



20	3.4	Conversions . . . . .	27
21	3.5	The class interface . . . . .	30
22	3.5.1	Inline functions . . . . .	30
23	3.5.2	Argument passing and return values . . . . .	30
24	3.5.3	const correctness . . . . .	35
25	3.5.4	Overloading and default arguments . . . . .	36
26	3.5.5	Comparisons . . . . .	37
27	3.6	new and delete . . . . .	38
28	3.7	Static and global objects . . . . .	40
29	3.8	Object-oriented programming . . . . .	41
30	3.9	Notes on the use of library functions. . . . .	44
31	3.10	Thread friendliness and thread safety . . . . .	46
32	3.11	Formatted output . . . . .	52
33	3.12	Assertions and error conditions . . . . .	53
34	3.13	Error handling . . . . .	53
35	3.14	Parts of C++ to avoid . . . . .	58
36	3.15	Readability and maintainability . . . . .	64
37	3.16	Portability . . . . .	66
38	<b>4</b>	<b>Style</b>	<b>70</b>
39	4.1	General aspects of style . . . . .	70
40	4.2	Comments . . . . .	73
41	<b>5</b>	<b>Changes</b>	<b>76</b>
42	5.1	Version 2.1 (Jan 1, 2026) . . . . .	76
43	5.2	Version 2.0 (March 6, 2024) . . . . .	76
44	5.3	Version 0.7 (Sep 18, 2019) . . . . .	77
45	5.4	Version 0.6 (Dec 20, 2017) . . . . .	77
46	5.5	Version 0.5 (Nov 21, 2017) . . . . .	77
47	5.6	Version 0.4 (Nov 16, 2017) . . . . .	77
48	5.7	Version 0.3 (Aug 23, 2017) . . . . .	77
49	5.8	Version 0.2 (Aug 9, 2017) . . . . .	78
50	<b>References</b>		<b>78</b>

51 

## 1 Introduction

52 This note gives a set of guidelines and recommendations for coding in C++ for  
53 the ATLAS experiment.

54 There are several reasons for maintaining and following a set of programming  
55 guidelines. First, by following some rules, one can avoid some common errors  
56 and pitfalls in C++ programming, and thus have more reliable code. But even  
57 more important: a computer program should not only tell the machine what to  
58 do, but it should also tell *other people* what you want the machine to do. (For  
59 much more elaboration on this idea, look up references on “literate programming,”  
60 such as [1].) This is obviously important any time when you have many people  
61 working on a given piece of software, and such considerations would naturally  
62 lead to code that is easy to read and understand. Think of writing ATLAS code as  
63 another form of publication, and take the same care as you would writing up an  
64 analysis for colleagues.

65 This document is derived from the original ATLAS C++ coding standard, [ATL-](#)  
66 [SOFT-2002-001](#) [2], which was last revised in 2003. This itself derived from work  
67 done by the CERN “Project support team” and SPIDER project, as documented  
68 in CERN-UCO/1999/207 [3]. These previous guidelines have been significantly  
69 revised to take into account the evolution of the C++ language [4], current  
70 practices in ATLAS, and the experience gained over the past decade.

71 Some additional useful information on C++ programming may be found in the  
72 references [5–13].

73 This note is not intended to be a fixed set of rigid rules. Rather, it should evolve  
74 as experience warrants.

75 

## 2 Naming

76 This section contains guidelines on how to name objects in a program.

77 

### 2.1 Naming of files

78 • **Each class should have one header file, ending with “.h”, and one**  
79 **implementation file, ending with “.cxx”. [source-naming]**

80 Some exceptions: Small classes used as helpers for another class should generally not go in their own file, but should instead be placed with the larger class. Sometimes several very closely related classes may be grouped together in a single file; in that case, the files should be named after whichever is the “primary” class. A number of related small helper classes (not associated with a particular larger class) may be grouped together in a single file, which should be given a descriptive name. An example of the latter could be a set of classes used as exceptions for a package.

88 For classes in a namespace, the namespace should not be included in the file name. For example, the header for `Trk::Track` should be called `Track.h`.

90 Implementation (“.cxx”) files that would be empty may be omitted.

91 The use of the “.h” suffix for headers is long-standing ATLAS practice; 92 however, it is unfortunate since language-sensitive editors may then default 93 to using “C” rather than “C++” mode for these files. For Emacs, it can help 94 to put a line like this at the start of the file:

```
1 // This file is really -*- C++ -*-.
```

## 95 2.2 Meaningful names

- 96 **Choose names based on pronounceable English words, common  
97 abbreviations, or acronyms widely used in the experiment, except  
98 for loop iteration variables.** [\[use-meaningful-names\]](#)

99 For example, `nameLength` is better than `nLn`.

100 Use names that are English and self-descriptive. Abbreviations and/or  
101 acronyms used should be of common use within the community.

- 102 **Do not create very similar names.** [\[no-similar-names\]](#)

103 In particular, avoid names that differ only in case. For example, `track` /  
104 `Track`; `c1` / `c1`; `XO` / `X0`.

## 105 2.3 Required naming conventions:

106 Generally speaking, you should try to match the conventions used by whatever  
107 package you’re working on. But please try to always follow these rules:

108 • **Use only ASCII characters in identifier names** [\[ascii-identifiers\]](#)

109 This is what C++ calls the basic character set. Specifically, identifiers should  
110 use only the characters a-z, A-Z, 0-9, and underscore.

111 Handling of non-ASCII characters is implementation-defined. While many  
112 compilers can indeed handle extended (unicode) characters, not all tools  
113 may process them correctly. Some characters may not display correctly,  
114 depending on a user's local installation. Further, it is often not obvious how  
115 to type an arbitrary unicode character that one sees displayed, especially  
116 since there exist distinct characters that look very similar or identical.

117 • **Use prefix `m_` for private/protected data members of classes.** [\[data-  
118 member-naming\]](#)

119 Use a lowercase letter after the prefix `m_`.

120 An exception for this is xAOD data classes, where the member names are  
121 exposed via ROOT for analysis.

122 • **Do not start any other names with `m_`.** [\[m-prefix-reserved\]](#)

123 • **Do not start names with an underscore. Do not use names that  
124 contain anywhere a double underscore.** [\[system-reserved-names\]](#)

125 Such names are reserved for the use of the compiler and system libraries.

126 The precise rule is that names that contain a double underscore or which  
127 start with an underscore followed by an uppercase letter are reserved  
128 anywhere, and all other names starting with an underscore are reserved in  
129 the global namespace. However, it's good practice to just avoid all names  
130 starting with an underscore. An exception is the use of a single underscore  
131 to indicate something that's structurally required but ignored.

132 **2.4 Recommended naming conventions**

133 If there is no already-established naming convention for the package you're work-  
134 ing on, the following guidelines are recommended as being generally consistent  
135 with ATLAS usage.

136 • **Use prefix `s_` for private/protected static data members of classes.** [\[static-members\]](#)

138 Use a lowercase letter after the prefix s\_.

139

- 140 • **The choice of namespace names should be agreed to by the communities concerned.** [\[namespace-naming\]](#)

141 Don't proliferate namespaces. If the community developing the code has a  
142 namespace defined already, use it rather than defining a new one. Examples  
143 include Trk:: for tracking and InDet:: for inner detector.

144

- 145 • **Use namespaces to avoid name conflicts between classes.** [\[use-namespaces\]](#)

146 A name clash occurs when a name is defined in more than one place. For  
147 example, two different class libraries could give two different classes the  
148 same name. If you try to use many class libraries at the same time, there  
149 is a fair chance that you will be unable to compile and link the program  
150 because of name clashes. To solve the problem you can use a namespace.

151 New code should preferably be put in a namespace, unless typical ATLAS  
152 usage is otherwise. For example, ATLAS classes related to the calorimeter  
153 conventionally start with "Calo" rather than being in a C++ namespace.

154

- 155 • **Start class and enumeration types with an uppercase letter.** [\[class-naming\]](#)

```
1 class Track;
2 enum State { green, yellow, red };
```

156

- 157 • **Type alias (`typedef`) names should start with an uppercase letter if they are public and treated as classes.** [\[typedef-naming\]](#)

```
1 using TrackVector =
2     std::vector<MCParticleKinematics*>;
```

158

- 159 • **Alternatively, a type alias (`typedef`) name may start with a lowercase letter and end with `_t`.** [\[typedef-naming-2\]](#)

160 This form should be reserved for type names which are not treated as classes  
161 (e.g., a name for a fundamental type) or names which are private to a class.

```
1 using mycounter_t = unsigned int;
```

162 • **Start names of variables, members, and functions with a lowercase**  
163 **letter.** [variable-and-function-naming]

```
1 double energy;  
2 void extrapolate();
```

164 Names starting with `s_` and `m_` should have a lowercase letter following  
165 the underscore.

166 Exceptions may be made for the case where the name is following standard  
167 physics or mathematical notation that would require an uppercase letter;  
168 for example, uppercase E for energy.

169 • **In names that consist of more than one word, write the words to-  
170 gether, and start each word that follows the first one with an upper-  
171 case letter.** [compound-names]

```
1 class OuterTrackerDigit;  
2 double depositedEnergy;  
3 void findTrack();
```

172 Some ATLAS packages also use the convention that names are entirely low-  
173 ercase and separated by underscores. When modifying an existing package,  
174 you should try to be consistent with the existing naming convention.

175 • **All package names in the release must be unique, independent of  
176 the package's location in the hierarchy.** [unique-package-names]

177 If there is a package, say “A/B/C”, already existing, another package may  
178 not have the name “D/E/C” because that “C” has already been used. This is  
179 required for proper functioning of the build system.

180 • **Underscores should be avoided in package names.** [no-underscores-  
181 in-package-names]

182 The old ATLAS rule was that a `_` should be used in package names when they  
183 are composites of one or more acronyms, e.g. `TRT_Tracker`, `AtlasDB_*`.  
184 Underscores should be avoided unless they really help with readability and  
185 help in avoiding spelling mistakes. `TRTTracker` looks odd because of the  
186 double “T”. Using underscores in package names will also add to confusion  
187 in the multiple-inclusion protection lines.

188 • Acronyms should be written as all uppercase. [uppercase-acronyms]

1 METReconstruction, not MetReconstruction  
2 MuonCSCValidation, not MuonCscValidation

189 Unfortunately, existing code widely uses both forms.

190 

## 3 Coding

191 This section contains a set of items regarding the “content” of the code. Orga-  
192 nization of the code, control flow, object life cycle, conversions, object-oriented  
193 programming, error handling, parts of C++ to avoid, portability, are all examples  
194 of issues that are covered here.

195 The purpose of the following items is to highlight some useful ways to exploit the  
196 features of the programming language, and to identify some common or potential  
197 errors to avoid.

198 

### 3.1 Organizing the code

199 • Header files must begin and end with multiple-inclusion protection.  
200 [header-guards]

1 `#ifndef PACKAGE_CLASS_H`  
2 `#define PACKAGE_CLASS_H`  
3 `// The text of the header goes in here ...`  
4 `#endif // PACKAGE_CLASS_H`

201 Header files are often included many times in a program. Because C++  
202 does not allow multiple definitions of a class, it is necessary to prevent the  
203 compiler from reading the definitions more than once.

204 The include guard should include both the package name and class name,  
205 to ensure that is unique.

206 Don’t start the include guard name with an underscore; such names are  
207 reserved to the compiler.

208 Be careful to use the same string in the `ifndef` and the `define`. It’s useful  
209 to get in the habit of using copy/paste here rather than retyping the string.

210 Some compilers support an extension `#pragma once` that has similar  
 211 functionality. A long time ago, this was sometimes faster, as it allowed  
 212 the compiler to skip reading headers that have already been read. How-  
 213 ever, modern compilers will automatically do this optimization based on  
 214 recognizing header guards. As `#pragma once` is nonstandard and has no  
 215 compelling advantage, it is best avoided.

216 In some rare cases, a file may be intended to be included multiple times, and  
 217 thus not have an include guard. Such files should be explicitly commented  
 218 as such, and should usually have an extension other than “.h”; “.def” is  
 219 sometimes used for this purpose.

220 • **Use forward declaration instead of including a header file, if this is  
 221 sufficient.** [\[forward-declarations\]](#)

```
1  class Line;
2  class Point
3  {
4  public:
5    // Distance from a line
6    Number distance(const Line& line) const;
7 }
```

222 Here it is sufficient to say that `Line` is a class, without giving details which  
 223 are inside its header. This saves time in compilation and avoids an apparent  
 224 dependency upon the `Line` header file.

225 Be careful, however: this does not work if `Line` is actually an alias (as is  
 226 the case, for example, with many of the xAOD classes).

227 • **Each header file must contain the declaration for one class only,  
 228 except for embedded or very tightly coupled classes or collections  
 229 of small helper classes.** [\[one-class-per-source\]](#)

230 This makes your source code files easier to read. This also improves the  
 231 version control of the files; for example the file containing a stable class  
 232 declaration can be committed and not changed any more.

233 Some exceptions: Small classes used as helpers for another class should gen-  
 234 erally not go in their own file, but should instead be placed with the larger  
 235 class. Sometimes several very closely related classes may be grouped to-

236      gether in a single file; in that case, the files should be named after whichever  
 237      is the “primary” class. A number of related small helper classes (not associated  
 238      with a particular larger class) may be grouped together in a single file,  
 239      which should be given a descriptive name. An example of the latter could  
 240      be a set of classes used as exceptions for a package.

- 241      **Implementation files must hold the member function definitions for the class(es) declared in the corresponding header file.**  
 242      [implementation-file]

244      This is for the same reason as for the previous item.

- 245      **Ordering of #include statements.** [include-ordering]

246      #include directives should generally be listed according to dependency  
 247      ordering, with the files that have the most dependencies coming first. This  
 248      implies that the first #include in a “.cxx” file should be the corresponding  
 249      “.h” file, followed by other #include directives from the same package.  
 250      These would then be followed by #include directives for other packages,  
 251      again ordered from most to least dependent. Finally, system #include  
 252      directives should come last.

```

1 // Example for CaloCell.cxx
2 // First the corresponding header.
3 #include "CaloEvent/CaloCell.h"
4 // The headers from other ATLAS packages,
5 // from most to least dependent.
6 #include "CaloDetDescr/CaloDetDescrElement.h"
7 #include "SGTools/BaseInfo.h"
8 // Headers from external packages.
9 #include "CLHEP/Geometry/Vector3D.h"
10 #include "CLHEP/Geometry/Point3D.h"
11 // System headers.
12 #include <cmath>

```

253      Ordering the #include directives in this way gives the best chance of  
 254      catching problems where headers fail to include other headers that they  
 255      depend on.

256      Some old guides recommended testing on the C++ header guard around the  
 257      #include directive. This advice is now obsolete and should be avoided.

```

1 // Obsolete --- don't do this anymore.
2 #ifndef MYPACKAGE_MYHEADER_H
3 # include "MyPackage/MyHeader.h"
4 #endif

```

258 The rationale for this was to avoid having the preprocessor do redundant  
 259 reads of the header file. However, current C++ compilers do this optimization  
 260 on their own, so this serves only to clutter the source.

261 • **Do not use “using” directives or declarations in headers or prior to**  
 262 **an #include.** [\[no-using-in-headers\]](#)

263 A using directive or declaration imports names from one namespace into  
 264 another, often the global namespace.

265 This does, however, lead to pollution of the global namespace. This can  
 266 be manageable if it's for a single source file; however, if the directive is in  
 267 a header file, it can affect many different source files. In most cases, the  
 268 author of these sources won't be expecting this.

269 Having using in a header can also hide errors. For example:

```

1 // In first header A.h:
2 using namespace std;
3
4 // In second header B.h:
5 #include "A.h"
6
7 // In source file B.cxx
8 #include "B.h"
9 ...
10 vector<int> x; // Missing std!

```

270 Here, a reference to std::vector in B.cxx is mistakenly written without  
 271 the std:: qualifier. However, it works anyway because of the using  
 272 directive in A.h. But imagine that later B.h is revised so that it no longer  
 273 uses anything from A.h, so the #include of A.h is removed. Suddenly,  
 274 the reference to vector in B.cxx no longer compiles. Now imagine there  
 275 are several more layers of #include and potentially hundreds of affected

276 source files. To try to prevent problems like this, headers should not use  
 277 using outside of classes. (Within a class definition, using can have a  
 278 different meaning that is not covered by this rule.)

279 For similar reasons, if you have a using directive or declaration in a “.cxx”  
 280 file, it should come after all #include directives. Otherwise, the using  
 281 may serve to hide problems with missing namespace qualifications in the  
 282 headers.

283 This rule does not apply when using is used to define a type alias (similarly  
 284 to typedef).

## 285 3.2 Control flow

- 286 • **Do not change a loop variable inside a for loop block.** [do-not-modify-  
 287 for-variable]

288 When you write a for loop, it is highly confusing and error-prone to change  
 289 the loop variable within the loop body rather than inside the expression exe-  
 290 cuted after each iteration. It may also inhibit many of the loop optimizations  
 291 that the compiler can perform.

- 292 • **Prefer range-based for loops.** [prefer-range-based-for]

293 Prefer a range-based for to a loop with explicit iterators. That is, prefer:

```
1 std::vector<int> v = ...;
2 for (int x : v) {
3     doSomething (x);
4 }
```

294 to

```
1 std::vector<int> v = ...;
2 for (std::vector<int>::const_iterator it = v.begin();
3      it != v.end();
4      ++it)
5 {
6     doSomething (*it);
7 }
```

295        In some cases you can't make this replacement; for example, if you need to  
 296        call methods on the iterator itself, or you need to manage multiple iterators  
 297        within the loop. But most simple loops over STL ranges are more simply  
 298        written with a range-based for.

299        As of C++20, you can initialize additional variables in a range-based for:

```

1  void foo (const std::vector<float>& v) {
2      for (int i = 0; float f : v) {
3          std::cout << i++ << " " << f << "\n";
4      }
5  }
```

300        • **Switch statements should have a `default` clause.** [\[switch-default\]](#)

301        A switch statement should have a default clause, rather than just falling  
 302        off the bottom, as a cue to the reader that this case was expected.

303        In some cases, a switch statement may be on a enum and includes case  
 304        clauses for all possible values of the enum. In such cases, a default cause  
 305        is not required. Recent compilers will generate warnings if some elements  
 306        of an enum are not handled in a switch. This mitigates the risk that a  
 307        switch does not get updated after a new enum value is added.

308        • **Each clause of a `switch` statement must end with `break`.** [\[switch-break\]](#)

310        If you must “fall through” from one switch clause to another (excluding  
 311        the trivial case of a clause with no statements), this should be explicitly  
 312        indicated using the `fallthrough` attribute. This should, however, be a  
 313        rare case.

```

1  switch (case) {
2      case 1:
3          doSomething();
4          [[fallthrough]];
5      case 2:
6          doSomethingMore();
7          break;
8      ...
```

314 Recent compilers will warn about such constructs unless you use the attribute or a special comment. For new code, using the attribute is preferred.  
315

- 316 • **An if-statement which does not fit in one line must have braces  
317 around the conditional statement.** [\[if-bracing\]](#)

318 This makes code much more readable and reliable, by clearly showing the  
319 flow paths.

320 The addition of a final else is particularly important in the case where  
321 you have if/else-if. To be safe, even single statements should be explicitly  
322 blocked by {}.

```
1  if (val == thresholdMin) {  
2      statement;  
3  }  
4  else if (val == thresholdMax) {  
5      statement;  
6  }  
7  else {  
8      statement; // handles all other (unforeseen) cases  
9  }
```

- 323 • **Do not use goto.** [\[no/goto\]](#)

324 Use break or continue instead.

325 This statement remains valid also in the case of nested loops, where the use  
326 of control variables can easily allow to break the loop, without using goto.

327 goto statements decrease readability and maintainability and make testing  
328 difficult by increasing the complexity of the code.

329 If goto statements must be used, it's better to use them for forward branching  
330 than backwards, and the functions involved should be kept short.

### 331 3.3 Object life cycle

#### 332 3.3.1 Initialization of variables and constants

- 333 • **Declare each variable with the smallest possible scope and initialize  
334 it at the same time.** [\[variable-initialization\]](#)

335 It is best to declare variables close to where they are used. Otherwise you  
 336 may have trouble finding out the type of a particular variable.

337 It is also very important to initialize the variable immediately, so that its  
 338 value is well defined.

```
1 int value = -1;    // initial value clearly defined
2 int maxValue;      // initial value undefined ...
3                      // NOT recommended
```

339 • **Avoid use of “magic literals” in the code. [no-magic-literals]**

340 If some number or string has a particular meaning, it's best to declare a  
 341 symbol for it, rather than using it directly. This is especially true if the same  
 342 value must be used consistently in multiple places.

343 Bad example:

```
1 class A
2 {
3     ...
4     TH1* m_array[10];
5 }
6
7 void A::foo()
8 {
9     for (int i = 0; i < 10; i++) {
10         m_array[i] = dynamic_cast<TH1*>
11             (gDirectory()->Get (TString ("hist_") +
12             TString::Itoa(i, 10)));
13 }
```

344 Better example:

```
1 class A
2 {
3     ...
4
5     static const s_numberOfHistograms = 10;
6     static TString s_histPrefix;
```

```

7   TH1* m_array[s_numberOfHistograms];
8 }
9
10 TString s_histPrefix = "hist_";
11
12 void A::foo()
13 {
14   for (int i = 0; i < s_numberOfHistograms; i++) {
15     TString istr = TString::Itoa (i, 10); // base 10
16     m_array[i] = dynamic_cast<TH1*>
17       (gDirectory()->Get (s_histPrefix + istr));
18 }

```

345 It is not necessary to turn *every* literal into a symbol. For example, the  
 346 ‘10’ in the example above in the `Itoa` call, which gives the base for the  
 347 conversion, would probably not benefit from being made a symbol, though  
 348 a comment might be helpful. Similarly, sometimes `reserve()` is called on  
 349 a `std::vector` before it is filled with a value that is essentially arbitrary.  
 350 It probably also doesn’t help to make this a symbol, but again, a comment  
 351 would be helpful. Strings containing text to be written as part of a log  
 352 message are also best written literally.

353 In general, though, if you write a literal value other than ‘0’, ‘1’, `true`,  
 354 `false`, or a string used in a log message, you should consider defining a  
 355 symbol for it.

- 356 • **Use the `<numbers>` header for mathematical constants.** [\[math-  
 357 constants\]](#)

358 Basic mathematical constants are available in the header `<numbers>`. Use  
 359 these in preference to the `M_` constants from `math.h` or explicit definitions:

```

1 #include <numbers>
2 #include <cmath>
3 double f (double x) {
4   return std::sin (x * std::numbers::pi);
5 }

```

- 360 • **Declare each type of variable in a separate declaration statement, and**

361       **do not declare different types (e.g. `int` and `int*`) in one declaration**  
 362       **statement.** [\[separate-declarations\]](#)

363       Declaring multiple variables on the same line is not recommended. The  
 364       code will be difficult to read and understand.

365       Some common mistakes are also avoided. Remember that when you declare  
 366       a pointer, a unary pointer is bound only to the variable that immediately  
 367       follows.

```
1  int i, *ip, ia[100], (*ifp)();    // Not recommended
2
3  // recommended way:
4  LoadModule* oldLm = 0; // pointer to the old object
5  LoadModule* newLm = 0; // pointer to the new object
```

368       Bad example: both `ip` and `jp` were intended to be pointers to integers, but  
 369       only `ip` is – `jp` is just an integer!

```
1  int* ip, jp;
```

370       • **Do not use the same variable name in outer and inner scope.** [\[no-variable-shadowing\]](#)

372       Otherwise the code would be very hard to understand; and it would certainly  
 373       be very error prone.

374       Some compilers will warn about this.

375       • **Be conservative in using `auto`.** [\[using-auto\]](#)

376       The `auto` keyword allows one to omit explicitly writing types that the  
 377       compile can deduce. Examples:

```
1  auto x = 10; // Type int deduced
2  auto y = 42ul; // Type unsigned long deduced.
3  auto it = vec.begin(); // Iterator type deduced
```

378       Some authorities have recommended using `auto` pretty much everywhere  
 379       you can (calling it “auto almost always”). However, our experience has  
 380       been that this adversely affects the readability and robustness of the code.  
 381       It generally helps a reader to understand what the code is doing if the type

382 is apparent, but with `auto`, it often isn't. Using `auto` also makes it more  
 383 difficult to find places where a particular type is used when searching the  
 384 code with tools like LXR. It can also make it more difficult to track errors  
 385 back to their source:

```

1 const Foo* doSomething();
2 ... a lot of code here ...
3 auto foo = doSomething();
4 // What is the type of foo here? You have to look up
5 // doSomething() in order to find out! Makes it much
6 // harder to find all places where the type Foo
7 // gets used.
8
9 // If the return type of doSomething() changes, you'll
10 // get an error here, not at the doSomething() call.
11 foo->doSomethingElse();

```

386 `auto` has also been observed to be a frequent source of errors leading to  
 387 unwanted copies of objects. For example, in this code:

```

1 std::vector<std::vector<int> > arr = ...;
2 for (auto v : arr) {
3     for (auto elt : v) { ...

```

388 each element of the outermost vector will be copied, as the assignment to  
 389 `v` will be done by value. One would probably want:

```

1 std::vector<std::vector<int> > arr = ...;
2 for (const auto& v : arr) {
3     for (auto elt : v) { ...

```

390 but having to be aware of the type like this kind of obviates the motivation  
 391 for using `auto` in the first place. Using the type explicitly makes this sort  
 392 of error much more difficult.

393 The current recommendation is to generally not use `auto` in place of a  
 394 (possibly-qualified) simple type:

```

1 // Use these
2 int x = 42;

```

```

3  const Foo* foo = doSomething();
4  for (const CaloCell* cell : caloCellContainer) ...
5  Foo foo (x);
6
7  // Rather than these
8  auto x = 42;
9  auto foo = doSomething();
10 for (auto cell : caloCellContainer) ...
11 auto foo = Foo {x};

```

395 There are a few sorts of places where it generally makes sense to use `auto`.

396 – When the type is already evident in the expression and the declaration  
 397 would be redundant. This is usually the case for expressions with `new`  
 398 or `make_unique`.

```

1  // auto is fine here.
2  auto foo = new Foo;
3  auto ufoo = std::make_unique<Foo>();

```

399 – When you need a declaration for a complicated derived type, where  
 400 the type itself isn't of much interest.

```

1  // Fine to use auto here; the full name of the
2  // type is too cumbersome to be useful.
3  std::map<int, std::string> m = ...;
4  auto ret = m.insert (std::make_pair (1, "x"));
5  if (ret.second) ....

```

401 – In the case where a class method returns a type defined within the  
 402 class, using the `auto` syntax to write the return type at the end of the  
 403 signature can make things more readable when the method is defined  
 404 out-of-line:

```

1  template <class T> class C {
2  public:
3      using ret_t = int;
4      ret_t foo();
5  };

```

```

6   // Verbose: the return type is interpreted at the
7   // global scope, so it needs to be qualified with
8   // the class name.
9   template <class T>
10  typename C<T>::ret_t C<T>::foo() ...
11
12
13  // With this syntax, the return type is
14  // interpreted within the class scope.
15  template <class T>
16  auto C<T>::foo() -> ret_t ...

```

405 – `auto` may also be useful in writing generic template code.

406 In some cases, C++20 allows declaring a template function without the  
407 `template` keyword when the argument is declared as `auto`:

```
1 auto fn (auto x) { return x + 1; }
```

408 It is recommended to avoid this syntax for public interfaces.

409 In general, the decision as to whether or not to use `auto` should be made  
410 on the basis of what makes the code easier to *read*. It is bad practice to use  
411 it simply to save a few characters of typing.

412 **3.3.2 Constructor initializer lists**

413 • **Let the order in the initializer list be the same as the order of the**  
414 **declarations in the header file: first base classes, then data members.**  
415 [\[member-initializer-ordering\]](#)

416 It is legal in C++ to list initializers in any order you wish, but you should  
417 list them in the same order as they will be called.

418 The order in the initializer list is irrelevant to the execution order of the  
419 initializers. Putting initializers for data members and base classes in any order  
420 other than their actual initialization order is therefore highly confusing  
421 and can lead to errors.

422 Class members are initialized in the order of their declaration in the class;  
423 the order in which they are listed in a member initialization list makes no

424 difference whatsoever! So if you hope to understand what is really going on  
 425 when your objects are being initialized, list the members in the initialization  
 426 list in the order in which those members are declared in the class.

427 Here, in the bad example, `m_data` is initialized first (as it appears in the  
 428 class) *before* `m_size`, even though `m_size` is listed first. Thus, the `m_data`  
 429 initialization will read uninitialized data from `m_size`.

430 Bad example:

```

1 class Array
2 {
3 public:
4     Array(int lower, int upper);
5 private:
6     int* m_data;
7     unsigned m_size;
8     int m_lowerBound;
9     int m_upperBound;
10    } ;
11    Array::Array(int lower, int upper) :
12        m_size(upper-lower+1),
13        m_lowerBound(lower),
14        m_upperBound(upper),
15        m_data(new int[m_size])
16    { . . .

```

431 Correct example:

```

1 class Array
2 {
3 public:
4     Array(int lower, int upper);
5 private:
6     unsigned m_size;
7     int m_lowerBound;
8     int m_upperBound;
9     int* m_data;
10    } ;

```

```

11  Array::Array(int lower, int upper) :
12    m_size(upper-lower+1),
13    m_lowerBound(lower),
14    m_upperBound(upper),
15    m_data(new int[m_size]) { ...

```

432 Virtual base classes are always initialized first, then base classes, data  
 433 members, and finally the constructor body for the derived class is run.

```

1  class Derived : public Base // Base is number 1
2  {
3    public:
4      explicit Derived(int i);
5      // The keyword explicit prevents the constructor
6      // from being called implicitly.
7      // int x = 1;
8      // Derived dNew = x;
9      // will not work
10
11    Derived();
12
13    private:
14      int m_jM; // m_jM is number 2
15      Base m_bM; // m_bM is number 3
16    } ;
17
18    Derived::Derived(int i) : Base(i), m_jM(i), m_bM(i) {
19      // Recommended order           1           2           3
20      ...
21    }

```

434 **3.3.3 Copying of objects**

435 • **A function must never return, or in any other way give access to,**  
 436 **references or pointers to local variables outside the scope in which**  
 437 **they are declared.** [no-refs-to-locals]

438 Returning a pointer or reference to a local variable is always wrong because

439 it gives the user a pointer or reference to an object that no longer exists.

440 Bad example:

441 You are using a complex number class, Complex, and you write a method  
442 that looks like this:

```

1 Complex&
2 calculateC1 (const Complex& n1, const Complex& n2)
3 {
4     double a = n1.getReal() - 2*n2.getReal();
5     double b = n1.getImaginary() * n2.getImaginary();
6
7     // Create local object.
8     Complex C1(a, b);
9
10    // Return reference to local object.
11    // The object is destroyed on exit from this
12    // function: trouble ahead!
13    return C1;
14 }
```

443 In fact, most compilers will spot this and issue a warning.

444 This particular function would be better written to return the result by  
445 value:

```

1 Complex calculateC1 (const Complex& n1,
2                         const Complex& n2)
3 {
4     double a = n1.getReal() - 2*n2.getReal();
5     double b = n1.getImaginary() * n2.getImaginary();
6
7     return Complex(a, b);
8 }
```

446 • If objects of a class should never be copied, then the copy con-  
447 structor and the copy assignment operator should be deleted. [copy-  
448 protection]

449 Ideally the question whether the class has a reasonable copy semantic will

450 naturally be a result of the design process. Do not define a copy method for  
 451 a class that should not have it.

452 By deleting the copy constructor and copy assignment operator, you can  
 453 make a class non-copyable.

```

1  // There is only one ATLASExperimentalHall,
2  // and that should not be copied
3  class ATLASExperimentalHall
4  {
5    public:
6      ATLASExperimentalHall();
7      ~ATLASExperimentalHall();
8
9      // Delete copy constructor to disallow copying.
10     ATLASExperimentalHall(const ATLASExperimentalHall& )
11     = delete;
12
13     // Delete assignment operator to disallow assignment.
14     ATLASExperimentalHall&
15     operator=(const ATLAS.ExperimentalHall&) = delete;
16   } ;

```

454 In older versions of the language, this was achieved by declaring the deleted  
 455 methods as **private** (and not implementing them). For new code, prefer  
 456 explicitly deleting the functions.

```

1  // There is only one ATLASExperimentalHall,
2  // and that should not be copied
3  class ATLASExperimentalHall
4  {
5    public:
6      ATLASExperimentalHall();
7      ~ATLASExperimentalHall();
8
9    private:
10      // Disallow copy constructor and assignment.
11      ATLASExperimentalHall(const ATLASExperimentalHall& );
12      ATLASExperimentalHall& operator=

```

```

13     (const ATLAS_ExperimentalHall&);
14 }
```

- **If a class owns memory via a pointer data member, then the copy constructor, the assignment operator, and the destructor should all be implemented.** [\[define-copy-and-assignment\]](#)

The compiler will generate a copy constructor, an assignment operator, and a destructor if these member functions have not been declared. A compiler-generated copy constructor does memberwise initialization and a compiler-generated copy assignment operator does memberwise assignment of data members and base classes. For classes that manage resources (examples: memory (new), files, sockets) the generated member functions probably have the wrong behavior and must be implemented by the developer. You have to decide if the resources pointed to must be copied as well (deep copy), and implement the correct behavior in the operators. Of course, the constructor and destructor must be implemented as well.

Bad Example:

```

1  class String
2  {
3  public:
4      String(const char *value=0);
5      ~String(); // Destructor but no copy constructor
                // or assignment operator.
6
7  private:
8      char *m_data;
9 }
10
11 String::String(const char *value)
12 { // Correct behavior implemented in constructor.
13     m_data = new char[strlen(value)]; // Fill m_data
14 }
15 String::~String()
16 { // Correct behavior implemented in destructor
17     delete m_data;
18 }
```

```

20  . . .
21
22
23 // Declare and construct a. m_data points to "Hello"
24 String a("Hello");
25
26 // Open new scope
27 { // Declare and construct b.
28   // m_data points to "World"
29   String b("World");
30
31   b=a;
32   // Execute default op= as synthesized by the compiler.
33   // This is simply memberwise assignment.
34   // For pointers like m_data, this is a bitwise copy
35   // ==> m_data of "a" and "b" now point to the
36   // same string "Hello"
37   // ==> 1) Memory b used to point to never deleted:
38   //           a possible memory leak!
39   // 2) When either a or b goes out of scope,
40   //           its destructor will delete the memory
41   //           still pointed to by the other.
42 }
43
44 // Close scope: b's destructor called;
45 // memory still pointed to by `a' deleted!
46 String c=a;
47 // But m_data of a is undefined!!

```

- 471 • **Assignment member functions must work correctly when the left and right operands are the same object. [self-assign]**

472 This requires some care when writing assignment code, as this case (when left and right operands are the same) may require that most of the code is bypassed.

```

1 A& A::operator=(const A& a)
2 {

```

```

3  if (this != &a) {
4      // ... implementation of operator=
5  }
6 }
```

## 476 3.4 Conversions

477 • **Use explicit rather than implicit type conversion.** [avoid-implicit-  
478 conversions]

479 Most conversions are bad in some way. They can make the code less  
480 portable, less robust, and less readable. It is therefore important to use only  
481 explicit conversions. Implicit conversions are almost always bad.

482 • **Use the C++ cast operators (`dynamic_cast` and `static_cast`)**  
483 **instead of the C-style casts.** [use-c++-casts]

484 In general, casts should be strongly discouraged, especially the old style C  
485 casts.

486 The new cast operators give the user a way to distinguish between different  
487 types of casts, and to ensure that casts only do what is intended and nothing  
488 else.

489 The C++ `static_cast` operator allows explicitly requesting allowed im-  
490 plicit conversions and between integers and enumerations. It also allows  
491 casting pointers up and down a class hierarchy (as long as there's no vir-  
492 tual inheritance), but no checking is done when casting from a less- to a  
493 more-derived type.

494 The C++ `dynamic_cast` operator is used to perform safe casts down or  
495 across an inheritance hierarchy. One can actually determine whether the  
496 cast succeeded because failed casts are indicated either by a `bad_cast`  
497 exception or a null pointer. The use of this type of information at run time  
498 is called Run-Time Type Identification (RTTI).

```

1  int n = 3;
2  double r = static_cast<double>(n) * a;
3
4  class Base { };
```

```

5  class Derived : Base { };
6  void f(Derived* d_ptr)
7  {
8      // if the following cast is inappropriate
9      // a null pointer will be returned!
10     Base* b_ptr = dynamic_cast<Base*>(d_ptr);
11     // ...
12 }
```

499 • **Do not convert `const` objects to non-`const`.** [no-const-cast]

500 In general you should never cast away the constness of objects.

501 If you have to use a `const_cast` to remove `const`, either you're writing  
 502 some low-level code that that's deliberately subverting the C++ type system,  
 503 or you have some problem in your design or implementation that the  
 504 `const_cast` is papering over.

505 Sometimes you're forced to use a `const_cast` due to problems with external  
 506 libraries. But if the library in question is maintained by ATLAS, then  
 507 try to get it fixed in the original library before resorting to `const_cast`.

508 The keyword `mutable` allows data members of an object that have been  
 509 declared `const` to remain modifiable, thus reducing the need to cast away  
 510 constness. The `mutable` keyword should only be used for variables which  
 511 are used for caching information. In other words, the object appears not to  
 512 have changed but it has stored something to save time on subsequent use.

513 • **Do not use `reinterpret_cast`.** [no-reinterpret-cast]

514 `reinterpret_cast` is machine-, compiler- and compile-options-  
 515 dependent. It is a way of forcing a compiler to accept a type conversion  
 516 which is dependent on implementation. It blows away type-safety, violates  
 517 encapsulation and more importantly, can lead to unpredictable results.

518 `reinterpret_cast` has legitimate uses, such as low-level code which  
 519 deliberately goes around the C++ type system. Such code should usually  
 520 be found only in the core and framework packages.

521 Exception: `reinterpret_cast` is required in some cases if one is not  
 522 using old-style casts. It is required for example if you wish to convert a

523       callback function signature (X11, expat, Unix signal handlers are common  
 524       causes). Some external libraries (X11 in particular) depend on casting  
 525       function pointers. If you absolutely have to work around limitations in  
 526       external libraries, you may of course use it.

527       One particularly nasty case to be aware of and to avoid is *pointer aliasing*.  
 528       If two pointers have different types, the compiler may assume that they  
 529       cannot point at the same object. For example, in this function:

```
1 int convertAndBuffer (int* buf, float x)
2 {
3     float* fbuf = reinterpret_cast<float*>(buf);
4     *fbuf = x;
5     return *buf;
6 }
```

530       the compiler is entitled to rewrite it as

```
1 int convertAndBuffer (int* buf, float x)
2 {
3     int ret = *buf;
4     float* fbuf = reinterpret_cast<float*>(buf);
5     *fbuf = x;
6     return ret;
7 }
```

531       (As a special case, you can safely convert any pointer type to or from a  
 532       char\*.) The proper way to do such a conversion is with a std::bit\_cast:

```
1 #include <bit>
2 int convertAndBuffer (int* buf, float x)
3 {
4     *buf = std::bit_cast<int> (x);
5     return *buf;
6 }
```

533       Prior to C++20, the recommended way to do this was with a union, but that  
 534       should not be used for new code.

535 **3.5 The class interface**536 **3.5.1 Inline functions**

537 • **Header files must contain no implementation except for small func-**  
538 **tions to be inlined. These inlined functions must appear at the end**  
539 **of the header after the class definition.** [\[inline-functions-impls\]](#)

540 If you have many inline functions, it is usually better to split them out into  
541 a separate file, with extension “.icc”, that is included at the end of the  
542 header.

543 Inline functions can improve the performance of your program; but they  
544 also can increase the overall size of the program and thus, in some cases,  
545 have the opposite result. It can be hard to know exactly when inlining is  
546 appropriate. As a rule of thumb, inline only very simple functions to start  
547 with (one or two lines). You can use profiling information to identify other  
548 functions that would benefit from inlining.

549 Use of inlining makes debugging hard and, even worse, can force a complete  
550 release rebuild or large scale recompilation if the inline definition needs to  
551 be changed.

552 **3.5.2 Argument passing and return values**

553 • **Pass an unmodifiable argument by value only if it is of built-in type**  
554 **or small; otherwise, pass the argument by `const` reference (or by**  
555 **`const` pointer if it may be null).** [\[large-argument-passing\]](#)

556 An object is considered small if it is a built-in type or if it contains at most  
557 one small object. Built-in types such as `char`, `int`, and `double` can be  
558 passed by value because it is cheap to copy such variables. If an object is  
559 larger than the size of its reference (typically 64 bits), it is not efficient to  
560 pass it by value. Of course, a built-in type can be passed by reference when  
561 appropriate.

```
1 void func(char c);    // OK
2 void func(int i);    // OK
3 void func(double d); // OK
4 void func(complex<float> c); // OK
5
```

```

6  void func(Track t); // not good, since Track is large,
7          // so there is an overhead in
8          // copying t

```

562 Arguments of class type are often costly to copy, so it is preferable to pass  
 563 a `const` reference to such objects; in this way the argument is not copied.  
 564 Const access guarantees that the function will not change the argument.

565 In terms of efficiency, passing by pointer is the same as passing by reference.  
 566 However, passing by reference is preferred, unless it is possible to the object  
 567 to be missing from the call.

```

1  void func(const LongString& s); // const reference

```

568

- 569 • **If an argument may be modified, pass it by `non-const` reference and clearly document the intent.** [\[modifiable-arguments\]](#)

570 For example:

```

1  // Track @c t is updated by the addition of hit @c h.
2  void updateTrack(const Hit& h, Track& t);

```

571 Again, passing by references is preferred, but a pointer may be used if the  
 572 object can be null.

573

- 574 • **Use `unique_ptr` to pass ownership of an object to a function.** [\[pass-ownership\]](#)

575 To pass ownership of an object into a function, use `unique_ptr` (by value):

```

1  void foo (std::unique_ptr<Object> obj);
2
3  ...
4
5  auto obj = std::make_unique<Object>();
6  ...
7  foo (std::move (obj));

```

576 In most cases, `unique_ptr` should be passed by *value*. There are however  
 577 a few possible use cases for passing `unique_ptr` by reference:

578           – The called function may replace the object passed in with another one.  
 579            In this case, however, consider returning the new object as the value  
 580            of the function.

581           – The called function may only conditionally take ownership of the  
 582            passed object. This is likely to be confusing and error-prone and  
 583            should probably be avoided. Consider if a `shared_ptr` would be  
 584            better in this case.

585       There is basically no good case for passing `unique_ptr` as a `const` refer-  
 586       ence.

587       If you need to interoperate with existing code, object ownership may be  
 588       passed by pointer. The fact that ownership is transferred should be clearly  
 589       documented.

590       *Do not* pass ownership using references.

591       Here are a some additional examples to illustrate this. Assume that class C  
 592       contains a member `Foo* m_owning_pointer` which the class deletes.  
 593       (In modern C++, it would of course usually be better for this to be a  
 594       `unique_ptr`.)

```

1  // --- Best
2  void C::takesOwnership (std::unique_ptr<Foo> foo)
3  {
4      delete m_owning_pointer;
5      m_owning_pointer = foo.release();
6  }
7
8  // --- OK if documented.
9  // Takes ownership of the @c foo pointer.
10 void C::takesOwnership (Foo* foo)
11 {
12     delete m_owning_pointer;
13     m_owning_pointer = foo;
14 }
15
16 // --- Don't do this!
17 void C::takesOwnership (Foo& foo)

```

```

18  {
19      delete m_owning_pointer;
20      m_owning_pointer = &foo;
21  }

```

595 • **Return basic types or new instances of a class type by value.** [return-by-value]

597 Returning a class instance by value is generally preferred to passing an  
598 argument by non-const reference:

```

1  // Bad
2  void getVector (std::vector<int>& v)
3  {
4      v.clear();
5      for (int i=0; i < 10; i++) v.push_back(v);
6  }
7
8  // Better
9  std::vector<int> getVector()
10 {
11     std::vector<int> v;
12     for (int i=0; i < 10; i++) v.push_back(v);
13     return v;
14 }

```

599 The return-value optimization plus move semantics will generally mean  
600 that there won't be a significant efficiency difference between the two.

601 • **Use unique\_ptr to return ownership.** [returning-ownership]

602 If a function is returning a pointer to something that is allocated off the heap  
603 which the caller is responsible for deleting, then return a unique\_ptr.

604 If compatibility with existing code is an issue, then a plain pointer may be  
605 used, but the caller takes ownership should be clearly documented.

606 *Do not* return ownership via a reference.

```

1  // Best
2  std::unique_ptr<Foo> makeFoo()

```

```

3  {
4      return std::make_unique<Foo> ( . . . );
5  }
6
7  // OK if documented
8  // makeFoo() returns a newly-allocated Foo;
9  // caller must delete it.
10 Foo* makeFoo()
11 {
12     return new Foo ( . . . );
13 }
14
15 // NO!
16 Foo& makeFoo()
17 {
18     Foo* foo = new Foo ( . . . );
19     return *foo;
20 }
```

- 607 • Have **operator=** return a reference to **\*this**. [assignment-return-  
608 value]

609 This ensures that

```
1  a = b = c;
```

610 will assign c to b and then b to a as is the case with built-in objects.

- 611 • Use **std::span** to represent and pass a bounded region of memory.  
612 [span]

613 In particular, use **std::span** instead of passing a pointer with a sepa-  
614 rate element count (or even worse, a pointer to an array with no bounds  
615 information).

616 So you can use this:

```
1  #include <span>
2  int sum (const std::span<int>& s)
3  {
```

```

4     int ret = 0;
5     for (int i : s) ret += i;
6     return ret;
7 }
```

617 instead of

```

1 int sum (const int* p, size_t n)
2 {
3     int ret = 0;
4     for (size_t i = 0; i < n; i++) ret += p[i];
5     return ret;
6 }
```

618 One might expect that `std::span` would include an `at()` method, to  
 619 allow indexing with bounds checking, but that is only available in C++23.  
 620 In the meantime, `CxxUtils::span` is very similar to `std::span` but does  
 621 implement `at()`.

622 **3.5.3 const correctness**

- 623 • **Declare a pointer or reference argument, passed to a function, as**  
 624 **const if the function does not change the object bound to it.** [const-  
 625 arguments]

626 An advantage of `const`-declared parameters is that the compiler will ac-  
 627 tually give you an error if you modify such a parameter by mistake, thus  
 628 helping you to avoid bugs in the implementation.

```

1 // operator<< does not modify the String parameter
2 ostream& operator<<(ostream& out, const String& s);
```

- 629 • **The argument to a copy constructor and to an assignment operator**  
 630 **must be a const reference.** [copy-ctor-arg]

631 This ensures that the object being copied is not altered by the copy or  
 632 assign.

- 633 • **In a class method, do not return pointers or non-const references**  
 634 **to private data members.** [no-non-const-refs-returned]

635 Otherwise you break the principle of encapsulation.

636 If necessary, you can return a pointer to a const or const reference.

637 This does not mean that you cannot have methods returning an iterator  
638 if your class acts as a container.

639 An allowed exception to this rule is the use of the singleton pattern. In  
640 that case, be sure to add a clear explanation in a comment so that other  
641 developers will understand what you are doing.

- 642 • **Declare as const a member function that does not affect the state  
643 of the object.** [\[const-members\]](#)

644 Declaring a member function as const has two important implications.  
645 First, only const member functions can be called for const objects; and  
646 second, a const member function will not change data members

647 It is a common mistake to forget to const declare member functions that  
648 should be const.

649 This rule does not apply to the case where a member function which does  
650 not affect the state of the object overrides a non-const member function  
651 inherited from some super class.

- 652 • **Do not let const member functions change the state of the program.**  
653 [\[really-const\]](#)

654 A const member function promises not to change any of the data members  
655 of the object. Usually this is not enough. It should be possible to call a  
656 const member function any number of times without affecting the state  
657 of the complete program. It is therefore important that a const member  
658 function refrains from changing static data members or other objects to  
659 which the object has a pointer or reference.

#### 660 3.5.4 Overloading and default arguments

- 661 • **Use function overloading only when methods differ in their argu-  
662 ment list, but the task performed is the same.** [\[function-overloading\]](#)

663 Using function name overloading for any other purpose than to group  
664 closely related member functions is very confusing and is not recommended.

## 665 3.5.5 Comparisons

666 • Define comparisons for custom types using `operator==` and  
667 `operator<=`. [comparison-operators]

668 Comparisons of for a custom class should be written using `operator==`  
669 (for equality/inequality) and `operator<=` (for ordering). The compiler  
670 will supply the other comparison operators (`operator!=`, `operator<`,  
671 etc.) automatically. Where possible, `operator<=` is best defined using  
672 the same operator on the members involved. Examples:

```

1 #include <compare>
2 #include <tuple>
3
4 class S
5 {
6 public:
7     bool operator== (const S& other)
8     {
9         return m_key == other.m_key;
10    }
11    std::strong_ordering operator<= (const S& other)
12    {
13        return m_key <= other.m_key;
14    }
15 private:
16     int m_key;
17 };
18
19
20 class Version
21 {
22 public:
23     bool operator== (const Version& other)
24     {
25         return m_major == other.m_major &&
26                 m_minor == other.m_minor;
27     }
28 }
```

```

29     std::strong_ordering
30     operator<=> (const Version& other)
31     {
32         return
33             std::make_tuple (m_major, m_minor) <=>
34             std::make_tuple (other.m_major, other.m_minor);
35     }
36 private:
37     int m_major;
38     int m_minor;
39 };

```

## 673 3.6 new and delete

- 674 • **Do not use new and delete where automatic allocation will work.**  
[\[auto-allocation-not-new-delete\]](#)

676 Suppose you have a function that takes as an argument a pointer to an  
 677 object, but the function does not take ownership of the object. Then suppose  
 678 you need to create a temporary object to pass to this function. In this case,  
 679 it's better to create an automatically-allocated object on the stack than it  
 680 is to use new / delete. The former will be faster, and you won't have the  
 681 chance to make a mistake by omitting the delete.

```

1 // Not good:
2 Foo* foo = new Foo;
3 doSomethingWithFoo (foo);
4 delete foo;
5
6 // Better:
7 Foo foo;
8 doSomethingWithFoo (&foo);

```

- 682 • **Match every invocation of new with one invocation of delete in**  
**683 all possible control flows from new.** [\[match-new-delete\]](#)

684 A missing delete would cause a memory leak.

685 However, in the Gaudi/Athena framework, an object created with new

686 and registered in StoreGate is under control of StoreGate and must not be  
 687 deleted.

688 In new code, you should generally use `make_unique` for this.

```

1 #include <memory>
2
3 ...
4 DataVector<C>* dv = ...;
5 auto c = std::make_unique<C>("argument");
6 ...
7 if (test) {
8     dv->push_back (std::move (c));
9 }
```

689 `auto_ptr` was an attempt to do something similar to `unique_ptr` in older  
 690 versions of the language. However, it has some serious deficiencies and  
 691 should not be used in new code.

- 692 • **A function should explicitly document if it takes ownership of a  
 693 pointer passed to it as an argument. [explicit-ownership]**

694 The default expectation for a function should be that it does *not* take own-  
 695 ership of pointers passed to it as arguments. In that case, the function must  
 696 *not* invoke `delete` on the pointer, nor pass it to any other function that  
 697 takes ownership.

698 However, if the function is clearly documented as taking ownership of the  
 699 pointer, then it *must* either delete the pointer or pass it to another function  
 700 which will ensure that it is eventually deleted.

701 Rather than simply documenting that a function takes ownership of a  
 702 pointer, it is recommended that you use `std::unique_ptr` to explicitly  
 703 show the transfer of ownership.

```

1 void foo (std::unique_ptr<C> ptr);
2
3 ...
4 std::unique_ptr<C> p (new C);
5 ...
6 foo (std::move (p));
```

```

7  // The argument of foo() is initialized by move.
8  // p is left as a null pointer.

```

704 • **Do not access a pointer or reference to a deleted object.** [deleted-  
705 objects]

706 A pointer that has been used as argument to a `delete` expression should  
707 not be used again unless you have given it a new value, because the language  
708 does not define what should happen if you access a deleted object. This  
709 includes trying to delete an already deleted object. You should assign  
710 the pointer to `nullptr` or a new valid object after the `delete` is called;  
711 otherwise you get a “dangling” pointer.

712 • **After deleting a pointer, assign it to `nullptr`.** [deleted-objects-2]

713 C++ guarantees that deletion of null pointers is safe, so this gives some  
714 safety against double deletes.

```

1 X* myX = makeAnX();
2 delete myX;
3 myX = nullptr;

```

715 This is of course not needed if the pointer is about to go out of scope, or  
716 when objects are deleted in a destructor (unless it's particularly complicated).  
717 But this is a good practice if the pointer persists beyond the block of code  
718 containing the `delete` (especially if it's a member variable).

## 719 3.7 Static and global objects

720 • **Do not declare variables in the global namespace.** [no-global-variables]

721 If necessary, encapsulate those variables in a class or in a namespace. Global  
722 variables violate encapsulation and can cause global scope name clashes.  
723 Global variables make classes that use them context-dependent, hard to  
724 manage, and difficult to reuse.

725 For variables that are used only within one “.cxx” file, put them in an  
726 anonymous namespace.

```

1 namespace {
2   // This variable is visible only in the file

```

```

3  // containing this declaration, and is guaranteed
4  // not to conflict with any declarations from
5  // other files.
6  int counter;
7  }
```

727 • **Do not put functions into the global namespace.** [no-global-functions]

728 Similarly to variables, functions declarations should be put in a namespace.  
 729 If they are used only within one “.cxx” file, then they should be put in an  
 730 anonymous namespace.

731 In a few cases, it might be necessary to declare a function in the global  
 732 namespace to have overloading work properly, but this should be an excep-  
 733 tion.

734 **3.8 Object-oriented programming**

735 • **Do not declare data members to be public.** [no-public-data-members]

736 This ensures that data members are only accessed from within member  
 737 functions. Hiding data makes it easier to change implementation and  
 738 provides a uniform interface to the object.

```

1  class Point
2  {
3  public:
4    Number x() const; // Return the x coordinate
5  private:
6    Number m_x;
7    // The x coordinate (safely hidden)
8  };
```

739 The fact that the class Point has a data member `m_x` which holds the x  
 740 coordinate is hidden.

741 An exception is objects that are intended to be more like C-style structures  
 742 than classes. Such classes should usually not have any methods, except  
 743 possibly a constructor to make initialization easier.

744

745

746

- **If a class has at least one virtual method then it must have a public virtual destructor or (exceptionally) a protected destructor.** [virtual-destructor]

747

748

749

750

751

The destructor of a base class is a member function that in most cases should be declared virtual. It is necessary to declare it virtual in a base class if derived class objects are deleted through a base class pointer. If the destructor is not declared virtual, only the base class destructor will be called when an object is deleted that way.

752

753

754

755

756

757

758

759

760

761

762

There is one case where it is not appropriate to use a virtual destructor: a mix-in class. Such a class is used to define a small part of an interface, which is inherited (mixed in) by subclasses. In these cases the destructor, and hence the possibility of a user deleting a pointer to such a mix-in base class, should normally not be part of the interface offered by the base class. It is best in these cases to have a nonvirtual, nonpublic destructor because that will prevent a user of a pointer to such a base class from claiming ownership of the object and deciding to simply delete it. In such cases it is appropriate to make the destructor protected. This will stop users from accidentally deleting an object through a pointer to the mix-in base-class, so it is no longer necessary to require the destructor to be virtual.

763

764

- **Always re-declare virtual functions as virtual in derived classes.** [redclare-virtual]

765

766

767

768

769

770

771

772

773

This is just for clarity of code. The compiler will know it is virtual, but the human reader may not. This, of course, also includes the destructor, as stated in item [virtual-destructor, page 42]. Virtual functions in derived classes which override methods from the base class should also be declared with the `override` keyword. If the signature of the method is changed in the base class, so that the declaration in the derived class is no longer overriding it, this will cause the compiler to flag an error. (As an exception, `override` is not required for destructors. Since there is only one possible signature for a destructor, `override` doesn't add anything.)

```
1  class B
2  {
3  public:
4      virtual void foo(int);
5  };
```

```

6
7 class D : public B
8 {
9 public:
10    // Declare foo as a virtual method that overrides
11    // a method from the base class.
12    virtual void foo(int) override;
13 }

```

774 • **Avoid multiple inheritance, except for abstract interfaces.** [no-  
775 multiple-inheritance]

776 Multiple inheritance is seldom necessary, and it is rather complex and error  
777 prone. The only valid exception is for inheriting interfaces or when the  
778 inherited behavior is completely decoupled from the class's responsibility.

779 For a detailed example of a reasonable application of multiple inheritance,  
780 see [12], item 43.

781 • **Avoid the use of friend declarations.** [no-friend]

782 Friend declarations are almost always symptoms of bad design and they  
783 break encapsulation. When you can avoid them, you should.

784 Possible exceptions are the streaming operators and binary operators on  
785 classes. Other possible exceptions include very tightly coupled classes and  
786 unit tests.

787 • **Avoid the use of protected data members.** [no-protected-data]

788 Protected data members are similar to friend declarations in that they  
789 allow a controlled violation of encapsulation. However, it is even less well-  
790 controlled in the case of protected data, since any class may derive from  
791 the base class and access the protected data.

792 The use of protected data results in one class depending on the internals  
793 of another, which is a maintenance issue should the base class need to  
794 change. Like friend declarations, the use of protected member data should  
795 be avoided except for very closely coupled classes (that should generally be  
796 part of the same package). Rather, you should define a proper interface for  
797 what needs to be done (parts of which may be protected).

798 **3.9 Notes on the use of library functions.**

799 • Use `std::abs` to calculate an absolute value. [\[std-abs\]](#)

800 The return type of `std::abs` will conform to the argument type; other  
801 variants of `abs` may not do this.

802 In particular, beware of this:

```
1 #include <cstdlib>
2 float foo (float x)
3 {
4     return abs(x);
5 }
```

803 which will truncate `x` to an integer. (Clang will warn about this.)

804 Conversely, in this example:

```
1 #include <cmath>
2 int (int x)
3 {
4     return fabs(x);
5 }
```

805 the argument will first be converted to a float, then the result converted  
806 back to an integer.

807 Using `std::abs` uniformly should do the right thing in almost all cases  
808 and avoid such surprises.

809 • Use C++20 ranges with caution. [\[ranges\]](#)

810 C++20 adds `ranges`, an abstraction an abstraction of something that can be  
811 iterated over. Essentially, a range is something that can return `begin()` and  
812 `end()` iterators. The `ranges` library allows composing and transforming  
813 ranges. For example:

```
1 #include <ranges>
2 ...
3 auto even = [](int i) { return (i%2) == 0; };
4 auto sq = [](int i) { return i*i; };
```

```

5  using namespace std::views;
6  auto r = iota(0, 6) | filter(even) | transform(sq);
7  for (int i : r) std::cout << i << " ";

```

814 Ranges can be very useful. However, they need to be used with caution.

815 – Do not reimplement missing functionality yourself.

816 Much of that C++20 ranges library originated from an external library,  
 817 range-v3 [14]. However, many useful operations from that library  
 818 did not make it into the C++20 standard (some are added in later  
 819 versions of the standard). For example, in C++20 ranges, there is no  
 820 straightforward way to initialize a `std::vector` from a range. If  
 821 such additional functionality is needed, it should be added centrally  
 822 in `CxxUtils` rather than being reimplemented where it is needed.  
 823 In that way, it can be shared with other parts of Athena. This also  
 824 makes it easier to replace any such reimplemented functionality with  
 825 versions from the standard library when they become available.

826 – Functions used to define ranges should not have side effects.

827 One can define a range in terms of functions that filter and transform  
 828 the range, as in the example above. However, it may be difficult  
 829 to predict under exactly what circumstances these functions may be  
 830 called, as this depends on the implementation of the range components.  
 831 Therefore, functions used with ranges should not have side-effects  
 832 (and should generally execute quickly).

833 – Beware of dangling ranges.

834 Ranges are often references to other objects. Like any references, they  
 835 must not outlive the object that they reference.

```

1  auto squares()
2  {
3      auto sq = [](int i) { return i*i; };
4      std::vector<int> v {1, 2, 3, 4};
5      return v | std::views::transform(sq);
6      // BAD: returns a range with a dangling
7      //       reference to a deleted vector.

```

8      | }

836    – Do not modify containers referenced by ranges.

837    Similarly, do not modify a container referenced by a range. Some  
 838    of the range components may cache results internally; changing the  
 839    underlying container may cause these to return incorrect results.

```
1  std::vector<int> v {1, 2, 3, 4};
2  auto sq = [] (int i) { return i*i; };
3  auto r = v | std::views::transform(sq);
4  v.insert (v.begin(), 5); // BAD: may invalidate
5                                // the range r.
```

840    In general, C++20 view objects should be used directly after they are defined,  
 841    and not saved in, say, member variables.

## 842    3.10 Thread friendliness and thread safety

843    Code that is to be run in AthenaMT as part of an `AthAlgorithm` must be “thread-friendly.” While the framework will ensure that no more than one thread is  
 844    executing a given `AthAlgorithm` instance at one time, the code must ensure  
 845    that it doesn’t interfere with *other* threads. Some guidelines for this are outlined  
 846    below; but in brief: don’t use static data, don’t use `mutable`, and don’t cast away  
 847    `const`. Following these rules will keep you out of most potential trouble.

848    Code that runs as part of an `AthService`, an `AthReentrantAlgorithm`, a data  
 849    object implementation, or other common code may need to be fully “thread-safe;”  
 850    that is, allow for multiple threads to operate simultaneously on the *same* object.  
 851    The easiest way to ensure this is for the object to have no mutable internal state,  
 852    and only `const` methods. If, however, some threads may be modifying the state  
 853    of the object, then some sort of locking or other means of synchronization will  
 854    likely be required. A full discussion of this is beyond the scope of these guidelines.

855    To run successfully in a multithreaded environment, algorithmic code must also  
 856    respect the rules imposed by the framework on event and conditions data access.  
 857    This is also beyond the scope of these guidelines.

- 859    • **Follow C++ thread-safety conventions for data objects.** [\[mt-follow-c++-conventions\]](#)

861 The standard C++ container objects follow the rule that methods declared  
 862 as `const` are safe to call simultaneously from multiple threads, while no  
 863 non-`const` method can be called simultaneously with any other method  
 864 (`const` or non-`const`) on the same object.

865 Classes meant to be data objects should generally follow the same rules,  
 866 unless there is a good reason to the contrary. This will generally happen  
 867 automatically if the rules outlined below are followed: briefly, don't use  
 868 static data, don't use `mutable`, and don't cast away `const`.

869 Sometimes it may be useful to have data classes for which non-`const` meth-  
 870 ods may be called safely from multiple threads. If so, this should be indicated  
 871 in the documentation of the class, and perhaps hinted from its name (maybe  
 872 like `ConcurrentFoo`).

873 • **Do not use non-`const` static variables** [mt-no-nonconst-static]

874 Do not use non-`const` static variables in thread-friendly code, either global  
 875 or local.

```

1  int a;
2  int foo() {
3      if (a > 0) { // Bad use of global static.
4          static int count = 0;
5          return ++count; // Bad use of local static.
6      }
7      return 0;
8  }
9
10 struct Bar
11 {
12     static int s_x;
13     int x() { return s_x; } // Bad use of static
14                           // class member.
15 }
```

876 A `const static` is, however, perfectly fine:

```

1  static const std::string s = "a string"; // OK, const
```

877

It's generally OK to have static mutex or thread-local variables:

```

1  static std::mutex m; // OK. It's a mutex,
2                                // so it's meant to be accessed
3                                // from multiple threads.
4  static thread_local int a; // OK, it's thread-local.

```

878

879

880

(Be aware, though, that thread-local variables can be quite slow.) A static `std::atomic<T>` variable may be OK, but only if it doesn't need to be updated consistently with other variables.

881

- **Do not cast away `const`** [mt-no-const-cast]

882

883

884

885

886

887

888

This rule was already mentioned above. However, it deserves particular emphasis in the context of thread safety. The usual convention for C++ is that a `const` method is safe to call simultaneously from multiple threads, while if you call a non-`const` method, no other threads can be simultaneously accessing the same object. If you cast away `const`, you are subverting these guarantees. Any use of `const_cast` needs to be analyzed for its effects on thread-safety and possibly protected with locking.

889

For example, consider this function:

```

1  void foo (const std::vector<int>& v)
2  {
3      ...
4      // Sneak this in.
5      const_cast<std::vector<int>&>(v).push_back(10);
6  }

```

890

891

892

893

894

Someone looking at the signature of this function would see that it takes only a `const` argument, and therefore conclude that that it is safe to call this simultaneously with other code that is also reading the same vector instance. But it is not, and the `const_cast` is what causes that reasoning to fail.

895

- **Avoid mutable members.** [mt-no-mutable]

896

897

898

The use of `mutable` members has many of the same problems as `const_cast` (as indeed, `mutable` is really just a restricted version of `const_cast`). A `mutable` member can generally not be changed

899 from a non-const method without some sort of explicit locking or other  
 900 synchronization. It is best avoided in code that should be used with  
 901 threading.

902 `mutable` can, however, be used with objects that are explicitly intended to  
 903 be accessed from multiple threads. These include mutexes and thread-local  
 904 pointers. In some cases, members of `atomic` type may also be safely made  
 905 `mutable`, but only if they do not need to be updated consistently with  
 906 other members.

907 • **Do not return non-const member pointers/references from const  
 908 methods** [\[mt-const-consistency\]](#)

909 Consider the following fragment:

```
1 class C
2 {
3     public:
4         Impl* impl() const { return m_impl; }
5     private:
6         Impl* m_impl;
7 }
```

910 This is perfectly valid according to the C++ const rules. However, it allows  
 911 modifying the `Impl` object following a call to the `const` method `impl()`.  
 912 Whether this is actually a problem depends on the context. If `m_impl` is  
 913 pointing at some unrelated object, then this might be OK; however, if it  
 914 is pointing at something which should be considered part of `C`, then this  
 915 could be a way around the `const` guarantees.

916 To maintain safety, and to make the code easier to reason about, do not  
 917 return a non-const pointer (or reference) member from a `const` member  
 918 function.

919 • **Be careful returning const references to class members.** [\[mt-const-  
 920 references\]](#)

921 Consider the following example:

```
1 class C
2 {
```

```

3  public:
4      const std::vector<int>& v() const { return m_v; }
5      void append (int x) { m_v.push_back (x); }
6  private:
7      std::vector<int> m_v;
8  };
9
10 int getSize (const C& c)
11 {
12     return c.v().size();
13 }
14
15 int push (C& c)
16 {
17     c.append (1);
18 }
```

922 This is a fairly typical example of a class that has a large object as a member,  
 923 with an accessor the returns the member by const reference to avoid having  
 924 to do a copy.

925 But suppose now that one thread calls `getSize()` while another thread  
 926 calls `push()` at the same time on the same object. It can happen that first  
 927 `getSize()` gets the reference and starts the call to `size()`. At that point,  
 928 the `push_back()` can run in the other thread. If `push_back()` runs at  
 929 the same time as `size()`, then the results are unpredictable – the `size()`  
 930 call could very well return garbage.

931 Note that it doesn't help to add locking within the class `C`:

```

1  class C
2  {
3  public:
4      const std::vector<int>& v() const
5      {
6          std::lock_guard<std::mutex> lock (m_mutex);
7          return m_v;
8      }
9      void append (int x)
```

```

10  {
11      std::lock_guard<std::mutex> lock (m_mutex);
12      m_v.push_back (x);
13  }
14 private:
15     mutable std::mutex m_mutex;
16     std::vector<int> m_v;
17 }
```

932 This is because the lock is released once `v()` returns — and at that point,  
 933 the caller can call (const) methods on the vector instance unprotected  
 934 by the lock.

935 Here are a few ways in which this could possibly be solved. Which is  
 936 preferable would depend on the full context in which the class is used.

- 937 – Change the `v()` accessor to return the member by value instead of by  
 938 reference.
- 939 – Remove the `v()` accessor and instead add the needed operations to  
 940 the `C` class, with appropriate locking. For the above example, we could  
 941 add something like:

```

1 size_t C::vSize() const
2 {
3     std::lock_guard<std::mutex> lock (m_mutex);
4     return m_v.size();
5 }
```

- 942 – Change the type of the `m_v` member to something that is inherently  
 943 thread-safe. This could mean replacing it with a wrapper around  
 944 `std::vector` that does locking internally, or using something like  
 945 `concurrent_vector` from TBB.
- 946 – Do locking externally to class `C`. For example, introduce a mutex  
 947 that must be acquired in both `getSize()` and `push()` in the above  
 948 example.

## 949 3.11 Formatted output

950 • Prefer `std::format` to `printf` or `iostream` formatting. [use-  
951 format]

952 For new code, use the C++20 formatting library to format values to a string  
953 rather than using `printf`-style formatting or using `iostream` manipula-  
954 tors.

955 Example:

```

1 #include <format>
2 ...
3 const char* typ = "ele";
4 float energy = 14.2345;
5 int mask = 323;
6
7 std::cout << std::format
8     ("A {1:.2f} GeV {0} mask {2:#06x}.\n",
9      typ, energy, mask);
10 // prints: A 14.23 GeV ele mask 0x0143.

```

956 Compare using `printf`-style formatting:

```

1 #include "CxxUtils/StrFormat.h"
2 ...
3 std::cout << CxxUtils::strformat
4     ("A %.2f GeV %s mask %#06x.\n",
5      energy, typ, mask);

```

957 or `iostream`:

```

1 #include <iomanip>
2 ...
3 const int default_precision = std::cout.precision();
4 const std::ios_base::fmtflags default_flags =
5     std::cout.flags();
6 const char default_fill = std::cout.fill();
7 std::cout << "A " << std::fixed << std::setprecision(2)
8             << energy << std::defaultfloat

```

```

9      << std::setprecision(default_precision)
10     << " GeV " << typ << " mask "
11     << std::hex << "0x" << std::setfill('0')
12     << std::setw(4) << mask
13     << std::setfill(default_fill)
14     << ".\n";
15 std::cout.flags(default_flags);

```

Like the streaming operator, `std::format` has a way of customizing how a given type is formatted. However, it is somewhat more involved than `operator<<`; in addition, `std::format` will not use existing custom streaming operators. Therefore, for generating printable representations of class instances, it is probably better in most cases to use the `iostream` mechanism.

## 3.12 Assertions and error conditions

- **Pre-conditions and post-conditions should be checked for validity.** [\[pre-post-conditions\]](#)

You should validate your input and output data whenever an invalid input can cause an invalid output.

- **Don't use assertions in place of exceptions.** [\[assertion-usage\]](#)

Assertions should only be used to check for conditions which should be logically impossible to occur. Do not use them to check for validity of input data. For such cases, you should raise an exception (or return a Gaudi error code) instead.

Assertions may be removed from production code, so they should not be used for any checks which must always be done.

## 3.13 Error handling

- **Use the standard error printing facility for informational messages. Do not use `cerr` and `cout`.** [\[no-cerr-cout\]](#)

The “standard error printing facility” in Athena/Gaudi is `MsgStream`. No production code should use `cout`. Classes which are not Athena-aware

981 could use `cerr` before throwing an exception, but all Athena-aware classes  
 982 should use `MSG::FATAL` and/or throw an exception. In addition, it is ac-  
 983 ceptable to use writes to `cout` in unit tests.

984 When using `MsgStream`, note that a call to, e.g., `msg() << MSG::VERBOSE` that is suppressed by the output level has a higher  
 985 runtime cost than a call suppressed by `if (msgLvl <= MSG::VERBOSE)`.  
 986 The `ATH_MSG` macros (`ATH_MSG_INFO` and `ATH_MSG_DEBUG` etc) wrap  
 987 `msg()` calls in appropriate `if` statements and are preferred in general  
 988 for two reasons: they take up less space in the source code and indicate  
 989 immediately that the message is correctly handled.

991 **• Check for all errors reported from functions.** [\[check-return-status\]](#)

992 It is important to always check error conditions, regardless of how they are  
 993 reported.

994 **• Use exceptions to report fatal errors from non-Gaudi components.**  
 995 [\[exceptions\]](#)

996 Exceptions in C++ are a means of separating error reporting from error  
 997 handling. They should be used for reporting errors that the calling code  
 998 should not be expected to handle. An exception is “thrown” to an error  
 999 handler, so the treatment becomes non-local.

1000 If you are writing a Gaudi component, or something that returns a Gaudi  
 1001 `StatusCode`, then you should usually report an error by posting a message  
 1002 to the message service and returning a status code of `ERROR`.

1003 However, if you are writing a non-Gaudi component and you need to report  
 1004 an error that should stop event processing, you should raise an exception.

1005 If your code is throwing exceptions, it is helpful to define a separate class  
 1006 for each exception that you throw. That way, it is easy to stop in the  
 1007 debugger when a particular exception is thrown by putting a breakpoint in  
 1008 the constructor for that class.

```

1 #include <stdexcept>
2
3 class ExcMyException
4   : public std::runtime_error
5 {
```

```

1 public:
2     // Constructor can take arguments to pass through
3     // additional information.
4     ExcMyException (const std::string& what)
5         : std::runtime_error ("My exception: " : what)
6     {}
7
8     ...
9
10
11
12
13
14
15
16     throw MyException ("You screwed up.");

```

- **Do not throw exceptions as a way of reporting uncommon values from a function.** [\[exception-usage\]](#)

If an error *can* be handled locally, then it should be. Exceptions should not be used to signal events which can be expected to occur in a regular program execution. It is up to programmers to decide what it means to be exceptional in each context.

Take for example the case of a function `find()`. It is quite common that the object looked for is not found, and it is certainly not a failure; it is therefore not reasonable in this case to throw an exception. It is clearer if you return a well-defined value.

- **Do not use exception specifications.** [\[no-exception-specifications\]](#)

Exception specifications were a way to declare that a function could throw one of only a restricted set of exceptions. Or rather, that's what most people wanted it to do; what it actually did was require the compiler to check, at runtime, that a function did not throw any but a restricted set of exceptions.

Experience has shown that exception specifications are generally not useful and non-empty exception specifications are now an error [\[15\]](#). They should not be used in new code, and are not allowed in C++20.

There is also the keyword `noexcept`. The motivation for this was really to address a specific problem with move constructors and exception-safety, and it is not clear that it is generally useful [\[16\]](#). For now, it is not recommended to use `noexcept`, unless you have a specific situation where you know it

1031 would help.

1032 • **Do not catch a broad range of exceptions outside of framework code.**  
 1033 [no-broad-exception-catch]

1034 The C++ exception mechanism allows catching a thrown exception, giving  
 1035 the program the chance to continue execution from the point where the  
 1036 exception was caught. This can be used some specific cases where you  
 1037 know that some specific exception isn't really a problem. However, you  
 1038 should catch only the particular exception involved here. If you use an  
 1039 overly-broad catch specification, you risk hiding other problems. Example:

```

1  try {
2   return getObject ("foo");
3   // getObject may throw ExcNotFound if the "foo"
4   // object is not found. In that case we can just
5   // return 0.
6 }
7 catch (ExcNotFound&) {
8   return 0;
9 }
10
11 // But one would not want to do this, since that would
12 // hide other errors:
13 catch (...) {
14   return 0;
15 }
```

1040 • **Prefer to catch exceptions as `const` reference, rather than as value.**  
 1041 [catch-const-reference]

1042 Classes used for exceptions can be polymorphic just like data classes, and  
 1043 this is in fact the case for the standard C++ exceptions. However, if you  
 1044 catch an exception and name the base class by value, then the object thrown  
 1045 is copied to an instance of the base class.

1046 For example, consider this program:

```

1  #include <stdexcept>
2  #include <iostream>
```

```

3
4 class myex : public std::exception {
5 public:
6     virtual const char* what() const noexcept
7     { return "Mine!"; }
8 }
9
10 void foo()
11 {
12     throw myex();
13 }
14
15 int main()
16 {
17     try {
18         foo();
19     }
20     catch (std::exception ex) {
21         std::cout << "Exception: " << ex.what() << "\n";
22     }
23     return 0;
24 }
```

1047 It looks like the intention here is to have a custom message printed when the  
 1048 exception is caught. But that's not what happens — this program actually  
 1049 prints:

1   Exception: std::exception

1050 That's because in the catch clause, the myex instance is copied to a  
 1051 std::exception instance, so any information about the derived myex  
 1052 class is lost. If we change the catch to use a reference instead:

1   **catch** (**const** std::exception **ex**&) {

1053 then the program prints what was probably intended.

1   Exception: Mine!

1054 Recent versions of gcc will warn about this.

## 1055 3.14 Parts of C++ to avoid

1056 Here a set of different items are collected. They highlight parts of the language  
1057 which should be avoided, either because there are better ways to achieve the  
1058 desired results or because the language features are still immature. In particular,  
1059 programmers should avoid using the old standard C functions, where C++ has  
1060 introduced new and safer possibilities.

- 1061 • **Do not use C++ modules.** [\[no-modules\]](#)

1062 Modules were introduced in C++20 as a better alternative to `#include`.  
1063 If a module is referenced via `import`, it avoids repeatedly parsing the  
1064 code as well as avoiding issues that arise due to interference between  
1065 headers. However, building modules requires significant support from the  
1066 build system, and the support in compilers and associated tools is still  
1067 very immature. Even using the standard library as a module is not fully  
1068 functional with C++20.

1069 For now, avoid any use of modules. With C++23, it may be possible to use  
1070 standard libraries as modules, but building ATLAS code as modules will  
1071 require significant additional development.

- 1072 • **Do not use C++ coroutines.** [\[no-coroutines\]](#)

1073 Coroutines allow for a non-linear style of control flow, where one can return  
1074 from the middle of a function and then resume execution from that point at  
1075 a later time. However, the coroutine interfaces available in C++20 are quite  
1076 low-level: they are intended to be used as building blocks for other library  
1077 components rather than for direct use by user code. Further, uncontrolled  
1078 use of the type of control flow made possible by coroutines has the potential  
1079 to be terribly confusing.

1080 For now, avoid use of coroutines. If you have a use case that would greatly  
1081 benefit from using coroutines, please consult with software coordination.  
1082 This recommendation will be revisited for new versions of C++ which may  
1083 include easier mechanisms for using coroutines.

- 1084 • **Do not use `malloc`, `calloc`, `realloc`, and `free`. Use `new` and  
1085 `delete` instead.** [\[no-malloc\]](#)

1086 You should avoid all memory-handling functions from the standard C-  
1087 library (`malloc`, `calloc`, `realloc`, and `free`) because they do not call  
1088 constructors for new objects or destructors for deleted objects.

1089 Exceptions may include aligned memory allocations, but this should gener-  
1090 ally not be done outside of low-level code in core packages.

- 1091 • **Do not use functions defined in `stdio`. Use the `iostream` func-  
1092 tions in their place.** [\[no-stdio\]](#)

1093 `scanf` and `printf` are not type-safe and they are not extensible. Use  
1094 `operator>>` and `operator<<` associated with C++ streams instead,  
1095 along with `std::format` to handle formatting (see [use-format, page 52](#)).  
1096 `iostream` and `stdio` functions should never be mixed.

1097 Example:

```
1 // type safety
2 char* aString("Hello Atlas");
3 printf("This works: %s \n", aString);
4 cout <<"This also works :"<<aString<<endl;
5 char aChar('!');
6 printf("This does not %s \n", aChar);
7     // and you get a core dump
8 cout <<"But this is still OK :"<<aChar<<endl;
9
10 //extensibility
11 std::string aCPPString("Hello Atlas");
12 printf("This does not work: %s \n", aCPPString);
13     //Core dump again
```

1098 It is of course acceptable to use `stdio` functions if you're calling an external  
1099 library that requires them.

1100 If you need to use `printf` style formatting, see “`CxxUtils/StrFormat.h`.”  
1101 However, `std::format` is preferred for new code.

- 1102 • **Do not use the ellipsis notation for function arguments.** [\[no-ellipsis\]](#)

1103 Prior to C++ 11, functions with an unspecified number of arguments had  
1104 to be declared and used in a type-unsafe manner:

```

1  // avoid to define functions like:
2  void error(int severity, ...) // "severity" followed
3                                // by a zero-terminated
4                                // list of char*s

```

1105 This method should be avoided.

1106 As of C++11, one can accomplish something similar using variadic tem-  
1107 plates:

```

1  template<typename ...ARGS>
2  void error(int severity, ARGS...)

```

1108 This is fine, but should be used judiciously. It's appropriate for forwarding  
1109 arguments through a template function. For other cases, it's worth thinking  
1110 if there might be a simpler way of doing things.

1111 An ellipsis can also occur in a catch clause to catch any exception:  
1112 `catch(...)`. This is acceptable, but should generally be restricted to  
1113 framework-like code.

- 1114 • **Do not use preprocessor macros to take the place of functions, or  
1115 for defining constants.** [\[no-macro-functions\]](#)

1116 Use templates or inline functions rather than the pre-processor macros.

```

1  // NOT recommended to have function-like macro
2  #define SQUARE(x) x*x

3
4  // Better to define an inline function:
5  inline int square(int x) {
6      return x*x;
7  }

```

- 1117 • **Do not declare related numerical values as `const`. Use `enum` decla-  
1118 rations.** [\[use-enum\]](#)

1119 The `enum` construct allows a new type to be defined and hides the numerical  
1120 values of the enumeration constants.

```
1 enum State {halted, starting, running, paused};
```

- 1121 • **Do not use `NULL` to indicate a null pointer; use the `nullptr` keyword instead.** [\[nullptr\]](#)

1123 Older code often used the constant 0. `NULL` is appropriate for C, but not  
1124 C++.

- 1125 • **Do not use `const char*` or built-in arrays “[ ]”; use `std::string` instead.** [\[use-std-string\]](#)

1127 One thing to be aware of, though. C++ will implicitly convert a `const`  
1128 `char*` to a `std::string`; however, this may add significant overhead if  
1129 used in a loop. For example:

```
1 void do_something (const std::string& s);
2 ...
3 for (int i=0; i < lots; i++) {
4 ...
5 do_something ("hi there!");
```

1130 Each time through the loop, this will make a new `std::string` copy of  
1131 the literal. Better to move the conversion to `std::string` outside of the  
1132 loop:

```
1 std::string myarg = "hi there!";
2 for (int i=0; i < lots; i++) {
3 ...
4 do_something (myarg);
```

- 1133 • **Avoid using union types.** [\[avoid-union-types\]](#)

1134 Unions can be an indication of a non-object-oriented design that is hard to  
1135 extend. The usual alternative to unions is inheritance and dynamic binding.  
1136 The advantage of having a derived class representing each type of value  
1137 stored is that the set of derived class can be extended without rewriting any  
1138 code. Because code with unions is only slightly more efficient, but much  
1139 more difficult to maintain, you should avoid it.

1140 Unions may be used in some low-level code and in places where efficiency

1141 is particularly important. Unions may also be used in low-level code to  
 1142 avoid pointer aliasing (see [no-reinterpret-cast, page 28](#)).

1143 • **Avoid using bit fields.** [\[avoid-bitfields\]](#)

1144 Bit fields are a feature that C++ inherited from C that allow one to specify  
 1145 that a member variable should occupy only a specified number of bits, and  
 1146 that it can be packed together with other such members.

```

1 class C
2 {
3 public:
4   unsigned int a : 2; // Allocated two bits
5   unsigned int b : 3; // Allocated three bits
6 }
```

1147 It may be tempting to use bit fields to save space in data written to disk, or  
 1148 in packing and unpacking raw data. However, this usage is not portable.  
 1149 The C++ standard has this to say:

1150 Allocation of bit-fields within a class object is implementation-  
 1151 defined. Alignment of bit-fields is implementation-defined. Bit-  
 1152 fields are packed into some addressable allocation unit. [ Note:  
 1153 Bit-fields straddle allocation units on some machines and not on  
 1154 others. Bit-fields are assigned right-to-left on some machines,  
 1155 left-to-right on others. – end note ]

1156 Besides portability issues, there are other other potential issues with bit  
 1157 fields that could be confusing: bit fields look like class members but obey  
 1158 subtly different rules. For example, one cannot form a reference to a bit  
 1159 field or take its address. There is also an issue of data races when writing  
 1160 multithreaded code. It is safe to access two ordinary class members sim-  
 1161 ultaneously from different threads, but not two adjacent bit fields. (Though  
 1162 it is safe to access simultaneously two bit field members separated by an  
 1163 ordinary member. This leads to the possibility that thread-safety of bit field  
 1164 access could be compromised by the removal of an unrelated member.)  
 1165 Access to bit fields also incurs a CPU penalty.

1166 In light of this, it is best to avoid bit fields in most cases. Exceptions would  
 1167 be cases where saving memory is very important and the internal structure

1168 of the class is not exposed.

1169 For some cases, `std::bitset` can be a useful, portable replacement for  
1170 bit fields.

1171 • **Do not use `asm` (the assembler macro facility of C++). [no-asm]**

1172 Many special-purpose instructions are available through the use of compiler  
1173 intrinsic functions. For those rare use cases where an `asm` might be needed,  
1174 the use of the `asm` should be encapsulated and made available in a low-level  
1175 package (such as `CxxUtils`).

1176 • **Do not use the keyword `struct` for types used as classes. [no-struct]**

1177 The `class` keyword is identical to `struct` except that by default its con-  
1178 tents are private rather than public. `struct` may be allowed for writing  
1179 non-object-oriented PODs (plain old data, i.e. C structs) on purpose. It is a  
1180 good indication that the code is on purpose not object-oriented.

1181 • **Do not use static objects at file scope. Use an anonymous namespace  
1182 instead. [anonymous-not-static]**

1183 The use of `static` to signify that something is private to a source file is  
1184 obsolete; further it cannot be used for types. Use an anonymous namespace  
1185 instead.

1186 For entities which are not public but are also not really part of a class, prefer  
1187 putting them in an anonymous namespace to putting them in a class. That  
1188 way, they won't clutter up the header file.

1189 • **Do not declare your own alias for booleans. Use the `bool` type of  
1190 C++ for booleans. [use-bool]**

1191 The `bool` type was not implemented in C. Programmers usually got around  
1192 the problem by `typedefs` and/or `const` declarations. This is no longer needed,  
1193 and must not be used in ATLAS code.

1194 • **Avoid pointer arithmetic. [no-pointer-arithmetic]**

1195 Pointer arithmetic reduces readability, and is extremely error prone. It  
1196 should be avoided outside of low-level code.

1197 • **Do not declare variables with `register`. [no-register]**

1198 The `register` keyword was originally intended as a hint to the compiler  
1199 that a variable will be used frequently, and therefore it would be good to  
1200 assign a dedicated register to that variable. However, compilers have long  
1201 been able to do a good job of assigning values to registers; this is anyway  
1202 highly-machine dependent.

1203 Use of the `register` keyword now an error.

## 1204 3.15 Readability and maintainability

- 1205 • **Code should compile with no warnings.** [\[no-warnings\]](#)

1206 Many compiler warnings can indicate potentially serious problems with  
1207 your code. But even if a particular warning is benign, it should be fixed, if  
1208 only to prevent other people from having to spend time examining it in the  
1209 future.

1210 Warnings coming from external libraries should be reported to whomever is  
1211 maintaining the ATLAS wrapper package for the library. Even if the library  
1212 itself can't reasonably be fixed, it may be possible to put a workaround in  
1213 the wrapper package to suppress the warning.

1214 See [\[17\]](#) for help on how to get rid of many common types of warning. If it  
1215 is really impossible to get rid of a warning, that fact should be documented  
1216 in the code.

- 1217 • **Keep functions short.** [\[short-functions\]](#)

1218 Short functions are easier to read and reason about. Ideally, a single function  
1219 should not be bigger than can fit on one screen (i.e., not more than 30–40  
1220 lines).

- 1221 • **Avoid excessive nesting of indentation.** [\[excessive-nesting\]](#)

1222 It becomes difficult to follow the control flow in a function when it becomes  
1223 deeply nested. If you have more than 4–5 indentation levels, consider  
1224 splitting off some of the inner code into a separate function.

- 1225 • **Avoid duplicated code.** [\[avoid-duplicate\]](#)

1226 This statement has a twofold meaning.

1227 The first and most evident is that one must avoid simply cutting and pasting  
1228 pieces of code. When similar functionalities are necessary in different  
1229 places, they should be collected in methods, and reused.

1230 The second meaning is at the design level, and is the concept of code reuse.  
1231 Reuse of code has the benefit of making a program easier to understand  
1232 and to maintain. An additional benefit is better quality because code that is  
1233 reused gets tested much better.

1234 Code reuse, however, is not the end-all goal, and in particular, it is less  
1235 important than encapsulation. One should not use inheritance to reuse a  
1236 bit of code from another class.

1237 **• Document in the code any cases where clarity has been sacrificed  
1238 for performance. [\[document-changes-for-performance\]](#)**

1239 Optimize code only when you know you have a performance problem. This  
1240 means that during the implementation phase you should write code that is  
1241 easy to read, understand, and maintain. Do not write cryptic code, just to  
1242 improve its performance.

1243 Very often bad performance is due to bad design. Unnecessary copying of  
1244 objects, creation of large numbers of temporary objects, improper inheri-  
1245 tance, and a poor choice of algorithms, for example, can be rather costly  
1246 and are best addressed at the architecture and design level.

1247 **• Avoid creating type aliases for classes. [\[avoid-typedef\]](#)**

1248 Type aliases (typedefs) are a serious impediment in large systems. While  
1249 they simplify code for the original author, a system filled with aliases can  
1250 be difficult to understand. If the reader encounters a class A, he or she can  
1251 find an `#include` with “A.h” in it to locate a description of A; but aliases  
1252 carry no context that tell a reader where to find a definition. Moreover,  
1253 most of the generic characteristics obtained with aliases are better handled  
1254 by object oriented techniques, like polymorphism.

1255 Aliases are acceptable where they provide part of the expected interface for  
1256 a class, for example `value_type`, etc. in classes used with STL algorithms.  
1257 They are often indispensable in template programming and metaprogram-  
1258 ming, and are also part of how xAOD classes and POOL converters are  
1259 typically defined.

1260 In other contexts, they should be used with care, and should generally be  
 1261 accompanied with a comment giving the rationale for the alias.

1262 Aliases may be used as a “customization point;” that is, to allow the pos-  
 1263 sibility of changing a type in the future. For example, the auxiliary store  
 1264 code uses integers to identify particular auxiliary data items. But rather  
 1265 than declaring these as an integer type directly, an alias `auxid_t` is used.  
 1266 This allows for the possibility of changing the type in the future without  
 1267 having to make changes throughout the code base. It also makes explicit  
 1268 that variables of that type are meant to identify auxiliary data items, rather  
 1269 than being random integers.

1270 An alias may also be used inside a function body to shorten a cumbersome  
 1271 type name; however, this should be used sparingly.

- 1272 • **Code should use the standard ATLAS units for time, distance, energy,  
 1273 etc.** [\[atlas-units\]](#)

1274 As a reminder, energies are represented as MeV and lengths as mm. Please  
 1275 use the symbols defined in `GaudiKernel/SystemOfUnits.h`.

```
1 #include "GaudiKernel/SystemOfUnits.h"
2
3 float pt_thresh = 20 * Gaudi::Units::GeV;
4 float ip_cut = 0.1 * Gaudi::Units::cm;
```

## 1276 3.16 Portability

- 1277 • **All code must comply with the 2020 version of the ISO C++ standard  
 1278 (C++20).** [\[standard-cxx\]](#)

1279 A draft of the standard which is essentially identical to the final version may  
 1280 be found at [\[4\]](#). However, the standards documents are not very readable.  
 1281 A better reference for most questions about what is in the standard is the  
 1282 `cppreference.com` website [\[5\]](#).

1283 At some point, compatibility with C++23 will also be required.

- 1284 • **Make non-portable code easy to find and replace.** [\[limit-non-portable-  
 1285 code\]](#)

1286 Non-portable code should preferably be factored out into a low-level pack-  
 1287 age in Control, such as CxxUtils. If that is not possible, an `#ifdef` may  
 1288 be used.

1289 However, `#ifdef`s can make a program completely unreadable. In addition,  
 1290 if the problems being solved by the `#ifdef` are not solved centrally by the  
 1291 release tool, then you resolve the problem over and over. Therefore, the  
 1292 using of `#ifdef` should be limited.

1293 **• Headers supplied by the implementation (system or standard li-  
 1294 braries header files) must go in `<>` brackets; all other headers must  
 1295 go in `" "` quotes. [system-headers]**

```
1 // Include only standard header with <>
2 #include <iostream> // OK: standard header
3 #include <MyFyle.hh> // NO: nonstandard header
4
5 // Include any header with ""
6 #include "stdlib.h" // NO: better to use <>
7 #include "MyPackage/MyFyle.h" // OK
```

1296 **• Do not specify absolute directory names in include directives. In-  
 1297 stead, specify only the terminal package name and the file name.  
 1298 [include-path]**

1299 Absolute paths are specific to a particular machine and will likely fail  
 1300 elsewhere.

1301 The ATLAS convention is to include the package name followed by the file  
 1302 name. Watch out: listing the package name twice is wrong, but some build  
 1303 systems don't catch it.

```
1 #include "/atlas/sw/dist/1.2/Foo/Bar/Qux.h"
2 // Wrong
3 #include "Foo/Bar/Qux.h" // Wrong
4 #include "Bar/Bar/Qux.h" // Wrong
5 #include "Bar/Qux.h" // Right
```

1304 **• Always treat include file names as case-sensitive. [include-case-  
 1305 sensitive]**

1306 Some operating systems, e.g. Windows NT, do not have case-sensitive  
 1307 file names. You should always include a file as if it were case-sensitive.  
 1308 Otherwise your code could be difficult to port to an environment with  
 1309 case-sensitive file names.

```
1 // Includes the same file on Windows NT,  

2 // but not on UNIX  

3 #include <Iostream>    //not correct  

4 #include <iostream>     //OK
```

1310 • **Do not make assumptions about the size or layout in memory of an**  
 1311 **object.** [no-memory-layout-assumptions]

1312 The sizes of built-in types are different in different environment. For ex-  
 1313 ample, an int may be 16, 32, or even 64 bits long. The layout of objects is  
 1314 also different in different environments, so it is unwise to make any kind of  
 1315 assumption about the layout in memory of objects.

1316 If you need integers of a specific size, you can use the definitions from  
 1317 <cstdint>:

```
1 #include <cstdint>  

2  

3 int16_t a;           // A 16-bit signed int  

4 uint8_t b;           // A 8-bit unsigned int  

5 int_fast16_t c;     // Fastest available signed int type  

6 // at least 16 bits wide.
```

1318 The C++ standard requires that class members declared with no intervening  
 1319 access control keywords (public, protected, private) be laid out in  
 1320 memory in the order in which they are declared in the class. However, if  
 1321 there is an access control keyword between two member declarations, their  
 1322 relative ordering in memory is unspecified. In any case, the compiler is free  
 1323 to insert arbitrary padding between members.

1324 • **Take machine precision into account in your conditional statements.**  
 1325 **Do not compare floats or doubles for equality.** [float-precision]

1326 Have a look at the `std::numeric_limits<T>` class, and make sure your  
 1327 code is not platform-dependent. In particular, take care when testing float-  
 1328 ing point values for equality. For example, it is better to use:

```

1 const double tolerance = 0.001;
2
3 ...
4
5 #include <cmath>
6
7 if (std::abs(value1 - value2) < tolerance) ...

```

1329 than

```

1 if ( value1 == value2 ) ...

```

1330 Also be aware that on 32-bit platforms, the result of inequality operations  
 1331 can change depending on compiler optimizations if the two values are very  
 1332 close. This can lead to problems if an STL sorting operation is based on this.  
 1333 A fix is to use the operations defined in `CxxUtils/fpcompare.h`.

- 1334 • **Do not depend on the order of evaluation of arguments to a function;  
 1335 in particular, never use the increment and decrement operators in  
 1336 function call arguments.** [\[order-of-evaluation\]](#)

1337 The order of evaluation of function arguments is not specified by the  
 1338 C++ standard, so the result of an expression like `foo(a++, vec(a))`  
 1339 is platform-dependent.

```

1 func(f1(), f2(), f3());
2 // f1 may be evaluated before f2 and f3,
3 // but don't depend on it!

```

1340 Beware in particular if you're using random numbers. The result of some-  
 1341 thing like

```

1 atan2 (static_cast<double>(rand()),
2 static_cast<double>(rand()));

```

1342 can change depending on how it's compiled.

- 1343 • **Do not use system calls if there is another possibility (e.g. the C++  
 1344 run time library).** [\[avoid-system-calls\]](#)

1345 For example, do not forget about non-Unix platforms.

- 1346 • **Prefer `int` / `unsigned int` and `double` types.** [preferred-types]

1347 The default type used for an integer value should be either `int` or `unsigned`  
 1348 `int`. Use other integer types (`short`, `long`, etc.) only if they are actually  
 1349 needed.

1350 For floating-point values, prefer using `double`, unless there is a need to save  
 1351 space and the additional precision of a `double` vs. `float` is not important.

- 1352 • **Do not call any code that is not in the release or is not in the list of  
 1353 allowed external software.** [no-new-externals]

1354 

## 4 Style

1355 This section concerns the style, as opposed to the functionality, of the code.

1356 

### 4.1 General aspects of style

- 1357 • **The `public`, `protected`, and `private` sections of a class must be  
 1358 declared in that order. Within each section, nested types (e.g. `enum`  
 1359 or `class`) must appear at the top.** [class-section-ordering]

1360 The public part should be most interesting to the user of the class, and  
 1361 should therefore come first. The private part should be of no interest to the  
 1362 user and should therefore be listed last in the class declaration.

```

1  class Path
2  {
3  public:
4      Path();
5      ~Path();
6
7  protected:
8      void draw();
9
10 private:
11     class Internal {
12         // Path::Internal declarations go here ...

```

```
13 } ;
14 }
```

- **Keep the ordering of methods in the header file and in the source files identical.** [\[method-ordering\]](#)

This makes it easier to go back and forth between the declarations and the definitions.

- **Statements should not exceed 100 characters (excluding leading spaces). If possible, break long statements up into multiple ones.** [\[long-statements\]](#)
- **Limit line length to 120 character positions (including white space and expanded tabs).** [\[long-lines\]](#)
- **Include meaningful dummy argument names in function declarations. Any dummy argument names used in function declarations must be the same as in the definition.** [\[dummy-argument-names\]](#)

Although they are not compulsory, dummy arguments make the class interface much easier to read and understand.

For example, the constructor below takes two Number arguments, but what are they?

```
1 class Point
2 {
3 public:
4     Point (Number, Number);
5 }
```

The following is clearer because the meaning of the parameters is given explicitly.

```
1 class Point
2 {
3 public:
4     Point (Number x, Number y);
5 }
```

1381 • **The code should be properly indented for readability reasons.**

1382 [indenting]

1383 The amount of indentation is hard to regulate. If a recommendation were  
1384 to be given then two to four spaces seem reasonable since it guides the eye  
1385 well, without running out of space in a line too soon. The important thing  
1386 is that if one is modifying someone else's code, the indentation style of the  
1387 original code should be adopted.

1388 It is strongly recommended to use an editor that automatically indents code  
1389 for you.

1390 Whatever style is used, if the structure of a function is not immediately  
1391 visually apparent, that should be a cue that that function is too complicated  
1392 and should probably be broken up into smaller functions.

1393 • **Do not use spaces in front of [] and to either side of . and ->.**

1394 [spaces]

```
1 a->foo() // Good
2 x[1] // Good
3 b . bar() // Bad
```

1395 Spacing in function calls is more a matter of taste. Several styles can be  
1396 distinguished. First, not using spaces around the parentheses (K&R, Linux  
1397 kernel):

```
1 foo()
2 foo(1)
3 foo(1, 2, 3)
```

1398 Second, always putting a space before the opening parenthesis (GNU):

```
1 foo ()
2 foo (1)
3 foo (1, 2, 3)
```

1399 Third, putting a space before the opening parenthesis unless there are no  
1400 arguments.

```
1 foo()
```

```

2  foo (1)
3  foo (1, 2, 3)

```

1401 Fourth, putting spaces around the argument list:

```

1  foo()
2  foo( 1 )
3  foo( 1, 2, 3 )

```

1402 In any case, if there are multiple arguments, they should have a space  
 1403 between them, as above. A parenthesis following a C++ control keyword  
 1404 with as if, for, while, and switch should always have a space before it.

- 1405 • **Keep the style of each file consistent within itself.** [\[style-consistency\]](#)

1406 Although standard appearance among ATLAS source files is desirable, when  
 1407 you modify a file, code in the style that already exists in that file. This means,  
 1408 leave things as you find them. Do not take a non-compliant file and adjust a  
 1409 portion of it that you work on. Either fix the whole thing, or code to match.

- 1410 • **Prefer using to typedef.** [\[prefer-using\]](#)

1411 To declare a type alias, prefer the newer using syntax:

```

1  using Int_t = int;

```

1412 to the typedef syntax:

```

1  typedef int Int_t;

```

1413 The using syntax makes it clearer what is being defined; it can also be  
 1414 used to declare templated aliases.

1415 **4.2 Comments**

- 1416 • **Use Doxygen style comments before class/method/data member  
 1417 declarations. Use “//” for comments in method bodies.** [\[doxygen-  
 1418 comments\]](#)

1419 ATLAS has adopted the Doxygen code documentation tool, which requires  
 1420 a specific format for comments. Doxygen comments either be in a block

1421 delimited by `/** */` or in lines starting with `///`. We recommend using  
 1422 the first form for files, classes, and functions/methods, and the second for  
 1423 data members.

```

1  /**
2  * @file MyPackage/MyClusterer.h
3  * @author J. R. Programmer
4  * @date April 2014
5  * @brief Tool to cluster particles.
6  */
7
8 #ifndef MYPACKAGE_MYCLUSTERER_H
9 #define MYPACKAGE_MYCLUSTERER_H
10
11
12 #include "MyPackage/ClusterContainer.h"
13 #include "xAODBase/IParticleContainer.h"
14 #include "AthenaBaseComps/AthAlgTool.h"
15
16
17 namespace MyStuff {
18
19
20 /**
21 * @brief Tool to cluster particles.
22 *
23 * This tool forms clusters using the method
24 * described in ...
25 */
26 class MyClusterer
27 {
28 public:
29     ...
30
31 /**
32 * @brief Cluster particles.
33 * @param particles List of particles to cluster.
34 * @param[out] clusters Resulting cluster list.

```

```

35     *
36     * Some additional description can go here.
37     */
38
39     StatusCode
40     cluster (const xAOD::IParticleContainer& particles,
41               ClusterContainer& clusters) const;
42
43     . . .
44
45     private:
46     /// Property: Cluster size.
47     float m_clusterSize;
48
49     . . .
50
51 }
52 // namespace MyStuff
53
54
55 #endif // MYPACKAGE_MYCLUSTERER_H

```

1424 See the ATLAS Doxygen page [\[18\]](#).

1425 Remember that the `/* */` style of comment does not nest. If you want to  
 1426 comment out a block of code, using `#if 0 / #endif` is safer than using  
 1427 comments.

1428 • **All comments should be written in complete (short and expressive) English sentences.** [\[english-comments\]](#)

1430 The quality of the comments is an important factor for the understanding  
 1431 of the code. Please do fix typos, misspellings, grammar errors, and the like  
 1432 in comments when you see them.

1433 • **In the header file, provide a comment describing the use of a declared  
 1434 function and attributes, if this is not completely obvious from its  
 1435 name.** [\[comment-functions\]](#)

1 **class** Point

```

2  {
3  public:
4  /**
5   * @brief Return the perpendicular distance
6   *          of the point from Line @c 1.
7   */
8  Number distance (Line 1);
9  } ;

```

1436 The comment includes the fact that it is the perpendicular distance.

## 1437 5 Changes

### 1438 5.1 Version 2.1 (Jan 1, 2026)

- 1439 Migrated source to pandoc-markdown. Produce mkdocs-compatible output.  
1440 Minor edits.

### 1441 5.2 Version 2.0 (March 6, 2024)

- 1442 Updated for C++20.
  - 1443 – Don't use modules or coroutines.
  - 1444 – Add recommendation to use `<numbers>`.
  - 1445 – Suggest using `auto` to move the return type to the end of a method signature when returning types defined within the class.
  - 1446 – Suggest not defining template functions without the `template` keyword.
  - 1447 – Recommend `std::format` for formatted output.
  - 1448 – Note that range-for can have init-statements.
  - 1449 – Mention `std::bit_cast`.
  - 1450 – Recommend using instead of `typedef`. Rephrase previous references to `typedef`.
  - 1451 – Comparisons should be defined in terms of `operator==` and `operator<=`.
  - 1452 – Mention `std::span`.
- 1453 Some additional references.
- 1454 Clarify that non-ASCII characters should not be used in identifier names.

1459 • Clarify that variable-length argument lists of variadic template functions  
1460 are OK.

1461 **5.3 Version 0.7 (Sep 18, 2019)**

1462 • Minor cleanups and updates to take into account that we now require  
1463 C++17.  
1464 • Use the `fallthrough` attribute, not a comment.  
1465 • Allow omitting the `default` clause in a `switch` statement on an `enum`  
1466 that handles all possible values. Recent compilers will warn if some values  
1467 are not handled, and it's better to get such a diagnostic at compile-time  
1468 rather than at runtime.  
1469 • Clarify `avoid-typedef` section.  
1470 • Mention preference for `ATH_MSG_` macros.  
1471 • Don't require `override` for destructors.  
1472 • Avoid using `#pragma once`.

1473 **5.4 Version 0.6 (Dec 20, 2017)**

1474 • The `register` keyword is an error in C++17.  
1475 • Dynamic exception specifications are errors in C++17.  
1476 • Exceptions should be caught using `const` references, not by value.  
1477 • Discourage using protected data.

1478 **5.5 Version 0.5 (Nov 21, 2017)**

1479 • Add an initial set of guidelines for AthenaMT.  
1480 • Add recommendation to prefer range-based `for`.

1481 **5.6 Version 0.4 (Nov 16, 2017)**

1482 • Minor updates: we're now using `c++14`. Add note about implicit `fallthrough`  
1483 warnings with `gcc7`. Add rule to use `std::abs()`.

1484 **5.7 Version 0.3 (Aug 23, 2017)**

1485 • Add recommendation to avoid bit fields.

1486 **5.8 Version 0.2 (Aug 9, 2017)**

1487 • Small typo fixes.  
1488 • Add a brief description of pointer aliasing.  
1489 • Add more details about argument passing to functions.  
1490 • Add recommendation on `auto`.

1491 **References**

1492 [1] D. Knuth, [Literate programming](#), The Computer Journal **27**, 97 (1984).  
1493 [2] ATLAS Quality Control Group, [ATLAS C++ Coding Standard](#), ATL-SOFT-2002-001, 2001.  
1494 [3] CERN Project Support Team, C++ Coding Standard, CERN-UCO/1999/207, 2000.  
1495 [4] Standard for the Programming Language C++, n4868, 2020.  
1496 [5] C++ reference, (n.d.).  
1497 [6] News, Status & Discussion about Standard C++, (n.d.).  
1498 [7] C++ Stories, (n.d.).  
1499 [8] R. Grimm, [Modernes C++](#), (n.d.).  
1500 [9] H. Sutter, [Guru of the week archive](#), (2008).  
1501 [10] H. Sutter, [Guru of the week archive](#), (2013).  
1502 [11] S. Meyers, *Effective C++*, 3rd Edition (Addison-Wesley, 2005).  
1503 [12] S. Meyers, *Effective STL* (Addison-Wesley, 2001).  
1504 [13] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns, Elements of Reusable Object-Oriented Software* (Addison-Wesley, 1994).  
1505 [14] E. Niebler, [Range library for C++](#), (n.d.).  
1506 [15] H. Sutter, [A Pragmatic Look at Exception Specifications](#), C++ Users Journal **20**, (2002).  
1507 [16] A. Krzemieński, [noexcept – what for?](#), (2014).  
1508 [17] [FaqCompileTimeWarnings ATLAS wiki page](#), (n.d.).  
1509 [18] [DoxygenDocumentation ATLAS wiki page](#), (n.d.).