting and getting elements	57	15
Resizing	76	19
Submatrices	81	21
Basic algebraic operations	94	25
Basic specializations and utilities	106	28
Algebraic operations	117	30
The sparse storage	135	36
The symmetric structure	146	38
File streams: input	153	40
File streams: output	181	45
The LU decomposition	196	48
Solving a linear system	202	50
Interfacing with LAPACK	204	51
The Cholesky decomposition	210	53
Solving a linear system	215	54
Interfacing with LAPACK	217	55
The QR decomposition	224	57
Solving linear equations	229	59
Least squares	230	59
Minimum norm	233	60
Matrix creation functions	236	62
Eye	237	62
Ones	239	63
Matrix functions	241	64
Determinant	242	64
Functions	244	65
The gaxpy function	250	66

Functionals	257	68
The gaxpy functional	263	69
The linear combination of functionals	270	70
The quadratic functional	282	73
The norm-2 error	289	75
The product of two functionals	297	77
The ratio of two functionals	306	79
The power functional	314	81
The entropy of a functional	323	83
Relative entropy	332	85
The Error Function	342	88
The exponential of a functional	352	90
The logarithm of a functional	361	91
Line Searching	370	94
The bisection algorithm	373	94
Backtracking	380	97
Computing a search direction	386	100
The gradient direction	389	100
The Newton direction	392	101
Enforcing equality constraints	396	103
Newton direction with equality constraints	404	105
Optimization algorithms	410	109
Functional minimization	411	109
Barrier functions	415	111
The log barrier function	425	113
Sequential unconstrained minimization	433	117
Index	442	121

Copyright © 1998–2000 César A. R. Crusius

Permission is granted to make and distribute verbatim copies of this document provided that the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this document under the conditions for verbatim copying, provided that the entire resulting derived work is given a different name and distributed under the terms of a permission notice identical to this one.