

# Three `<random>`-related Proposals, v2

Document #: WG21 N3742  
Date: 2013-08-30  
Revises: [N3547](#)  
Project: JTC1.22.32 Programming Language C++  
Reply to: Walter E. Brown <[webrown.cpp@gmail.com](mailto:webrown.cpp@gmail.com)>

---

## Contents

<b>1</b>	<b>Addition to <code>&lt;algorithm&gt;</code></b>	<b>1</b>
<b>2</b>	<b>Additions to <code>&lt;random&gt;</code></b>	<b>4</b>
<b>3</b>	<b>Deprecations in <code>&lt;cstdlib&gt;</code></b>	<b>5</b>
<b>4</b>	<b>Proposed wording</b>	<b>5</b>
<b>5</b>	<b>Summary and conclusion</b>	<b>9</b>
<b>6</b>	<b>Acknowledgments</b>	<b>9</b>
<b>7</b>	<b>Bibliography</b>	<b>9</b>
<b>8</b>	<b>Revision history</b>	<b>10</b>

---

## Abstract

This paper proposes (1) to add one function template to `<algorithm>`, (2) to add a few novice-friendly functions to `<random>`, and (3) to deprecate some related `<cstdlib>` legacy interfaces. The unifying factor in this tripartite proposal is the entities' respective connection with the random number component of the C++11 standard library.<sup>1</sup>

## 1 Addition to `<algorithm>`

*[Sampling is the first step if you are looking to move things forward.*

— TOM PIRKO

### 1.1 Proposal

In [Aus08a], Matt Austern proposed several additional algorithms (mostly taken from SGI's original STL implementation<sup>2</sup>) for the standard library. Among these were `random_sample` and `random_sample_n`, based respectively on the well-known algorithms S (“selection sampling technique”) and R (“reservoir sampling”) elucidated by Donald Knuth in [Knu97, §3.4.2]. The paper succinctly describes these algorithms as “two important versions of random sampling, one of which randomly chooses  $n$  elements from a range of  $N$  elements and the other of which randomly chooses  $n$  elements from an input range whose size is initially unknown. . . .”

After WG21 consideration at the Sophia-Antipolis meeting, Austern updated the proposal. Among other changes, he withdrew these two algorithms: “The LWG was concerned that they might not be well enough understood for standardization. . . . It may be appropriate to propose those algorithms for TR2” [Aus08b]. The wiki minutes of the discussion are equally brief and terse: “Bjarne feels rationale is insufficient. PJ worries we will get it wrong. Lawrence worries that

---

<sup>1</sup> Readers seeking greater familiarity with this component may find [Bro13] to be a helpful source of background information and tutorial guidance with numerous usage examples.

<sup>2</sup> See <http://www.sgi.com/tech/stl>.

people will roll their own and get it wrong. Good candidate for TR2” [LWG08]. The subsequent LWG vote regarding these proposed algorithms achieved a solid consensus (10-1, 2 abs.) in favor of their future inclusion in a Technical Report (now termed a Technical Specification).

Since WG21 is now preparing updates to C++11 and drafting new Technical Specifications, it is an appropriate time to reconsider adding these algorithms, which (as Austern pointed out) have long been part of SGI’s STL implementation. Unlike the SGI implementation on which Austern’s proposal was based, the present proposal features a unified (common) interface, `sample`, that selects between algorithms R and S based on the categories of the iterators supplied at the point of call.

## 1.2 Expository implementation<sup>3</sup>

First, here is the common interface. It relies on overload resolution with tag dispatch to select the appropriate algorithm based on its arguments’ iterator categories:

```

1  template< class PopIter, class SampleIter, class Size, class URNG >
2  inline SampleIter
3      sample( PopIter first, PopIter last
4              , SampleIter out
5              , Size n, URNG&& g
6              )
7  {
8      using pop_t  = typename iterator_traits<PopIter  >::iterator_category;
9      using samp_t = typename iterator_traits<SampleIter>::iterator_category;
10
11     return __sample( first, last, pop_t{}
12                    , out, samp_t{}
13                    , n, std::forward<URNG>(g)
14                    );
15 }

```

We next exhibit the reservoir sampling algorithm R.<sup>4</sup> It requires only input iterators for access to the population being sampled, but needs a random access iterator to (the start of) the reservoir that will ultimately hold the resulting sample:

<sup>3</sup> Beware of bugs in this code; I have only tried it, not proved it correct. [© Apologies to Knuth.]

<sup>4</sup> Knuth’s original algorithm R guarantees stability, but at the cost of an extra level of indirection, an additional sorting step, etc. The SGI version, as proposed here and in Austern’s earlier paper, sacrifices stability in the interest of performance, but retains all other characteristics, especially the all-important randomness of the resulting sample.

```

1  template< class PopIter, class SampleIter, class Size, class URNG >
2  SampleIter
3  __sample( PopIter first, PopIter last, input_iterator_tag
4            , SampleIter out, random_access_iterator_tag
5            , Size n, URNG&& g
6            )
7  {
8      using dist_t = uniform_int_distribution<Size>;
9      dist_t d{};
10
11     Size sample_sz{0};
12     while( first != last && sample_sz != n )
13         out[sample_sz++] = *first++;
14
15     for( Size pop_sz{sample_sz}; first != last; ++first, ++pop_sz ) {
16         using param_t = typename dist_t::param_type;
17         param_t const p{0, pop_sz};
18         Size k{ d(g, p) };
19         if( k < n )
20             out[k] = *first;
21     }
22     return out + sample_sz;
23 }

```

We conclude this section with the selection sampling algorithm S. It needs forward iterators to access the population, but only an output iterator to the resulting sample:

```

1  template< class PopIter, class SampleIter, class Size, class URNG >
2  SampleIter
3  __sample( PopIter first, PopIter last, forward_iterator_tag
4            , SampleIter out, output_iterator_tag
5            , Size n, URNG&& g
6            )
7  {
8      using dist_t = uniform_int_distribution<Size>;
9      dist_t d{};
10
11     Size unsampled_sz = distance(first, last);
12     for( n = min(n, unsampled_sz); n != 0; ++first ) {
13         using param_t = typename dist_t::param_type;
14         param_t const p{0, --unsampled_sz};
15         if( d(g, p) < n )
16             *out++ = *first, --n;
17     }
18     return out;
19 }

```

## 2 Additions to <random>

*Random numbers should not be generated with a method chosen at random.*

— DONALD E. KNUTH

### 2.1 Proposal

We propose to add to <random> the following modest toolkit of novice-friendly functions:<sup>5</sup>

- **global\_urng()**  
Grants access to a URNG object of implementation-specified type.
- **randomize()**  
Sets the above URNG object to an unpredictable state.<sup>6</sup>
- **pick\_a\_number(from, thru)**  
Returns an **int** variate uniformly distributed in the closed **int** range [**from**, **thru**].
- **pick\_a\_number(from, upto)**  
Returns a **double** variate uniformly distributed in the half-open **double** range [**from**, **upto**).

Taking advantage of this toolkit, we further propose that **global\_urng()** become the default argument<sup>7</sup> for the respective **UniformRandomNumberGenerator**&& parameters of the existing algorithm **shuffle** and of the above-proposed algorithm **sample**.

### 2.2 Representative implementation

Here, first, is a sample implementation of the **global\_urng()** and **randomize()** functions, demonstrating an implementor's choice of an engine type for the shared URNG:

```

1 auto& global_urng( )
2 {
3     static std::default_random_engine u{};
4     return u;
5 }

7 void randomize( )
8 {
9     static std::random_device rd{};
10    global_urng().seed( rd() );
11 }
```

If the implementor chooses differently, the same functions might be implemented this way:

```

1 auto& global_urng( )
2 {
3     static std::random_device u{};
4     return u;
5 }

7 void randomize( ) { }
```

In contrast, the **pick\_a\_number** functions can be written so as to be agnostic with respect to the implementor's choice of URNG type:

<sup>5</sup> This proposal is in response to the near-incessant user ~~whining kvetching~~ feedback, from even respected C++ *cognoscenti*, that <random>'s contents are insufficiently "friendly" to novices and other casual users.

<sup>6</sup> Unpredictability is the ideal. Using a computer, we typically settle for very-very-very-very-hard-to-predict.

<sup>7</sup> Neither of these parameters currently specifies any default argument.

```

1 int pick_a_number( int from, int thru )
2 {
3     static std::uniform_int_distribution<> d{};
4     using parm_t = decltype(d)::param_type;
5     return d( global_urng(), parm_t{from, thru} );
6 }

8 double pick_a_number( double from, double upto )
9 {
10    static std::uniform_real_distribution<> d{};
11    using parm_t = decltype(d)::param_type;
12    return d( global_urng(), parm_t{from, upto} );
13 }

```

### 3 Deprecations in <cstdlib>

*If a feature is not deprecated [I] don't see any point in not using it.*

— HARIHARAN SUBRAMANIAN

By common consensus at several consecutive WG21 meetings during which the C++11 random number facility was being discussed and shaped into its final form, it has for a number of years been the long-term plan to excise the legacy C random number facility from the `std` namespace. Indeed, obliquely acknowledging the quality<sup>8,9</sup> of C++11's <random> header, WG21 voted several years ago to insert a Note<sup>10</sup> into [c.math]/5 as a head start on this plan: “The random number generation . . . facilities in this standard are often preferable to `rand`.”

In light of our above proposal (§2) for a more user-friendly interface to <random>, we concurrently propose to execute the next step of the plan to discourage the use of the traditional C function `rand()` as well as its associated seeding function `srand()` and upper limit macro `RAND_MAX`. (These are declared in the traditional C header <stdlib.h> and the corresponding C++ header <cstdlib>.) We propose to begin this transition by formally deprecating:

- `std::rand()`, `std::srand()`, and (when obtained via <cstdlib>) `RAND_MAX` and
- algorithm `random_shuffle()` (keeping `shuffle`, however).

### 4 Proposed wording

All proposed wording is relative to WG21 draft [DuT12]. It is recommended to apply the wording **additions** and **deletions** in the order shown. Editorial notes are displayed against a gray background.

Append the following new paragraph to [depr.c.headers]:

<sup>8</sup> “[B]y and large, I think it’s the best random number library design of all, by a mile. If I were a random number, I’d think I died and went to heaven” [Ale07].

<sup>9</sup> “The C++11 <random> is very STL-like in that it sets up requirements for random number generators. . . , and random distributions. . . , and then the client can mix and match the two. It’s a really very cool design [Hin12].

<sup>10</sup> The language for this Note was proposed by Beman Dawes in [Daw08]; [Bec08] was the first Working Paper to incorporate it.

Use of function `rand`, function `srand`, and macro `RAND_MAX` is deprecated when these names are introduced via the `<cstdlib>` header. Use of function `std::rand` and function `std::srand` is deprecated when these names are introduced into the `std` namespace via the header `<stdlib.h>`.

Copy all of the current [alg.random.shuffle] to a new section in Annex D, applying to the copy the changes shown below.

## Random shuffle

[depr.alg.random.shuffle]

The following templates are in addition to those specified in Clause [alg.random].

```
template<class RandomAccessIterator>
    void random_shuffle(RandomAccessIterator first, RandomAccessIterator last);
```

```
template<class RandomAccessIterator, class RandomNumberGenerator>
    void random_shuffle(RandomAccessIterator first, RandomAccessIterator last,
                       RandomNumberGenerator&& rand);
```

```
template<class RandomAccessIterator, class UniformRandomNumberGenerator>
    void shuffle(RandomAccessIterator first, RandomAccessIterator last,
                UniformRandomNumberGenerator&& g);
```

*Effects:* Permutes the elements in the range `[first, last)` such that each possible permutation of those elements has equal probability of appearance.

*Requires:* `RandomAccessIterator` shall satisfy the requirements of `ValueSwappable` (17.6.3.2). The random number generating function object `rand` shall have a return type that is convertible to `iterator_traits<RandomAccessIterator>::difference_type`, and the call `rand(n)` shall return a randomly chosen value in the interval `[0, n)`, for `n > 0` of type `iterator_traits<RandomAccessIterator>::difference_type`. ~~The type `UniformRandomNumberGenerator` shall meet the requirements of a uniform random number generator (26.5.1.3) type whose return type is convertible to `iterator_traits<RandomAccessIterator>::difference_type`.~~

*Complexity:* Exactly  $(last - first) - 1$  swaps.

*Remarks:* To the extent that the implementation of these functions makes use of random numbers, the implementation shall use the following sources of randomness:

The underlying source of random numbers for the first form of the function is implementation-defined. An implementation may use the `rand` function from the standard C library.

In the second form of the function, the function object `rand` shall serve as the implementation's source of randomness.

~~In the third `shuffle` form of the function, the object `g` shall serve as the implementation's source of randomness.~~

In the synopsis in [algorithms.general], (a) remove both overloads of `random_shuffle`, (b) insert `"= global_urng()"` as the default value for the last function parameter of `shuffle`, and (c) insert `"= decltype(global_urng())"` as the default value for the last template parameter of `shuffle`.

Change the heading of [alg.random.shuffle] as follows and then make the remaining indicated adjustments to the retitled section's text:

## 25.3.12 Random shuffling and sampling

[alg.random.shuffle]

```

template<class RandomAccessIterator>
    void random_shuffle(RandomAccessIterator first, RandomAccessIterator last);

template<class RandomAccessIterator, class RandomNumberGenerator>
    void random_shuffle(RandomAccessIterator first, RandomAccessIterator last,
        RandomNumberGenerator&& rand);

template<class RandomAccessIterator,
    class UniformRandomNumberGenerator=decltype(global_urng())>
    void shuffle(RandomAccessIterator first, RandomAccessIterator last,
        UniformRandomNumberGenerator&& g=global_urng());

```

*Effects:* Permutes the elements in the range `[first, last)` such that each possible permutation of those elements has equal probability of appearance.

*Requires:* `RandomAccessIterator` shall satisfy the requirements of `ValueSwappable` (17.6.3.2). ~~The random number generating function object `rand` shall have a return type that is convertible to `iterator_traits<RandomAccessIterator>::difference_type`, and the call `rand(n)` shall return a randomly chosen value in the interval `[0, n)`, for `n > 0` of type `iterator_traits<RandomAccessIterator>::difference_type`.~~ The type `UniformRandomNumberGenerator` shall meet the requirements of a uniform random number generator (26.5.1.3) type whose return type is convertible to `iterator_traits<RandomAccessIterator>::difference_type`.

*Complexity:* Exactly  $(last - first) - 1$  swaps.

*Remarks:* To the extent that the implementation of ~~these~~ `this` functions makes use of random numbers, ~~the implementation shall use the following sources of randomness:~~

~~The underlying source of random numbers for the first form of the function is implementation-defined. An implementation may use the `rand` function from the standard C library.~~

~~In the second form of the function, the function object `rand` shall serve as the implementation's source of randomness.~~

~~In the third `shuffle` form of the function,~~ the object `g` shall serve as the implementation's source of randomness.

Append the following text to the newly retitled [alg.random], also incorporating the new declaration into the header `<algorithm>` synopsis at the beginning of [algorithms].

```

template <class PopulationIterator, class SampleIterator,
    class Distance, class URNG=decltype(global_urng())>
    SampleIterator sample(PopulationIterator first, PopulationIterator last,
        SampleIterator out, Distance n,
        UniformRandomNumberGenerator&& g=global_urng());

```

*Requires:*

- `PopulationIterator` shall meet the requirements of an `InputIterator` type.
- If `PopulationIterator` meets the additional requirements of a `ForwardIterator` type, `SampleIterator` shall meet the requirements of an `OutputIterator` type. Otherwise, `SampleIterator` shall meet the requirements of a `RandomAccessIterator` type.
- `PopulationIterator`'s value type shall be `Assignable`, and shall be writable to `out`.
- `Distance` shall be an integer type.

- **UniformRandomNumberGenerator** shall meet the requirements of a uniform random number generator ([rand.req.urng]) type whose return type is convertible to **Distance**.
- **out** shall not be in the range [**first**, **last**).

*Effects:* Copies  $\min(\mathbf{last} - \mathbf{first}, \mathbf{n})$  elements (the *sample*) from [**first**, **last**) (the *population*) to **out** such that each possible sample has equal probability of appearance.

*Returns:* The end of the resulting sample range.

*Complexity:* No worse than  $\mathcal{O}(\mathbf{n})$ .

*Remarks:*

- Stable if and only if **PopulationIterator** meets the requirements of a **ForwardIterator**.
- To the extent that the implementation of this function makes use of random numbers, the object **g** shall serve as the implementation's source of randomness.

Append the following text as a new subsection under [rand.util], also incorporating the new declarations into the header <random> synopsis at the beginning of [rand.synopsis].

### Global URNG functions

[rand.util.global]

```
auto& global_urng();
```

*Effects:* Ensures the instantiation of a mutable object *u* that has static storage duration and whose implementation-specified type meets the requirements of a uniform random number generator ([rand.req.urng]) type. [ *Note:* The implementation may select this type on the basis of performance, size, quality, or any combination of such factors, so as to provide at least acceptable behavior for relatively casual, inexpert, and/or lightweight use. — *end note* ] The extent to which *u*'s initial state can be predicted or can ever be reproduced is unspecified.

*Synchronization:* It is implementation-defined whether this function may introduce data races ([res.on.data.races]).

*Returns:* An lvalue reference to *u*.

```
void randomize();
```

*Effects:* Sets the uniform random number generator given by **global\_urng()** to an (ideally) unpredictable state.

```
int pick_a_number(int from, int thru);
double pick_a_number(double from, double upto);
```

*Returns:* Returns a variate uniformly distributed (a) in the closed range [**from**, **thru**], for the first form of the function, and (b) in the half-open range [**from**, **upto**), for the second form of the function.

*Remarks:* To the extent that the implementation of these functions makes use of random numbers, the implementation shall use, as its source of randomness, the uniform random number generator given by `global_urng()`.

## 5 Summary and conclusion

This paper has proposed three modest changes to the C++11 standard, each related to the random number facility. While the proposals are largely independent of each other, we respectfully recommend their adoption together via approval of the paper in its entirety.

## 6 Acknowledgments

Many thanks to the readers of early drafts of this paper for their thoughtful comments, and to Matt Austern for the original proposal of the sampling algorithms.

## 7 Bibliography

- [Ale07] Andrei Alexandrescu: "Re: Conveniently generating random numbers with TR1 random." `comp.std.c++`, 2007-06-11.
- [Aus08a] Matt Austern: "More STL algorithms." ISO/IEC JTC1/SC22/WG21 document N2569 (post-Bellevue mailing), 2008-02-29.  
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2008/n2569.pdf>.
- [Aus08b] Matt Austern: "More STL algorithms (revision 2)." ISO/IEC JTC1/SC22/WG21 document N2666 (post-Sophia mailing), 2008-06-11.  
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2008/n2666.pdf>.
- [Bec08] Pete Becker: "Working Draft, Standard for Programming Language C++." ISO/IEC JTC1/SC22/WG21 document N2691 (post-Sophia mailing), 2008-06-27.  
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2008/n2691.pdf>.
- [Bro13] Walter E. Brown: "Random Number Generation in C++11." ISO/IEC JTC1/SC22/WG21 document N3551 (pre-Bristol mailing), 2013-03-12.  
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3551.pdf>.
- [Daw08] Beman Dawes et al.: "Thread-Safety in the Standard Library (Rev 2)." ISO/IEC JTC1/SC22/WG21 document N2669 (post-Sophia mailing), 2008-06-13.  
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2008/n2669.pdf>.
- [DuT12] Stefanus Du Toit: "Working Draft, Standard for Programming Language C++." ISO/IEC JTC1/SC22/WG21 document N3485 (post-Portland mailing), 2012-11-02.  
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2012/n3485.pdf>.
- [Hin12] Howard Hinnant: Untitled response to posted query. 2012-31-07.  
<http://stackoverflow.com/questions/11717433/tutorial-or-example-code-for-extending-c11-random-with-generators-and-distribu>.
- [Knu97] Donald E. Knuth: *The Art of Computer Programming, Volume 2: Seminumerical Algorithms (Third Edition)*. Addison-Wesley, 1997. ISBN 0-201-89684-2.
- [LWG08] "Library Working Group: Wednesday." Minutes of LWG meeting, Sophia-Antipolis, France. 2008-06.  
<http://wiki.edg.com/twiki/bin/view/Wg21sophiaAntipolis/LibraryWorkingGroup>.

## 8 Revision history

Version	Date	Changes
1	2013-03-12	<ul style="list-style-type: none"><li>• Published as N3547.</li></ul>
2	2013-08-30	<ul style="list-style-type: none"><li>• Corrected punctuation of WG21 bib entries.</li><li>• Relocated <code>++pop_sz</code> within reservoir algorithm.</li><li>• Clarified wording of editorial note.</li><li>• Relocated <code>--unsampled_sz</code> within selection algorithm.</li><li>• Improved wording proposing default arguments.</li><li>• Added footnote re unpredictability.</li><li>• Tweaked fonts and spacing.</li><li>• Added epigraphs.</li><li>• Added footnote and citation for background paper.</li><li>• Added new declarations to <code>random</code> synopsis.</li><li>• Recoded to use new C++14 features.</li><li>• Required <code>u</code> to be mutable.</li><li>• Published as N3742.</li></ul>