

Proposing a C++1Y Swap Operator, v2

Document #: WG21 N3746
Date: 2013-08-30
Revises: [N3553](#)
Project: Programming Language C++
Reply to: Walter E. Brown <webrown.cpp@gmail.com>

Contents

1	Introduction	1	5	Acknowledgments	5
2	Discussion	1	6	Addendum	5
3	Proposed syntax and semantics . .	3	7	Bibliography	5
4	Conclusion	5	8	Revision history	7

Abstract

We propose a new *swap* operator on native types for incorporation into the C++1Y core language. For user-defined types, we propose that this new operator be treated as a special member function, and that, when viable, its presence provide an alternative to traditional memberwise implementation of defaulted assignment semantics.

1 Introduction

As an algorithm in the C++ standard library, `std::swap` has been the focus of much discussion over the past years. Most of that discussion has centered on the means for other algorithms (e.g., `std::sort`) to call on user-supplied custom versions of the swap functionality, especially when that functionality is made available in foreign namespaces.

Although a general solution to those underlying lookup issues would be welcome, we believe that the swap primitive is itself sufficiently important¹ so as to warrant special treatment. Indeed swapping has been the subject of considerable programming language research over the years, and there are even several research results concluding that swap is a more fundamental operation than is (copy) assignment.² Independent of such issues, we believe it is time to incorporate a swap operator³ into C++.

2 Discussion

Fifty years ago, Barron et al. considered “Simultaneous assignment commands” as “the general form of an assignment command [whose operands are] two explicit lists [that] must contain the

¹ In [Krü09], D. Krügler considers `swap` to be one of the “rather fundamental cornerstones which are near to language facilities instead of library features.” The next section continues, “The swap function . . . is intimately related to the basics of copying and moving.”

² For example, Weide et al. state: “An analysis of the pros and cons of all options available for the built-in data movement operator in imperative languages shows that the swap operator is the best choice, while the assignment operator is the worst” [WPH02, Abstract]. Similarly, Harms and Weide argue “that a simple alternative to copying as a data movement primitive — swapping (exchanging) the value of two variables — has potentially significant advantages in the context of the design of generic reusable software components” [HW91, Introduction].

³ Due to euphony, we expect this operator to become known informally from time to time as the *swaperator*. ©

same number of members, and the command denotes a simultaneous assignment of each right-hand member to the corresponding left-hand member” [BBH⁺63, §24]. Even so long ago, swapping was envisioned: the third of the three examples following this description (albeit presented without explicit explanation) is $a, b := b, a$.

Over a decade later, R. Kieburtz cited this paper, explaining that “simultaneous multiple assignment allows one to perform an exchange of values by means of a single assignment, $\mathbf{I}, \mathbf{J} := \mathbf{J}, \mathbf{I}$.” He continues, “The main objection voiced to the use of simultaneous assignment is the potential for ambiguity. . . . For instance, the assignment, $\mathbf{A}[\mathbf{I}], \mathbf{A}[\mathbf{J}] := \mathbf{X}, \mathbf{Y}$ is ambiguous in case the index variables \mathbf{I} and \mathbf{J} happen to take the same value” [Kie76].

Although C++11’s `std::initializer_list` and `std::tuple` might be considered small steps in this direction, we are not proposing the full generality of unbounded simultaneous multiple assignment, as we have reservations as to its viability in a C++ context. Nonetheless, swapping (as if by simultaneous mutual assignment) has long played an important role in C++ programming.

In [PHA⁺09], S. Pike observes⁴ that “exchanging assignment has been used as the basis for the design of a research language, Resolve, and associated component libraries. Resolve, adapted as a discipline for C++ programming, has been evaluated in both educational and industrial software settings, including a 100K+ SLOC commercial software product. . . .” That evaluation concluded, in part:

The swapping paradigm works. It made it not only possible, but remarkably easy, to address the data movement dilemma in a way that preserved modular reasoning without sacrificing performance. Copying was only occasionally required in our application — and we think this would be the rule in many applications. . . . We attribute the remarkably clean bug report history of this product family to the relatively simple reasoning about behavior that resulted from this single most-important design decision [HBW00, Conclusion].

Far beyond such academic interest, exception-free swapping plays an important — even critical — role in a recognized C++ pattern usually termed *copy-and-swap*.⁵ This idiom allows programmers to achieve strong exception safety⁶ in assignment operators⁷ and other transactions⁸ when resources need to be managed. In a context that admits exceptions, swapping is considered a primary tool: once a type has a swap function which cannot fail, other functions can more easily provide the strong exception safety guarantee while correctly (a) supplanting old resources with new replacements and (b) disposing of the old resources in a timely fashion.

But the `swap` algorithm has demonstrated some persistent, frustrating issues with respect to C++ coding. For example, it took a long time before we understood how to call the algorithm in the presence of type-specific versions.⁹ We thought that that problem was solved¹⁰ but analogues

⁴ Citing both [SW94] and [HBW00].

⁵ See, for example, <http://stackoverflow.com/questions/3279543/what-is-the-copy-and-swap-idiom>. Note that in a C++11 context, copy-and-swap subsumes move-and-swap, in which a copy operation is replaced by a presumably less expensive move operation when the source is an rvalue.

⁶ According to M. Austern, a function \mathbf{f} offering such a guarantee either succeeds in its entirety, or else ensures that “the state of all objects that \mathbf{f} modifies will be restored to the state they had before \mathbf{f} was called” [Aus97]. Thus, in case of failure, there will be no observable side effects.

⁷ See H. Sutter’s [Sut02, Item 22], an expanded version of his on-line article “Exception-Safe Class Design, Part 1: Copy Assignment” found in its original version at <http://www.gotw.ca/gotw/059.htm>.

⁸ *Commit-or-rollback* semantics is the term by which the strong guarantee is often known in such a context.

⁹ See, for example, `comp.lang.c++.moderated` thread “Namespace issue with specialized swap,” 2000-03-12, at http://groups.google.ca/group/comp.lang.c++.moderated/browse_thread/thread/b396fedad7dcdc81.

¹⁰ Namely, by calling `swap` unqualified, following a `using` declaration.

keep popping up in new contexts, most recently while discussing `noexcept` behavior on the `std-proposals` reflector.¹¹

We believe this issue is so important as to warrant an addition to the C++ core language, the swap operator.

3 Proposed syntax and semantics

We propose:

3.1 Basics

- (a) that the swap operator be named `operator::=`: (a spelling selected, in part, for its symmetric appearance).
- (b) that the swap operator share the same near-lowest precedence and the same right associativity as the family of assignment operators.
- (c) that the swap operator be applied to two operands denoting mutable objects, almost always of identical type.¹²
- (d) that the swap operator's left operand be a modifiable lvalue expression and the right operand be a modifiable glvalue expression (so as to support the idiom `x := f(y)`).
- (e) that the swap operator return an lvalue referring to the left operand.¹³

3.2 Scalar and array types

- (f) that homogeneous `operator::=` be implicitly provided by the core language for every non-const scalar type¹⁴ `T`, as if by overloaded, non-throwing, namespace-scope functions declared:¹⁵

```
1 inline T& operator ::= (T& x, T&& y) { see below; return x; }
2 inline T& operator ::= (T& x, T& y) { return x ::= std::move(y); }
```

- (g) that applying the swap operator in an expression such as `x ::= y` produce side effects known as the *conventional swap semantics*:
 - If, after evaluating the operands (i.e., their unsequenced *value computation*), subexpression `x` denotes the same object as subexpression `y` denotes, there is no effect and the implementation may at its discretion elide the operation, simply returning a reference to the common object.
 - Otherwise,¹⁶ as if by simultaneous mutual assignment, (a) the (original) value of `y` replaces the value of the object denoted by `x` while (b) the (original) value of `x` replaces the value of the object denoted by `y`.

¹¹ See `std-proposals` thread "A proposal to add swap traits to the standard library," 2013-02-02, especially the contributions of Howard Hinnant.

¹² As is the case for `std::swap`, it seems axiomatic that `operator::=` applied twice consecutively to the same operands should be indistinguishable from zero applications. When the operands' types can be different (e.g., one `int` and one `double`), this axiom can be all too easily violated: if the double's initial value were 3.14, its value would be 3.0 after it has been swapped to the `int` and back. Nonetheless, we choose to permit heterogeneous user-defined `operator::=`: because some such cases, such as `operator::=(bool&, vector<bool>::reference)`, seem potentially useful. (After all, "being able to represent a value of type `BOOLEAN` on a single bit is, if not itself a pillar of civilization, one of the secrets of a happy life" [Mey10]. ☺)

¹³ This is unlike `std::swap`, whose return type is `void`.

¹⁴ "Arithmetic types, enumeration types, pointer types, pointer to member types, `std::nullptr_t`, and cv-qualified versions of these types are collectively called *scalar types*" [basic.types]/9 (cross-references elided).

¹⁵ We have not shown the `volatile` variants, as these are expected only rarely to occur in practice.

¹⁶ This merely describes the effect of a single call to algorithm `std::sort`, whose canonical implementation uses three copy operations (or, if available, three move operations instead): `T t{std::move(x)}; x = std::move(y); y = std::move(t);` Under the as-if rule, an equivalent implementation for distinct operands of certain types might avoid need for extra storage (beyond that of the operands) via three consecutive bitwise xor operations: `x ^= y; y ^= x; x ^= y;` Other possibilities include implementation via XCHG or similar op code, or via true concurrent mutual assignment.

(h) that if

- x and y have array types with identical extent n , and
- $x[0] := y[0]$ is a valid expression,

the core language evaluate $x := y$ as if by evaluating $x[i] := y[i]$ for $i = 0, 1, \dots, n-1$.

3.3 Class types

- (i) that `operator:=` be an overloadable operator, i.e., “given meaning when applied to expressions of class type” [expr]/2.
- (j) that a class-scope homogeneous swap operator be considered a special member function [special] of every class `C`, and that, in much the same way as is now done with assignment operators, such functionality be declared implicitly if `C`’s definition does not explicitly declare it.
- (k) that `C`’s homogeneous `operator:=` be declared in the form of overloaded member functions:¹⁷

```
1 inline C& C::operator :=( C&& y) & { see below; return *this; }
2 inline C& C::operator :=( C& y) & { return *this :=: std::move(y); }
```

- (l) that these `operator:=` member functions of `C` be also declared `noexcept(b)`, where $b = \text{is_nothrow_move_constructible}<C>::\text{value} \ \&\& \ \text{is_nothrow_move_assignable}<C>::\text{value}$.
- (m) that these implicitly-declared swap operators of `C` be defined as deleted if
- any subobject has a deleted, inaccessible, or ambiguous swap operator;
 - `C` fails to meet the requirements of a movable (i.e., `MoveConstructible` and `MoveAssignable`) type; or
 - `C` has a non-static data member of reference type or of const non-class type (or array thereof);

and that it be declared as defaulted otherwise.

- (n) that these swap operators of `C`, if defaulted and not declared as deleted, achieve the conventional swap semantics by applying `operator:=` to the operands’ subobjects in the usual order¹⁸ with each pair of corresponding subobjects swapped via `operator:=` in the manner appropriate to their type.

3.4 Type traits

- (o) that new type traits `is_swappable<>` and `is_nothrow_swappable<>`, with the obvious semantics, be added to [meta.unary.prop].¹⁹

3.5 Class assignment operator(s)

- (p) that, in a standard-layout class, the presence (or absence) of any swap operator have no impact with respect to C++11 rules regarding any implicitly declared assignment operator.
- (q) that if a non-standard-layout class `C`
- explicitly declares neither a copy nor a move assignment operator, and
 - has a `noexcept` swap operator that is not defined as deleted,
- then an assignment operator of the form:

¹⁷ We have again not shown `volatile` variants. but for a different reason this time: It should be up to the programmer to determine, on a per-class basis, whether such overloads are needed.

¹⁸ Direct base classes first, then immediate non-static data members of `C`, each in the order in which declared.

¹⁹ In [Mor13], A. Morrow motivates and proposes like-named traits. While our version presumes a proposed core language extension, Morrow’s seems based on contemporary library technology. There is no conflict for the moment, as LWG, upon reflection, has declined to pursue the library proposal: “Consensus that this is a good problem to solve, but not yet sure that this paper is the right solution. . . .” [LWG13].

```
1 C & operator = ( C c ) & noexcept(is_nothrow_swappable<C>::value);
```

be declared implicitly in **C**, and no other assignment operator(s) be implicitly declared in **C**.

- (r) that, if defaulted and not defined as deleted, this last form of assignment operator be implicitly defined, if and when needed, so as to call `operator::(c)` and return the result thereof.

4 Conclusion

This paper has proposed a swap operator, `operator::`, for addition to C++1Y and has further proposed its application, where viable, as an alternative implementation technique for defaulted class assignment operators. We invite feedback from WG21 participants and other knowledgeable parties, and especially invite implementors to collaborate with us in order to experiment and gain experience with this proposed new language feature.

5 Acknowledgments

Many thanks to the readers of early drafts of this paper for numerous helpful discussions and insightful reviews of this work and its predecessors.

6 Addendum

Since [Bro13], the first version of this paper, was published, J. Yasskin proposed [Yas13a, Yas13b, Yas13c] the following (lightly reformatted) **exchange** algorithm that “assigns a new value to the object and returns the old value”:

```
1 template< typename T, typename U = T >
2 T
3   exchange( T & obj, U && new_val )
4 {
5   T old_val = std::move(obj);
6   obj = std::forward<U>(new_val);
7   return old_val;
8 }
```

We are pleased that this proposal was adopted for the standard library. Moreover, we believe that this functionality is sufficiently primitive so as to make a very good candidate for integration into the core language. We mention this here to seek early feedback from WG21 as to the idea; with sufficient encouragement we will pursue a composite proposal encompassing both **swap** and **exchange**²⁰ operators for the core language.

7 Bibliography

- [Aus97] Matthew H. Austern: “Standard Library Exception Policy.” ISO/IEC JTC1/SC22/WG21 document N1077 (pre-London mailing) 1997-05-30.
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/1997/N1077.pdf>.
- [BBH⁺63] David W. Barron, John Noel Buxton, D. F. Hartley, E. Nixon, and Christopher Strachey: “The Main Features of CPL.”
<http://comjnl.oxfordjournals.org/content/6/2/134.full.pdf+html>.
 In *The Computer Journal*, 6(2):134–143, 1963.

²⁰ The name `operator::` seems appropriate.

- [Bro13] Walter E. Brown: “Proposing a C++1Y Swap Operator.” ISO/IEC JTC1/SC22/WG21 document N3553 (pre-Bristol mailing), 2013-03-12, <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3553.pdf>.
- [HBW00] Joseph E. Hollingsworth, Lori Blankenship, and Bruce W. Weide: “Experience Report: Using RESOLVE/C++ for Commercial Software.” <http://homepages.ius.edu/JHOLLY/pdf/FSE8.PDF>.
In *SIGSOFT Softw. Eng. Notes*, 25(6):11–19, November 2000. ISSN: 0163-5948.
Also in *Proceedings of the 8th ACM SIGSOFT International Symposium on Foundations of Software Engineering: Twenty-first Century Applications*, SIGSOFT ’00/FSE-8, pages 11–19. ACM Press, 2000. ISBN: 1-58113-205-0.
- [HW91] Douglas E. Harms and Bruce W. Weide: “Copying and swapping: Influences on the design of reusable software components.” <http://www.cse.ohio-state.edu/~weide/rsrg/documents/1991/91HW.pdf>.
In *IEEE Trans. Softw. Eng.*, 17(5):424–435, May 1991. ISSN: 0098-5589.
- [Kie76] Richard B. Kieburtz: “Programming without pointer variables.” <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.107.7746&rep=rep1&type=pdf>.
In *SIGPLAN Not.*, 11(SI):95–107, March 1976. ISSN: 0362-1340.
- [Krü09] Daniel Krüglér: “Moving Swap Forward (revision 1).” ISO/IEC JTC1/SC22/WG21 document N2979 (post-SantaCruz mailing), 2009-11-04. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2009/n2979.html>.
- [LWG13] Library Working Group: “Library Working Group Report.” 2013-04-19. <http://wiki.edg.com/wiki/bin/view/Wg21bristol/LwgReport>.
- [Mey10] Bertrand Meyer: “Reflexivity, and other pillars of civilization.” 2010-02-06. <http://bertrandmeyer.com/2010/02/06/reflexivity-and-other-pillars-of-civilization>.
- [Mor13] Andrew Morrow: “A proposal to add swappability traits to the standard library.” ISO/IEC JTC1/SC22/WG21 document N3619 (pre-Bristol mailing), 2013-03-15. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3619.html>.
- [PHA⁺09] Scott M. Pike, Wayne D. Heym, Bruce Adcock, Derek Bronish, Jason Kirschenbaum, and Bruce W. Weide: “Traditional Assignment Considered Harmful.” <http://www.cse.ohio-state.edu/rsrg/documents/2009/09PHABKW.pdf>.
In *Proceedings of the 24th ACM SIGPLAN Conference Companion on Object-Oriented Programming Systems, Languages, and Applications*, OOPSLA ’09, pages 909–916. ACM Press, 2009. ISBN: 978-1-60558-768-4.
- [SW94] Marulli Sitariman and Bruce W. Weide: “Component-based Software Using RESOLVE.” <http://www.cse.ohio-state.edu/~weide/rsrg/documents/1994/94SW/index.html>.
In *SIGSOFT Softw. Eng. Notes*, 19(4):21–22, October 1994. ISSN: 0163-5948.
- [Sut02] Herb Sutter: *More Exceptional C++: 40 New Engineering Puzzles, Programming Problems, and Solutions*. Addison-Wesley, 2002. ISBN: 0-201-70434-X.
- [WPH02] Bruce W. Weide, Scott M. Pike, and Wayne D. Heym: “Why Swapping?” <http://people.cs.vt.edu/~edwards/RESOLVE2002/proceedings/Weide2.html>.
In Stephen H. Edwards, editor: *Proceedings of the RESOLVE Workshop 2002*. Technical Report TR-02-11, Dept. of Computer Science, Virginia Tech, Blacksburg, VA. 2002-06.
- [Yas13a] Jeffrey Yasskin: “`exchange()` utility function.” ISO/IEC JTC1/SC22/WG21 document N3511 (mid-Portland/Bristol mailing), 2013-01-10. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3511.html>.
- [Yas13b] Jeffrey Yasskin: “`exchange()` utility function, revision 2.” ISO/IEC JTC1/SC22/WG21 document N3608 (pre-Bristol mailing), 2013-03-15. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3608.html>.
- [Yas13c] Jeffrey Yasskin: “`exchange()` utility function, revision 3.” ISO/IEC JTC1/SC22/WG21 document N3668 (post-Bristol mailing), 2013-04-19. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3668.html>.

8 Revision history

Revision	Date	Changes
1	2013-03-12	<ul style="list-style-type: none">• Published as N3553.
2	2013-08-30	<ul style="list-style-type: none">• Appended §6 re Yasskin's exchange papers.• Updated footnote and bibliography with latest information re Morrow's paper.• Performed numerous editorial tweaks.• Published as N3746.