

Doc. no.: P0939r0
Date: 2018-02-10
Programming Language C++
Audience: All WG21
Reply to: Bjarne Stroustrup (bs@ms.com)

Direction for ISO C++

B. Dawes, H. Hinnant, B. Stroustrup, D. Vandevorde, M. Wong

Revision History

- This is the initial version.

Main sections

- [History](#)
- [Long-term Aims \(decades\)](#)
- [Medium-term Aims \(3-10 years\)](#)
- [Priorities for C++20](#)
- [Process Issues](#)
- [The C++ Programmer's Bill of Rights](#)

Abstract

This is intended as a document updated as times go by and new facts and opinions emerge. It tries to articulate our aims for C++, the challenges faced, and some concrete suggestions for priorities.

This document differs from most papers by aiming for a global view and placing individual features in context, trying to avoid delving into technical details of individual proposals except where that “details” could affect the language as a whole or its tool environments.

To prevent this document from becoming toothless “pure philosophy,” we propose giving priority to specific areas of development and to specific proposals in those areas.

Suggested improvements are welcome.

1. [History](#)

The Direction Group (DWG) was created in response to calls for more thought about the direction of C++'s evolution (language and standard library) and comments about our self-established processes in WG21. The specific “call to action” was [Operating principles for C++](#) [Winkel,2017] by several Heads of National Delegations.

We see C++ in danger of losing coherency due to proposals based on differing and sometimes mutually contradictory design philosophies and differing stylistic tastes. For that reason, we recommend that you (re)read [Winkel,2017] before proposing a new feature (language or library).

[Winkel,2017] quoted heavily from D&E. However, during the presentation of [C++ Stability, Velocity, and Deployment Plans](#) [Winters,2017] at the June 2017 Toronto meeting to an almost complete WG21, Titus Winters asked for a show of hands of who had read “The Design and Evolution of C++” [Stroustrup,1994]. Only about a quarter of the hands went up. Assuming that another quarter was shy or distracted, that indicates that half of the committee have never read the articulation of key design principles and decisions for C++. Part of the reason for that is that copies of D&E can be hard to find. Consequently, we hope to be able to give every member of the committee a copy (on paper or electronically depending on what Addison Wesley agrees to).

We strongly recommend that someone who wants to push or oppose a proposal read [Stroustrup,1994] and its related HOPL papers ([Stroustrup,1993] and [Stroustrup,2007]). Our impression is that many members of WG21 have an inaccurate view of how and why C++ succeeded and some are pushing popular approaches that in the past have led to failure (e.g., by languages following those approaches and losing out to C++). Obviously aims, techniques, and ideals change over decades, progress happens, but it is dangerous to operate in ignorance of past successes and failures.

2. The Direction Group

The direction group’s (DWG) charter is [Sutter,2018]:

Direction group. The direction group is a small by-invitation group of experienced participants who are asked to recommend priorities for WG21. Currently, that group consists of: Beman Dawes, Howard Hinnant, Bjarne Stroustrup, Daveed Vandevoorde, and Michael Wong. Their charter includes setting forth a group opinion on:

- Evolution direction (language and library): This includes both language and library topics, and includes both proposals in hand and proposals we do not have but should solicit. The direction group maintains a list of the proposals it considers the most important for the next version of C++ or to otherwise make progress such as in a TS, and the design group chairs use that list to prioritize work at meetings. Typically, work on other topics will occur after there’s nothing further left to do at this meeting to advance the listed items.
- Providing an opinion on any specific proposal: This includes whether the proposal should be pursued or not, and if pursued any changes that should be considered. Design group participants are strongly encouraged to give weight to an opinion that the design group feel strongly enough about to suggest.

We are concerned with both the long-term and the shorter-term direction and handling of proposals, especially where they may intersect multiple WGs. When looking at direction and at individual proposals, we try to consider the interests of the larger C++ community, rather than just the narrower interests of WG21 members.

The direction group mailing list is direction@lists.isocpp.org. Opinions welcome.

We operate similar to a Board of Directors, with an annually rotating chairman. Bjarne lost a random number draw (written by Howard) and is therefore our chair for the first year:

- B. Stroustrup, (2018)
- H. Hinnant, (2019)
- B. Dawes, (2020)
- M. Wong (2021)
- D. Vandevoorde, (2022)

3. [Long-term Aims \(decades\)](#)

Assuming we succeed, many of us will be writing and maintaining C++ in 10 years' time and 20 years' time. What kind of aims would we think reasonable for such a time scale? Obviously, long-time aims cannot change every year.

As a community, we want lots of things for C++, but

- Some we can't get because they are impossible.
- Some we can't get because they are incompatible with something else we have or want to have.
- Some we don't yet know how to do.
- Some we don't have the resources to do.
- Some we don't agree to be good for C++.
- Some we don't agree to be worthwhile for C++.

We understand that we can't have everything we want and that some of the things we want take a lot of time and work. Given that, we need to clarify our current long-term goals.

No language can be everything for everybody. Trying to achieve that has killed many efforts. We have to decide what C++ is supposed to be and become.

What is C++?

- C++ is a language for defining and using light-weight abstractions
- C++ supports building resource constrained applications and software infrastructure
- C++ support large-scale software development

How do we want C++ to develop?

- Improve support for large-scale dependable software
- Improve support for high-level concurrency models
- Simplify language use
- Address major sources of dissatisfaction
- Address major sources of errors

We fundamentally need:

- Stability: Useful code "lives" for decades

- Evolution: The world changes and C++ must change to face new challenges.

There is an inherent tension here. We should be very reluctant to break compatibility. People always want a simpler language, a language without complicating “legacy features”, and get seriously angry if we break code that they depend on.

C++ is complicated, too complicated, yet we cannot remove significant facilities and changing them is very hard. Changing parts deemed insignificant can be risky (we need better analysis tools [Winter,2016]) and the potential gains would be insignificant. However, we badly need to simplify use of C++ and that leaves us with three alternatives:

- Provide a simpler alternative for simple uses
- Provide simplifying generalization
- Provide alternatives to complicated and/or error-prone features

Often, a significant improvement involves a combination of those three.

Calling C++ a “multi-paradigm language” is no excuse for adding support for incompatible programming styles. If features don’t interoperate in combination, we get de facto dialects. Today, some of the most powerful design techniques combine aspects of traditional object-oriented programming, aspects of generic programming, aspects of functional programming, and some traditional imperative techniques. Such combinations, rather than theoretical purity, are the ideal.

- Provide features that are coherent in style (syntax and semantics) and style of use

This applies to libraries, to language features, and combinations of the two.

Technically, C++ rests on two pillars:

- A direct map to hardware (initially from C)
- Zero-overhead abstraction in production code (initially from Simula, where it wasn’t zero-overhead)

Depart from those and the language is no longer C++. We should not depart from these principles and fall into the trap of:

- Abandoning the past which can seriously jeopardize compatibility. C++ is and will continue to be heavily used in long-lasting systems.
- Not addressing new challenges such as higher-level concurrency models. Should we fail, developers will need to switch to some other framework to gain the best performance.

Over the long term, we must strengthen those two pillars

- Better support for modern hardware (e.g., concurrency, GPUs, FPGAs, NUMA architectures, distributed systems, new memory systems)
- More expressive, simpler, and safer abstraction mechanisms (without added overhead)

Many useful abstractions should find their way into the standard library

- In principle, the C++ standard library can be implemented in C++ plus a few “intrinsic operations” for accessing low level machine facilities.

This is a variant of the “Don’t leave room for a language below C++ (except assembler)” rule of thumb from D&E. Where we depart from this principle, we get long-term problems because the semantics easily drifts apart from the rest of the language.

C++ relies critically on static type safety for expressiveness, performance, and safety. The ideal is

- Complete type-safety and resource-safety (no memory corruption and no leaks)

This is achievable without added overhead, in particular without adding a garbage collector, and without restricting expressibility (see [A brief introduction to C++'s model for type- and resource-safety](#) [Stroustrup,2015b] and [C++ Core Guidelines](#) [CG,2015]). We don’t consider a program littered with casts type-safe and we recognize that the ideal of complete type-safety and complete resource-safety cannot be achieved while accepting every legal C++ program (as we must for compatibility), so external tools will be needed (e.g., static analysis of lifetimes), but improved language and library support can help a lot.

C++ is now used in more domains than ever. We do not specifically recommend any specific domains, as every domain is important to somebody in the larger C++ community – even if their domain isn’t well-represented in the committee. But we do intend to broaden our support for domains that are well-represented by the C++ community, even non-traditional ones, as a high-level aim. However here are some areas of general concerns – often cutting across application domains – that we should not ignore:

- Safety and security: reduce vulnerability to intrusion and the ability to exploit intrusions (e.g., prevent unsafe input and eliminate type violations from misuse of free store)
- Simplification: make simple things simple to do (and thereby making C++ easier to learn).
- Interoperability: improve interoperability with important languages and systems (e.g., Python bindings). Modern C++ facilitates that with Standard Layout type and POD types.
- Support for demanding applications in important application areas, such as medical, finance, automotive, and games (e.g., key libraries, such as JSON, flat containers, and pool and stack allocators).
- Embedded systems: make C++ easier to use and more effective for large and small embedded systems, (that does not mean just simplified C-like C++; e.g., see [Quora1][Quora2][Saks,2016]).
- Alternatives for error-prone and unsafe facilities (like **std::variant** as an alternative to unions or pattern matching as an alternative for **std::variant**)

This is not an exhaustive list, but demonstrates high level aims that brackets our priority. The medium-term aims will bracket our priorities with more specific proposals that are currently in flight. Many of these goals cannot be met through changes in the standard alone. Some, such as bindings and avoidance of error-prone facilities, require improvements beyond the standard, but the standard can make such improvements easier.

We would like to see the software developments tools used for C++ (such as compilers, static analyzers, refactoring tools) significantly improved. Most of this is beyond the scope of WG21, but we should try hard not to make things more difficult, e.g. by significantly increasing compile times or the cost of tool building by adding significant complexity barriers (e.g., by encouraging coding styles increasing the use of macros or brittle SFINAE).

Conversely, to improve the predictability of our processes and maximize the likelihood of delivering significant improvements, we discourage

- Isolated “cute” proposals.
- “Change the World” papers for proposals already in flight

We change the language and standard library by gradually building on previous work or by providing a better alternative to an existing feature. Proposals to completely change direction for a proposal already in process should be treated with suspicion and subjected to at least as much scrutiny as the proposal already in flight has been. Looking at new proposals is more exciting than working out the obscure details of an old proposal, but eventually all successful proposals go through the “polish the last details” stage, so we can’t escape that through new and shiny proposals, just postpone the pain. A proposal for a radical change to a proposal already “in flight” (i.e., in its second or later discussion) should not be allowed to delay the latter unless it comes with a paper with a detailed discussion of design, use, and implementation.

4. Medium-term Aims (3-10 years)

Naturally, whatever we focus on in the medium term should reflect the aims, principles, and problems identified for the long term. What can we do to get closer to those ideals over the next few years? What are promising areas of development? What is urgent?

4.1. Learning and Teaching C++

C++ has a major problem with learning and teaching. This could send the C++ community into decline over the next few years. C++ is expert-friendly, but it cannot be just expert-friendly without losing important streams of new talent. This was the case before C++11 and C++11 reversed that negative trend. Currently, the C++ community is more active than it ever was, but this improvement must be sustained with solid improvements in the language and the standard library. It takes more still to sustain the enthusiasm, but here we are concerned with what WG21 can do. As individuals, we can encourage better tools and libraries. We suggest

- Establish a “Learning and Teaching” SG

We say “Learning and Teaching” because many (most?) new C++ programmers learn on their own. We much help those. To paraphrase D&E: C++ needs to be like a weed, not like a garden rose, in order to flourish. We can’t assume a reliable support system.

C++ is seriously underrepresented in academia and often very poorly taught. It has been conventional to start teaching C++ by first introducing the lowest level and most error-prone facilities. Naturally, that discourages students and increases the time needed to get to what students consider meaningful computing (graphics, networking, mathematics, data analysis, etc.). Often, teachers even go to the extreme of insisting on using a C compiler. If the ultimate aim is to teach C++, that’s like insisting people start learning English by reading Beowulf or the Canterbury Tales in their original early-English language versions. Those are great books, but Early English is incomprehensible to most native Modern English speakers and in addition to the linguistic difficulties, they present cultural conventions and idioms that

seem very peculiar today. Instead of C, maybe someone should teach Simula to prepare for learning C++? Why don't people do that? Because the historical approach to teaching language (natural or programming language) complicates and detracts from the end goal: good code.

Why do teachers use the C-first approach to teach C++? Part is tradition, curriculum inertia, and ignorance, but part of the reason is that C++ doesn't offer a smooth path to idiomatic, proper, modern use of C++. It is hard to bypass both the traps of low-level constructs and the complexities of advanced features and teach programming and proper C++ usage from the start.

C++ teaching is mostly stuck in a pre-graphics, pre-web world. This is a serious weakness, potentially fatal over the longer term. There are beginners' books trying to alleviate that (e.g., [Programming -- Principles and Practice Using C++ \(Second Edition\)](#) [Stroustrup,2014]) but they are a small minority and only partially successful. Teaching C++ to modern students will not be really effective until we can offer

- Simple standard graphics and simple use of browsers (or GUI).
- Simple mechanisms for packaging, distribution, and installation for libraries and programs

The latter is beyond the current scope of WG21, but it is essential for the continuing success of C++; maybe the new Tools SG can help. If WG21 cannot do something official, maybe its members can help establish widely accepted de facto standards (note that languages and systems competing with C++ tend not to have formal standards).

What in the language and libraries helps learning and teaching?

- The ability to start teaching using only what is supplied by every implementation.

Students cannot be expected to download and install libraries on day #1. Thus, the technique of providing a training-wheels library in the form of a header or precompiled libraries is at most second best. This is a major technical reason that courses start with pointers, arrays, C-style strings, unions, etc., rather than **std::string**, **std::vector**, **std::variant**, and algorithms.

- The ability to sequence teaching from the simple and fundamental to the advanced.

The snag is that the simple and fundamental typically involves abstractions. This implies that

- It should be trivial to include what is needed for teaching; modules should help here and we need an "all" module and/or a "foundation" module (see also [P0581R1]) so that novices don't have to search a textbook or the web to find out which header to **#include** to use a foundational feature.
- Errors coming from misuse of standard library facilities should be comprehensible; concepts should help here. All templates in the standard library intended for non-expert use should be "conceptualized."
- There should be support for preventing use of non-approved and or yet-to-be-taught features; this cannot be standardized, but tools for catching counterproductive styles should be encouraged; the Core Guidelines [CG,2015] is a start.

Note that this is effectively what teachers who restrict students to C do, with very undesirable effects for C++ programming.

4.2. Library and Language

During the early stages of the development of the first standard, we articulated some principle for what to include in the standard library, including

- Language support
- Facilities that everybody needs
- Facilities needed for communicating among separately developed libraries

We think these are reasonable criteria, but in practice they didn't have much effect. People worked on and overelaborated individual components (e.g., `std::string`) that didn't interoperate effectively, showed major stylistic differences (reflecting differing design conventions), and often didn't completely cover what programmers considered a topic. The results were complaints, incomplete adoption, and wide use of commercial alternatives. We were rescued from disaster by Alex Stepanov. The STL was a coherent design that was sufficiently complete for its domain to become widely adopted.

If we can supply a feature as a library, we should do so because it is easier to validate a design through experimentation, the facility will be available to users earlier, and a library is usually easier to specify in isolation. However, a library design should not be an excuse for inelegant interfaces, for irregular interfaces, for stylistic differences from built-in language features, or overelaboration. It is always easy to add another function to a class, so library components have a tendency to bloat. Note

- the trouble we have had from `std::vector` using unsigned subscripts while built-in arrays use signed indices (unsigned arithmetic is modular, leading to correctness, performance, and error-detection problems)
- the dramatic differences in the interfaces to `std::any`, `std::optional`, and `std::variant`

At the time, there were of course reasons for those design choices that looked good to many, but we must aim for coherence:

- among standard-library features
- between the standard-library features and the built-in language features

Or, stated differently:

- No feature should be added without someone explicitly considering its interaction with other features (language and library).

We suggest that WGs (in particular, EWG and LEWG) delegate a "consistency review" to a small ad hoc subgroup for each feature or set of related features to be done concurrently with advanced work on technical details.

4.3. Concrete suggestions

Here are some more concrete suggestions for what can be done over the next few years to bring C++ closer to its long-term ideals:

- **Pattern matching:** List comprehension, and not just for algebraic types. This could simplify code and strengthen type safety. It would eliminate a lot of ad-hoc selection, eliminate unsafe use of

unions, low-level use of variants, and eliminate the visitor pattern (which is an expensive workaround that confuses many).

- **Exception and error returns:** There are still people who argue against all use of exceptions and people who claim that exceptions should be used consistently instead of error codes. We'd like to see a clear articulation of where each alternative is ideal, but until we do that we should be careful about what error-handling strategies we support. The contracts proposal ([A Contract Design](#) [GDR,2016]) should simplify this by taking care of many nasty cases.
- **Static reflection:** C++'s lack of run-time reflection is an essential strength in some areas (such as small embedded systems), but lack of simple reflection in specific areas have let many to switch to alternative languages. In particular, the lack of ability to generate serializer/deserializer pairs and object layout maps has been painful in many areas.
- **Modern networking:** We need the networking as specified in the networking TS ([Wakely,2017]) urgently needed (on the C++20 timescale), but beyond that we need good simple support for interaction with modern browsers and networking (http, html5).
- **Modern hardware:** We need better support for modern hardware, such as executors/execution context, affinity support in C++ leading to heterogeneous/distributed computing support, SIMD/task blocks, more concurrency data structures, improved atomics/memory model/lock-free data structures support. The challenge is to turn this (incomplete) laundry list into a coherent set of facilities and to introduce them in a manner that leaves each new standard with a coherent subset of our ideal.
- **Simple graphics and interaction:** Command-line options and console I/O, while still useful in various contexts, haven't been the primary way of interacting with software for decades. Although modern interactive applications deal with broadly varying devices, almost all are controlled using an "event-driven" software model and most involve a graphical screen. We think the C++ standard library needs a component providing (a) a graphical canvas with basic geometric drawing, text rendering, and bitmap output primitives, (b) a simple input API for that graphical canvas abstract enough to deal with both "click" and "touch" events, and (c) a simple, highly-abstract "menus and dialogs" system (again, abstract enough to not force surfacing distinctions between devices as disparate as "phones" and "desktop computers"). Such a component will, of course, be "optional" to allow for platforms that don't have the needed I/O facilities. Practical uses of this component should be simple enough to be teachable to first-year students. Improved Unicode support.
- **Anything from the Priorities for C++20 that didn't make C++20:** see [Priorities for C++20](#).

We feel that work on a database interface is badly needed, but since there seem to be no current work on that in the committee and no critical mass of expertise and interest, we can't put it on our list.

5. Priorities for C++20

We are very sympathetic to Ville Voutilainen's proposal [To boldly suggest an overall plan for C++20](#) [Voutilainen,2017] to restrict new features to a limited set of significant proposals. Our list differs in detail:

- Concepts (a large part of the TS is already in the WP)
- Modules (offering proper modularity and dramatic compile-time improvements)

- Ranges (incl., some of the infinite sequence extensions to the TS)
- Networking ([Extensions for Networking](#) [Wakely,2017])
- Concepts in the standard library

Beyond that, if time allows while giving priority to the top issues for committee time, we recommend:

- An (optional) Graphics API TS (see [A Proposal to Add 2D Graphics Rendering and Display to C++](#) [McLaughlin,2017]).
- Simple Reflective Metaprogramming TS (relying on “ordinary C++” rather than template tricks).
- Contracts as specified in [GDR,2016], [GDR,2017].
- Calendar functions and time zones ([Extending <chrono> to Calendars and Time Zones](#); [Hinnant,2017]).
- Coroutines [Nishanov,2017].
- A high-level (and simple-to-use) API for common modern networking and browser use (HTML in particular and including Unicode support) TS (an example in current use: [WT](#))

Additions beyond that should be discouraged as time sinks and distractions. Proposals for minor features should be given priority if and only if they support the priority items. The last point is only feasible if someone can point to an existing system that could be turned into a proposal.

A significant part of the C++ community is panicked by fears of massive growth of features and complexity. This leads to disuse of novel features. The complexity of C++ and the difficulty of learning it are major barriers to increased use. No feature is cost free, there is always the implementation cost, the cost of producing teaching material, the time needed to learn, the opportunities for confusion, and the inevitable distraction from overselling.

We chose to prioritize these particular proposals because they

- relate directly to the long-term and medium-term aims
- are advanced with implementations and have user experience
- have people dedicated to their completion and improvement

When approaching these proposals try to make sure that a proposal

- actually offers the expected benefits
- handles the primary use cases simply and elegantly
- isn't delayed by a futile insistence of theoretical perfection
- isn't complicated to please every possible expert need

Once we know the approximate final goal, we should

- Use an incremental approach to design to benefit from actual experience

Our nightmares include a module system that doesn't offer modularity (e.g., where the order of import is significant or the compile-time cost of using modules approaches that of #includes) and a reflection system that is so integrated with general metaprogramming that nothing can be delivered until C++23 or later and significantly increases compile-times. Note that the cost of compilation is among the loudest reasonable complaints about C++ from its users.

6. Process Issues

We are “a bunch of volunteers.”

- Most are enthusiasts for some aspect of the language or other.
- Few have a global view (geographically, C++ community, C++ usage, C++ Language and standard library).
- Most have a strong interest in some subset of use, language, or library.
- Most are deeply engaged with a specific form of C++ as part of their daily work.
- Our levels and kinds of relevant computer-science, design, and programming language education vary dramatically.
- Many are clever people attracted to clever solutions.
- Some are devoted to ideas of perfection.

This implies that we can't rely on a common vocabulary, a core set of shared values, a common set of basic ideals, or a common understanding of what's a problem. Consequently, we must

- Spend more effort articulating rationales for proposals.
- Spend more effort providing facilities for the “average programmer,” who is seriously underrepresented on the committee.

Please pay special attention to that last point. We feel that C++'s utility and reputation suffer badly from the committee lacking attention to improvements for relative novices and developers with relatively mundane requirements. Remember:

- Most C++ programmers are not like the members of the committee

We, as a committee, have no mechanism of reward (except accepting someone's proposal) or punishment (except delaying or rejecting someone's proposal). To get something accepted, we need consensus (defined as large majority, but not necessarily unanimity). This has implications on what we can do and how we do it.

- Nothing gets done unless someone cares enough to do it
- A small vocal minority can stop any proposal at any stage of the process.

Currently, C++ is probably as popular as it has ever been and there are definitely more active members of the standards committees than ever before. One effect of this enthusiasm for C++ is to inundate us with a flood of proposals. The sheer volume of proposals leads to fewer proposals getting through the processes to get accepted. Some proposals become warped from the need to gain support in an environment where time to think and present ideas is limited. Many of these proposals come from people who are unacquainted with standardization. We understand that some new members find it hard to accept that

- Progress is less rapid than for their corporate and open source projects.
- Concerns that are essential to them are not given priority by the committee.
- Committee members don't all understand or accept their experience and design principles.
- Standardizing for decades differs dramatically from shipping the next release with the fewest number of bugs.

- There are millions of programmers who use C++ in ways that differ from what they consider normal and reasonable.
- “No bugs” does not imply that a proposal is good.
- “There are bugs” does not imply that a proposal is bad (there is a saying among mathematicians that the way to recognize an important result is by the number of errors in its initial proof – it is relatively simple to provide a flawless proof for a trivial result).
- “Good enough” isn’t always good enough because a second look at the problem might come up with a better, more general, or better integrated solution. We have to consider the long term.
- 90% or more of the work of getting a proposal into the standard is ensuring that it fits smoothly with other facilities (language and standard library).
- Stability/compatibility is an essential feature.

And still some improvements are urgent. However, in the context of the committee “urgent” still implies “years” of work/delay. Novices are not the only ones who are impatient, but we must try to channel our energies into constructive activities. We encourage members (old and new)

- to get acquainted with [C++’s history and design rules](#)
- to accept or contribute to [our long-term aims for C++](#)

The aim of most members, new or “vintage” is to improve C++ by having their favorite proposal accepted. However, the sheer number of proposals (constantly on the order of a hundred) and the number of papers (running at more than one hundred for each meeting), not even every **good** proposal can be accepted. We encourage everyone to ask themselves

- Is my proposal among the top-20 in terms of long-term benefit to the C++ community as a whole?
- Is my proposal essential for some important application domain?

Push only very gently if the answers are “no,” “sort of,” or “I don’t know” and redirect effort to proposals for which the answer is “Yes!” We encourage the WG chairs to

- Focus on the major high-level and intermediate-level goals
- Articulate the goals as they relate to their WG
- Prioritize proposals based on those goals
- Discourage proposals not in that scope
- Discourage proposals from being re-submitted with only minor changes after rejection (especially if the revised proposal does not include new insights into the problem to be solved).

We are a set of interrelated committees with about 130 members present at a meeting and more active via the Web. Thus some “design by committee”, or rather “design by committees,” is unavoidable. We need to consciously and systematically try to minimize those effects by building a shared sense of direction.

- We have no shared aims, no shared taste.

This is a major problem, possibly the most dangerous problem we face as a committee. For C++ to succeed, we must overcome that. For starters, we – as individuals, as SGs, as WGs, and as the committee as a whole – must devote more time to develop and articulate common understanding.

- The alternative is a dysfunctional committee producing an incoherent language.

We need to be more explicit about

- What general problems we are trying to address
- How a particular proposal serves those articulated aims

It would also be a great improvement if members didn't spend all of their time in a single WG. Doing that leads to lack of insight and of trust.

- Try to spend at least one day each meeting in a WG that isn't "your own"
- Try to read several papers from each mailing that is not aimed at "your WG"

Note that if "your" proposal progressed, you'll have to shepherd it through different WGs, so it is good to have some understanding of how they operate.

6.1. Trust

At the start of the formal standardization efforts (ANSI and ISO), P.J. Plauger and Tom Plum emphasized the need for trust in the process. The desired degree of trust is currently missing, and the consequences are dire.

A WG can spend meetings on a proposal, carefully crafting a compromise, balancing concerns, just to have the next WG spend several meetings reviewing the proposal, demanding changes, and rephrasing the wording with patchy understanding of the design rationale as it evolved over time. In particular

- large feature developments should be discussed in joint EWG+CWG, LEWG+LWG, and/or EWG+LEWG sessions as early as the overall design and the understanding of the underlying implementation alternatives allows.

Finally, after years of process, someone then stands up in full committee and raises issues that have been discussed for years stating "lack of comfort" with the proposal, suggesting alternative approaches, and demanding more time to consider or rejection. At this point, everybody unhappy with compromises made along the way chirps in with counter-points made over the years and the proposal is either withdrawn or defeated by a 20% minority, many of whom did not take part in previous discussions. We think that

- "lack of comfort" is not sufficient to block a proposal

Even when (as is common) there are minor remaining issues, typically, there are many months between a proposal being approved (by any part of the committee) and the final standard during which non-fatal flaws can be fixed. We think that a proposal that has passed a committee (say EWG or LEWG) should be accepted without undue delays by further groups. Obviously, new inputs should be considered, but in the absence of new information, "we considered and discussed that" by a WG chair should be conclusive:

- People who were not sufficiently motivated to take part in the discussion or conclusively defeated in WG should feel obliged to at least stay neutral.

Unfortunately, this principle is easily gamed by people who break it to block proposals they don't like. This has been observed in the committee and has led to resentment and erosion of trust.

At a first approximation,

- No proposal is ever perfect.
- Different people have very different ideas of what degree of perfection is needed for a proposal to be part of the standard (or part of a TS).

“The last bug” is a common programmer’s joke and our language specification reads like a program in a poorly defined language (English) without advanced control structures, without abstraction mechanisms, and without a compiler. Consequently, perfection will always elude us and we should take that into account when we decide what is “good enough” for a standard (or a TS).

We have seen proposal move forward with an issues list to be considered before a final vote or even after. We strongly encourage that approach to make progress

- If a proposal is fundamentally sound, it should be moved forward, even into the WP text
- Refinement of text can (and should) happen later
- There are problems that are unlikely to be found until after the wording for a proposal is integrated into the WP text (as in software development: “integrate early to allow testing”)
- Addition of desired but incompletely developed features can wait until later, even to a later standard
- As the shipping date for a standard approaches, the efforts to “polish” the text and drain the issue lists should be redoubled at the expense of effort on new proposals (as is currently done)

Through delays, a good proposal can get dated, failing to benefit from years improvements and progress in the community. Note the file system TS and networking TSs each has a decade of use behind them. Similarly, **std::variant**, **std::optional**, and **std::any** have a long history as independent proposals. That wouldn't be too bad if the reason was that significant improvements were added during the process. However, the current process tends to reduce proposals to their most conservative cores. Thus, some “novel” C++ features feel dated by the time they are accepted.

- Aim for prompt delivery followed by incremental improvements

Note that this approach can only succeed if the end-goals are reasonably clear.

- Articulate the end goals for a proposal

When triaging features for consideration, we propose that the WG chairs to put a higher priority on features that conform to these goals, but also keep a watch against features that

- Turn C++ into a radically different language
- Turn parts of C++ into a much significantly different language by providing a segregated sub-language
- Have C++ compete with every other language by adding as many as possible of their features
- Incrementally modify C++ to support a whole new “paradigm”
- Hamper C++'s use for the most demanding systems programming tasks

- Increase the complexity of C++ use for the 99% for the benefit of the 1% (us and our best friends)

The committee members and proposal authors are a group of volunteers who may have different aims, industry views, or company directives. As we propose features, we urge the group to maintain cordial discourse, and aim for what is best for C++ as a language and not merely what is best for your company's current direction. This sometimes means willing collaboration with people with different aims and priorities. For the most part, we have been successful. More specifically, don't oppose a proposal just because:

- it is seen as competing with your favorite proposal for time/resources
- it is not relevant to your current job
- you have not personally reviewed it
- it is not perfect (according to your principles)
- it is not coming from your friends
- it is coming from someone you have been at odds with

6.2. Proposal processing

WG21 does not lack for proposals today. The interest is high and the number of proposals of each mailing grow to more than 100 in a mailing. That's unmanageable for an individual with a day job. To focused on the priorities discussed in this document, we recommend that chairs focus at least 60% of their time on advancing these priorities. How that is done is entirely the chairs prerogative. For example:

- Raise the barrier for repeated presentation of and voting on failed proposals.
- Require a short tutorial for each proposal

For any non-trivial proposal, it is not obvious how to use the new facility well in combination with other features. It is easy to say "teaching this is easy" but such statements should be backed-up with examples and experience report (or be clearly marked as conjecture). In particular, for whom is learning and using a new feature easy?

- Never assume that the use of a proposed feature is easy and obvious to everybody.

The volume of proposals is such that people are not able to track all upcoming votes. This has led to "no" votes in WGs and plenary. We ask the committee to consider process improvements that improve this communication and notification, but also ask that members actively review and discuss with others their proposal progress, so that the burden is fairly spread. At one time, Alistair Meredith did a great job of summarizing the flow and progress of each paper prior to C++11. Today this is unmanageable for one person. We have noted other forms of process improvement that has helped improved communication, and understanding the issues as we review papers. These include use of githubs, and summary notes recorded for each proposal, to save readers the bother of going through each meeting minutes to trace the history of a proposal and find the related papers. We do not necessarily enforce any particular methods, but feel some method should be considered to allow other groups to track progress.

Remember the need for trust (§6.1):

- “I didn’t have time to read the document” is not sufficient reason to oppose (provided the paper was submitted on time)
- “I don’t understand” is not by itself sufficient reason to oppose (give reasons and examples why your lack of understanding is not just your own problem)

6.3. The role of TSs

Technical Specifications (TSs) have become popular. In theory, they provide an intermediate state where major new features can be discussed, specified, tested under fewer constraints than for the standard itself, and later be moved into the standard with changes based on the experiences gained. The move from a TS to an IS is supposedly simplified by the significant work to complete the TS and the experience gained.

In practice, this seems to work reasonably well for libraries, though we see examples where facilities are stripped from a proposal as being too advanced, rather than pushed forward for experimentation. For language features, the experience is less positive. It seems that the barrier to entry into a TS is not significantly lower than for an IS, and the effort devoted to complete a TS detracts from experimentation, freezing the language feature in time years before acceptance. After the TS, most of the design and specification issues are then revisited a second time. Thus, a TS becomes a method for delaying a proposal. Also, a TS doesn’t seem to be sufficient to encourage multiple implementations. This implies that large parts of the C++ community doesn’t get to use the facility, large organizations using multiple compilers can’t experiment with the TS at scale, and tool builders hold back waiting for an “proper standard” supported by all major implementers. Consequently, detractors can dismiss the experience reports and clamor for novel alternatives.

We recommend

- Use TSs for library components.
- Don’t use TSs for a language feature unless the feature is a mostly self-contained unit.
- Never use a TS simply to delay; it doesn’t simplify later decision making.

If a proposal isn’t ready for the standard let it be improved or rejected in the appropriate WG or SG (Study Group). Don’t add a formal TS process.

6.4. “Details”

The committee spends most of its time on details that the average programmer will never notice. A lot of that is necessary. However, it seems that most members spend most of their time on such details; that’s wasteful and weakens the language as experienced by users. Instead, try to spend the majority of time on design issues, such as:

- What problems are this feature meant to solve? (D&E recommends not to accept a feature that doesn’t solve two apparently unrelated problems)
- What alternative solutions are possible?
- What related features should be considered simultaneously with this one?
- How well does this feature fit with the general style of the language and standard library?
- Is this feature aimed only for a few language or library experts or will it be used in application code?

An individual proposal is essentially never useful in isolation, it will work in combination with other features:

- Consider every feature in the context of features with which it will be used
- Don't approve a set of related features one by one, evaluate them together

In particular, cluster related proposals (language and library) together.

- Be fair when comparing alternative proposals

In particular, try hard to give discussions of advantages and disadvantages equal weight for each alternative. Presenting only advantages for a "favored proposal" and only "disadvantages" for an unfavored alternative is not acceptable.

The text of the standard should be precise and comprehensive. It is not supposed to be a tutorial, but the ideal is that after some "acclimatization" an experienced programmer should be able to interpret the text with some confidence. This is not currently the case. When crafting WP text

- Consider readability by people not members of the WG crafting the text
- Non-normative notes can be useful
- Short examples can be very useful
- Try not to change the meaning of terms common in the C++ community
- Where possible, use terms common in the C++ community

Remember that compiler writers and standard-library implementers are not the only target audience for the standard (or a TS).

7. The C++ Programmers' Bill of Rights

This note was posted to reflectors, presented at the June 2017 Toronto meeting and discussed. We propose that it be formally adopted by WG21:

"The C++ Programmers' Bill of Rights."

We, the ISO C++ Standards Committee, promise to deliver the following to the best of our ability, assuming user code adheres to the current standard:

1. **Compile-time stability:** Every significant change in behavior in a new version of the standard is detectable by a compiler for the previous version.
2. **Link-time stability:** ABI breakage is avoided except in rare cases, which will be well documented and supported by a written rationale.
3. **Compiler performance stability:** Changes will not imply significant added compile-time costs for existing code.
4. **Run-time Performance stability:** Changes will not imply significant added run-time costs to existing code.
5. **Progress:** Every revision of the standard will offer improved support for some significant programming activity or community.

6. **Simplicity**: Every revision of the standard will offer simplification for some significant programming activity.

7. **Timeliness**: Every revision of the standard will be shipped on time according to a published schedule.

Note “to the best of our abilities”. These are ideals or guiding principles, rather than executable statements. For example, if a function is added to a header file, the compilation of code that includes that header will slow down imperceptibly. That’s acceptable. Adding enormous amounts of code to a header so that compilation slows noticeably would be another matter.

These are ideals. They are what we would like to see done. If we succeed, most users will be very happy. However, they are not a recipe we could blindly follow to deliver a new standard. As is typical for ideals, they can conflict: there are tensions among the desired goals. This is common: Ideally, we want quality, on-time delivery, and low cost of products, but we know from vast experience that it is very hard to get all three. We want freedom of speech and absence of verbal intimidation, but balancing those two can be very hard. It is the same for the ideals of the “The C++ Bill of Rights”; we want all, but the committee will have to make hard choices.

These are ideals. They are meant to be rather grand statements, rather than nit-picked long sentences watered down by adjectives and caveats. It will be up to the committee members to interpret the ideals in the light of real design and scheduling problems.

There are just 7 ideals listed and they are quite general. We could add many more, but a laundry list of specifics would dull the appeal and be impossible to remember in the heat of discussions about direction, what can be delivered when, and specific technical concerns.

8. Caveat

The members of the direction group have been members of the standards committee for decades. We were (we think) chosen because we have deep and broad experience with the C++ technology, use, and community. We all have well documented track records. In the role of members of the direction group, we try to serve the good of the C++ community as a whole. It is not possible to completely separate what we consider best for the C++ community from the specific proposals we work on, but we try not to unfairly favor our own proposals. Similarly, we try not to discourage proposals or actions just based on personal biases. Some of our recommendations may not be universally appreciated, but nothing we say is meant to insult anyone.

Some of the recommendations here could be seen as contradictory. We see such as fundamental tensions in a design, requiring tradeoffs, rather than contradictions. Design is hard, design by committees is even harder.

9. Acknowledgements

Many of the arguments and points of view have a long history in the committee. Thanks to all who contributed. In particular, thanks to the authors of the documents we reference here.

10. References

- [CG,2015] [C++ Core Guidelines](#).
- [Clow,2018] Marshall Clow, Beman Dawes, Gabriel Dos Reis, Stephan T. Lavavej, Billy O'Neal, Bjarne Stroustrup, and Jonathan Wakely: Standard Library Modules. P0581R1. 2018-2-6.
- [GDR,2016] G. Dos Reis, J. D. Garcia, J. Lakos, A. Meredith, N. Myers, B. Stroustrup: [A Contract Design](#). P0380R1. 2016-07-11.
- [GDR,2017] G. Dos Reis, J. D. Garcia, J. Lakos, A. Meredith, N. Myers, B. Stroustrup: [Support for contract based programming in C++](#). P0542R2. 2017-11-26.
- [Nishanov,2017] Gor Nishanov: [C++ Extensions for Coroutines TS Document](#). N4680. 2017-07-30.
- [Hinnant,2017] Howard Hinnant: [Extending <chrono> to Calendars and Time Zones](#). P0355R4. 2017-10-1.
- [McLaughlin,2017] Michael McLaughlin, Herb Sutter, Jason Zink, Guy Davidson: [A Proposal to Add 2D Graphics Rendering and Display to C++](#). P0267r6. 2017-07-30.
- [Quora1] [Can I use C++ for hard real-time embedded systems and kernel development?](#)
- [Quora2] [Does C-11 and C-14 make it easy to use -C++ in kernel development?](#)
- [Saks,2016] Dan Saks "extern c: Talking to C Programmers about C++". CppCon 2016 keynote.
- [Stroustrup,1994] B. Stroustrup: The Design and Evolution of C++. Addison-Wesley. 1994.
- [Stroustrup,1993] B. Stroustrup: [A History of C++: 1979-1991](#). Proc. ACM HOPL-2. 1993.
- [Stroustrup,2007] B. Stroustrup: [Evolving a language in and for the real world: C++ 1991-2006](#). Proc. ACM HOPL-III. 2007.
- [Stroustrup,2014] B. Stroustrup: [Programming -- Principles and Practice Using C++ \(Second Edition\)](#). May 2014. Addison-Wesley. ISBN 978-0321992789. 1st edition 2008.
- [Stroustrup,2015] B. Stroustrup: [Thoughts about C++17](#). N4492. 2015-05-15.
- [Stroustrup,2015b] B. Stroustrup, H. Sutter, and G. Dos Reis: [A brief introduction to C++'s model for type- and resource-safety](#). Isocpp.org. October 2015. Revised December 2015.
- [Sutter,2018] [WG21 Practices and Procedures](#). ISO/IEC JTC1/SC22/WG21/SD-4. 2018-01-17.
- [Winkel,2017] JC van Winkel, Jose Daniel Garcia, Ville Voutilainen, Roger Orr, Michael Wong, and Sylvain Bonnal: [Operating principles for C++](#). P0559R0. 2017-01-31.
- [Voutilainen,2017] Ville Voutilainen: [To boldly suggest an overall plan for C++20](#). P0592R0. 2017-02-05.
- [Wakely,2017] **Jonathan Wakely**: [Extensions for Networking](#). N471. 2017-11-27.
- [Winters,2017] Titus Winters, Bjarne Stroustrup, Daveed Vandevoorde, Beman Dawes, Michael Wong, and Howard Hinnant: [C++ Stability, Velocity, and Deployment Plans](#). P0684R0 2017-06-19.