International Journal of Parallel Programming (2022) 50:319–340
https://doi.org/10.1007/s10766-022-00726-5

A Deterministic Portable Parallel Pseudo-Random Number Generator for Pattern-Based Programming of Heterogeneous Parallel Systems

August Ernstsson[1] · Nicolas Vandenbergen[2] · Jörg Keller[3] · Christoph Kessler[1]

Received: 18 October 2021 / Accepted: 26 January 2022 / Published online: 22 March 2022
© The Author(s) 2022

Abstract
SkePU is a pattern-based high-level programming model for transparent program execution on heterogeneous parallel computing systems. A key feature of SkePU is that, in general, the selection of the execution platform for a skeleton-based function call need not be determined statically. On single-node systems, SkePU can select among CPU, multithreaded CPU, single or multi-GPU execution. Many scientific applications use pseudo-random number generators (PRNGs) as part of the computation. In the interest of correctness and debugging, deterministic parallel execution is a desirable property, which however requires a deterministically parallelized pseudo-random number generator. We present the API and implementation of a deterministic, portable parallel PRNG extension to SkePU that is scalable by design and exhibits the same behavior regardless where and with how many resources it is executed. We evaluate it with four probabilistic applications and show that the PRNG enables scalability on both multi-core CPU and GPU resources, and hence supports the universal portability of SkePU code even in the presence of PRNG calls, while source code complexity is reduced.

Keywords Skeleton programming · Parallelizable algorithmic pattern · Heterogeneous system · GPGPU · Deterministic parallel pseudo-random number generator

1 Introduction

For the foreseeable future, computer systems for performance-demanding application domains such as HPC, machine learning and image processing, will continue to be characterized by multi-/many-core parallelism and heterogeneity. Faced with the increasing slowdown of Moore's Law, a "Cambrian explosion" of computer

Extended author information available on the last page of the article

architectures is foreseen [16] that will continue to introduce new CPU and GPU architectures and entirely new accelerator types at a fast pace to sustain future hardware performance growth, while at the same time an increasing share of performance growth needs to come from both application and system software improvements. This imposes a challenge on the software side: How can we support the creation of truly portable, future-proof software that is high-level yet can efficiently leverage the hardware resources of today's and tomorrow's heterogeneous parallel architectures without permanent rewriting and re-optimization?

The skeleton programming approach [3, 4] is a powerful and programmer-friendly way to portable high-level parallel and heterogeneous programming, which has been demonstrated by a number of programming frameworks during the last decade [5, 7, 8, 21, 26, 28]. *Skeletons* are generic programming constructs based on higher-order functions such as *map*, *reduce*, *stencil* etc. used to express certain parallelism patterns, that can be parameterized in problem-specific code (the so-called *user functions*) and that come with parallel or accelerator-specific implementations (the so-called *backends*), which are hidden behind a portable high-level API, today usually based on C++. In short, skeletons expose *possible* application-level parallelism but not its implementation details to the programming framework and its runtime system, which might then be free to decide which skeleton instances in a program to use (and how), and which ones should better remain sequential.

SkePU [9] is a pattern-based high-level programming model for transparent program execution on heterogeneous parallel computing systems. A key feature of SkePU is that, in general, the selection of the backend, and thus, the execution platform for a skeleton-based function call need not be determined statically, i.e. prior to execution. On single-node systems, SkePU can select among CPU, multithreaded CPU, single or multi-GPU execution. For example, run-time selection of the expected fastest [6] backend (depending on operand size and location) can be tuned automatically based on training executions or manually set by a flag outside the program's source code. By careful API design, each SkePU program is a valid C++ program with sequential execution semantics if compiled with a standard C++(11 or later) compiler, and SkePU's design for portability aims at executions over multiple cores or one or several GPUs to show the same input-output behavior as this sequential view.

Many scientific applications, such as Monte-Carlo simulations, use pseudo-random number generators (PRNGs) as part of the computation. In the interest of correctness and debugging, deterministic parallel or heterogeneous execution of such a program that remains consistent with sequential execution also in terms of generated random numbers is a desirable property, which however requires a *deterministic* parallel pseudo-random number generator. This becomes a challenge with SkePU's design of late decision about sequential, parallel or accelerator execution.

In this paper, we present the principle, API and implementation of a deterministically parallelized portable PRNG extension to SkePU that exhibits the same behavior regardless where and with how many resources a SkePU program is executed. Our deterministic PRNG parallelization also relaxes the implicit

dependence structure of applications using the PRNG. We show that the implementation is scalable on both multi-core CPU and GPU resources, and hence supports the universal portability of SkePU code even in the presence of PRNG calls. It also leads to more compact source code. Core contributions are the determinism and the high-level language integration of our approach. While our solution is prototyped and evaluated in SkePU, where it is important due to the execution unit of a skeleton call being statically unknown, the approach could be adapted and integrated into other frameworks for high-level portable pattern-based parallel programming.

The remainder of this paper is organized as follows: Section 2 introduces background about SkePU and parallel random number generators, shows two motivating examples of previous workarounds used with SkePU to achieve deterministic parallel PRNG behavior, and discusses their drawbacks. Section 2.3 discusses related work. Section 3 explains three fundamental parallelization methods for PRNGs and presents the new API and implementation of the new built-in deterministic parallel PRNG in SkePU. Section 4 presents experimental results, and Sect. 5 concludes.

2 Background and Related Work

2.1 SkePU

In its current version [9], SkePU (https://skepu.github.io) provides 7 data-parallel skeletons: `Map` (elementwise transformation), `MapOverlap` (stencil updates in 1D...4D), `MapPairs` (generic outer product of vectors), `Reduce` (generic reduction), `Scan` (generic prefix sums), and the combinations `MapReduce` and `MapPairsReduce`. In general, the skeletons allow both element-wise accessed, random-access and scalar operands and are fully variadic within each of these categories. Most skeletons also allow multiple return operands. Array-based operands can have 1 to 4 dimensions.

Listing 1: A simple SkePU example: Dot product.

```
1  #include <skepu>
2  ...
3  double mult (double x, double y) { return x * y; }  // user function
4  double add (double x, double y) { return x + y; }  // user function
5  ...
6  void main( void )
7  {   ...
8      skepu::Vector<double> u(size), v(size);   // two 1D data-containers
9      ...
10     auto dotprod = skepu::MapReduce<2>(mult, add); // instantiate
11     std::cout << dotprod(u, v);  // call skeleton instance on 2 vectors
12 }
```

By instantiating a skeleton with one or several matching problem-specific user functions (detailed further below), a callable entity (a *skeleton instance*) is

generated, see Listing 1 for an example. The `MapReduce` instance `dotprod` can be used like any manually written function, but comes with multiple back-ends (implementations) for the different target platforms, such as sequential execution, OpenMP multithreaded execution, CUDA and OpenCL for GPUs. There exists also a cluster backend for SkePU that targets the MPI interface of the StarPU runtime system [9].

For passing array-type data into or out of skeleton instance calls, so-called *data containers* must be used, which transparently perform memory management, software caching and data transfers for contained array elements. SkePU 3 supports data containers for arrays in 1D (`Vector`), 2D (`Matrix`), 3D and 4D (`Tensor` X `D`). All data containers are generic in the element type.

User functions must be side-effect free and be written in a restricted subset of C++ (e.g., no dynamic memory allocation, no explicit parallelism, no skeleton instantiations or -calls, no global variable access) as they are translated into the various platform-specific programming models (e.g., OpenMP, CUDA, OpenCL) and may execute on an accelerator with a possibly separate address space. For array-based operands passed as arguments to user functions, the foreseen access pattern is specified by access proxy parameter objects such as `Vec` $<>$ for random-accessed vector, `Region` `XD` $<>$ for stencil halo regions in `MapOverlap` ($X \in \{1, ..., 4\}$) or `Index` X `D` for the index of the element operated on in `Map`-based skeletons; element-wise access is the default (no proxy parameter type required). Access to the proxy elements depends on where the user function will be executed and is thus entirely managed by SkePU's data containers. User functions can also be defined as C++ *lambda* (anonymous) functions, allowing for in-line skeleton instantiation.

2.2 Parallel Pseudo-Random Number Generation

A pseudo-random number generator is a finite state automaton. Each time it is invoked, its output function computes and outputs a pseudo-random number in a pre-defined range from the current inner state, and transitions the inner state via the state transition function (also called update function) into the follow-up state. The generator only receives input upon the time of seeding, when the seed is processed by the initialization function to produce an initial inner state. Thus, the generator only has a very limited amount of randomness, which is stretched over many outputs, i.e. pseudo-random. Still, current generators pass statistical tests such as Diehard. The complexity to achieve this may lie in the output function and/or the update function. For a complex output function, the update function can be as simple as a counter [17].

If an output of m bits is produced, the inner state comprises more than m bits. The state transition function mostly is non-bijective.[1] Thus, the state graph of the PRNG comprises one node for each state x, and a directed edge $(x, f(x))$ for the transition from x to its follow-up state $f(x)$, assuming f as the state transition function. Thus

[1] A notable exception is the linear congruential generator with transition function $f(x) = ax + b \bmod N$ for a, b chosen such that the period is maximum [18], e.g. a prime and $b = 0$

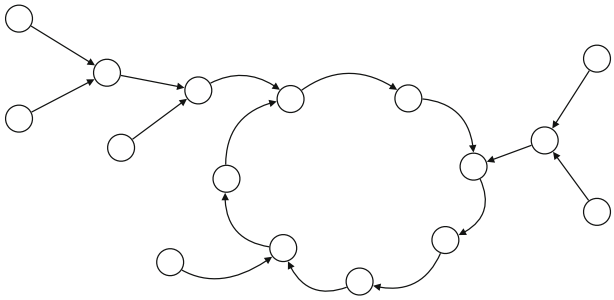

Fig. 1 State-space of a pseudo-random number generator

each node has an outdegree of exactly 1, but the indegree can vary. An example state graph is shown in Fig. 1.

Flajolet and Odlyzko [11] investigated the expected structure of state graphs if all possible transition functions are equally likely. The graph falls into a small number of weakly connected components, of which one comprises the majority of the nodes (about 75%). Each component comprises a cycle with a number of trees directed towards the cycle, where the largest tree is expected to comprise 50% of all nodes. The expected length of the longest cycle is less than $2\sqrt{N}$, where N is the number of nodes, i.e. quite short. Trees are ragged with depth about $\sqrt{N}$.

The sequence of generated pseudo-random numbers is only dependent on the seed. In a sequential program with a deterministic program flow, the calls to the pseudo-random number generator will produce exactly the same numbers at the same program place if the seed is fixed. If the program is parallelized, then the PRNG state becomes a shared resource. Moreover, the order of calls to the PRNG changes: consider e.g. a nested loop with one call per iteration of the inner loop, where the outer loop is parallelized, so that now the first iterations of all instances of the inner loop call the PRNG first. Still, a deterministic parallel execution with results similar to the sequential version (and independent of the number of threads used to parallelize the outer loop) demands that the sequence of PRNG outputs for each inner loop execution remains unchanged, e.g. to do debugging in the sequential version when the parallel version has an error. This calls for a deterministic PRNG implementation as part of the parallelization.

2.3 Related Work

Kneusel [17] has a chapter on parallel PRNGs, but with respect to deterministic execution only reports a manual construction of duplicating the state variable for each thread, plus skipping a number of states in order to achieve the same state as in

a sequential execution. He also explains counter-based PRNGs and their suitability for parallelization because they allow skipping any given number of states with constant effort. Fog [12] discusses requirements on PRNGs in parallel computations, but focuses on avoiding overlapping sequences in different threads by combining generators, while L'Ecuyer et al. [19] focus on providing independent streams and substreams. Salmon et al. [27] focus on output functions for counter-based PRNGs to provide fast skipping of states but still provide good statistical quality. All do not focus on deterministic execution independent of parallelization, and have static mapping of tasks to threads in mind.

Leiserson et al. [20] argue that SPRNG [22], which provides a deterministic parallel PRNG, shows poor performance on Cilk programs and thus is not suitable for massive parallelism. They propose pedigrees, a mechanism to achieve a kind of linearization (i.e. a kind of equivalence to a sequential execution) in a Cilk program independent of the Cilk scheduler. However, they do not address pattern-based parallelization.

Parallel PRNG specifically for GPU include the cuRAND library for CUDA, SYCL-PRNG for SYCL, and work by Ciglarič et al. [2] for OpenCL. The Thrust skeleton library for CUDA also includes a PRNG library. Passerat et al. [23] discuss general aspects of PRNG on GPGPUs. GASPRNG [14] is an early attempt at realizing the full SPRNG generator set on CUDA GPUs, including clusters of CUDA GPU nodes.

2.4 Previous Manual Parallelization of PRNG in SkePU Programs

With previous versions of SkePU, a deterministic parallel random number generator behavior had been achieved by the two workarounds described in the following. However, we will show that both have drawbacks.

2.4.1 Monte-Carlo Pi Calculation—Index-Based Scrambling

As a first example, we consider a simple Monte-Carlo simulation, namely probabilistic Pi approximation. This computation can be easily expressed as a `MapReduce` instance, see Listing 2, where the user function needs to generate two pseudo-random numbers, one per dimension. Here, a deterministic parallel PRNG was simulated by an *index-scrambling technique*, i.e., generation of pseudo-random numbers does not follow the automaton-based best-practice technique described above; instead, they are calculated independently of each other based on a transformation of the index in the parallelized main loop. In the code example in Listing 2, the `scramble` function itself has been extracted from a SPH (Smoothed Particle Hydrodynamics) simulation code. The drawback of the index scrambling method is that it may not really produce random numbers of high quality but can expose more regular patterns.

Listing 2: Ad-hoc deterministic pseudo-random number generation by index scrambling in a Monte-Carlo method for Pi calculation in SkePU.

```
#include <iostream>
#include <skepu>

// Define c, s, s2, s3, MY_RAND_MAX as preprocessor constants

float scramble(int in)
{
    return ((((int)(10*s*s2*in + 4*c*s3 + 5*in + 10*s*in)) % MY_RAND_MAX)
            / ((float)MY_RAND_MAX));
}

float monte_carlo_sample(skepu::Index2D index)
{
    float x = scramble(index.row);
    float y = scramble(index.col);
    // check if (x,y) is inside region:
    return ((x*x + y*y) < 1) ? 1.f : 0.f;
}

float add( float lhs, float rhs ) { return lhs + rhs; }

int main(int argc, char *argv[])
{
    auto montecarlo = skepu::MapReduce<0>(monte_carlo_sample, add);

    const size_t samples = atoi(argv[1]);
    montecarlo.setDefaultSize(samples, samples);

    float pi = montecarlo() / (samples * samples) * 4;
    std::cout << pi << "\n";
}
```

2.4.2 Markov Chain Monte Carlo methods in LQCD – PRNG with Explicit State

The code excerpt in Listing 3 is extracted from a Lattice QCD mini-application which computes the Yang-Mills theory of the $SU(3)$ group. This computation is typically done by applying the Metropolis algorithm, a common Markov Chain Monte Carlo (MCMC) based method. The Metropolis calculations are performed on a 4D tensor whose elements are structures of complex number arrays, with a 81-point $(3 \times 3 \times 3 \times 3)$ stencil computation required to evaluate the Metropolis acceptance function. For an in-depth introduction to MCMC methods in LQCD, see [15].

Unlike the conventional Monte Carlo method showcased in Listing 2, MCMC methods are inherently sequential. Thus, a PRNG for MCMC methods has to be stateful, i.e. a finite state automaton as outlined in Sect. 2.2. This conflicts with the requirement that SkePU user functions must be side-effect free. The chosen solution for the user functions of Listing 3 is a sequential PRNG which is algorithmically equivalent to POSIX `drand48` but has an explicit state argument instead of `drand48`'s internal state variable.

For such an approach, the PRNG state has to be explicitly managed. As a dedicated data container for PRNG state is not a viable solution due to syntactical constraints of the `MapOverlap` skeleton, the state is embedded directly in the data set. This has the drawback of having an unusually large memory footprint for a PRNG. Specifically, the memory requirement for storing the PRNG states grows by $O(L^4)$ where L is the side length of the 4D tensor, i.e. linearly with the problem size. Usage of the proposed new library PRNG inside SkePU is expected to lower the

memory footprint of PRNG state storage to $O(p)$ where p is the number of computational units used in the selected backend.

While it would be possible to adapt the index-based scrambling technique of Listing 2 to perform the initial seeding of the resulting parallel PRNG, Listing 3 contents itself with using the `Scan` skeleton to force a non-repeating state set into existence. While this is viable as a quick and dirty solution to deterministic parallel PRNG seeding, it is likely to produce random numbers of suboptimal quality; in that respect, a mathematically robust library solution is preferable.

Listing 3: Simplified SkePU code for an explicit-state parallel PRNG for Markov-chain-based LQCD applications (Sect. 2.4.2).

```
typedef uint64_t PRNGState;

// Seeding:
skepu::Tensor4<PRNGState> ones( L, L, L, L, 1), prngs( L, L, L, L);
auto seedPRNGs = skepu::Scan( []( PRNGState x, PRNGState y){ return x+y;});
seedPRNGs( prngs, ones);

// Extracting:
inline PRNGState statelessDrand48( PRNGState prng)
{
    return (0x5deece66d * prng + 11) % (1LL<<48);
}
inline double normalize( PRNGState prng)
{
    return (double)prng / (double)(1LL<<48);
}

// Parallel state management:
struct localGauge; // 36 double-precision complex numbers
struct localGaugeAndPRNG
{
    localGauge data;
    PRNGState prng;
};
skepu::Tensor4<localGaugeAndPRNG> gaugeField( L, L, L, L);

// Gauge randomization:
localGaugeAndPRNG randomizeGauge( PRNGState prng)
{
    localGaugeAndPRNG gaugeNew;
    for (int idx = 0; idx < 36; idx++) {
        prng = statelessDrand48( prng);
        gaugeNew.data.at(idx).re = normalize( prng);
        prng = statelessDrand48( prng);
        gaugeNew.data.at(idx).im = normalize( prng);
    }
    gaugeNew.prng = prng;
    return gaugeNew;
}

// Metropolis step:
localGaugeAndPRNG localUpdate( skepu::Region4D<localGaugeAndPRNG> stencil)
{
    localGaugeAndPRNG proposal = randomizeGauge( stencil( 0,0,0,0).prng);
    double limen = someDeterministicStencilArithmetic( stencil, proposal);
    stencil( 0,0,0,0).prng = statelessDrand48( proposal.prng);
    if (normalize( stencil( 0,0,0,0).prng) >= limen) {
        stencil( 0,0,0,0).data = proposal.data;
    }
    return stencil( 0,0,0,0);
}

auto metropolisUpdate = skepu::MapOverlap( localUpdate);
for (int iter = 0; iter < Niter; iter++) {
    metropolisUpdate( gaugeField, gaugeField);
}
```

3 Designing a Deterministic PRNG for SkePU

We will now introduce a more systematic approach that provides deterministic parallel random number generation for use in SkePU, together with an API extension of SkePU 3 that makes PRNG streams a fundamental part of the API. We will start by discussing inherent challenges to pseudo-random number generation in parallel programming and proceed step by step towards a deterministic PRNG implementation at the framework level.

3.1 Global Synchronization

A straightforward approach to random number generation in parallel applications is to consider the PRNG as a shared resource. As such, the PRNG needs to be protected by the appropriate synchronization operations during access, to avoid race conditions such as multiple threads reading the same random value, which would decrease the quality of the random number stream, or even the PRNG state itself being corrupted due to simultaneous writes.

This approach ensures a high-quality random number stream as each value is generated in the same manner as in a sequential program. Any random number generator can be used in this approach, including external entropy sources, since synchronization guarantees protected sequentialized access. This synchronization does however add significant overhead and is unfeasible in massively parallel accelerators such as GPUs. Only if the synchronization method guarantees a deterministic order of accesses to the critical section containing the PRNG state (which is usually not the case for ordinary lock-based synchronization), the random number stream generated from this method will be itself deterministic. We cannot predict in which order the threads will generate a value from the PRNG and update the state space.

3.2 Stream Splitting

With the goal of avoiding or minimizing global synchronization of the PRNG state, we consider a different approach [13]. As a PRNG state has to be considered a shared resource for proper operation, we can get around the synchronization requirement by assigning each individual thread its own PRNG state. A thread-private PRNG stream does not need protected access and will yield a perfect sequential series of random values by itself. However—aside from a large increase in memory space consumed by the replicated states—with several or many parallel threads in the system, the aggregate random number stream over all task invocations will differ greatly from a sequential program.

Whether data-parallel tasks are assigned in blocks or interleaved, we effectively have *split* the single PRNG stream into many shorter sequences distributed over the working set in the same pattern as the data-parallel tasks. The resulting pattern can be seen in Fig. 2a. This degrades the quality of the random values in aggregate, which is undesirable for sensitive applications.

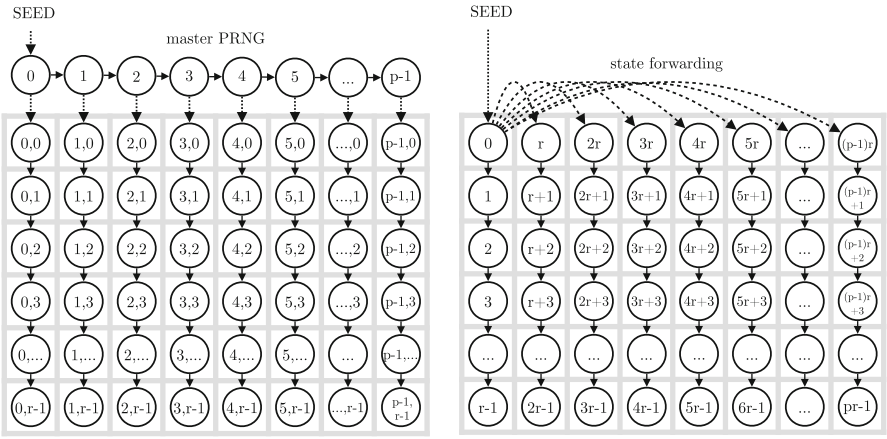

(a) Stream-splitting approach to parallel pseudo-random number generation.

(b) State-forwarding approach to parallel pseudo-random number generation.

Fig. 2 Approaches for parallelizing a PRNG sequence

There is another unfortunate consequence of this approach: ensuring determinism in the random value stream is possible, but with significant restrictions. Due to the aforementioned parallelization of the computation using the PRNG, the observed PRNG stream across the data set is a mangled mixture of (a potentially large number of) individual streams. This mangling has to be replicated in the sequential execution of the program to preserve determinism; and worse, all parallel backends have to observe the same such mapping. This can prove tricky when the parallel backends vary significantly in properties such as the available parallelism degree. A consequence of this behavior is also that in any execution of the program for which deterministic random values are desired, the maximal number of threads has to be known *a priori*, before even executing a sequential backend variant. If the degree of parallelism ever is increased, e.g. by moving to a larger processor, GPU, or cluster, the previous runs are invalidated with respect to the determinism criterion.

3.3 State Forwarding

The approach taken in this work is *state forwarding*. We attempt to side-step the issues of both the global synchronization as well as the stream splitting approaches. This is done by utilizing properties of the PRNG state spaces. A true sequential single-stream variant of the program is taken as the gold standard output, and the goal is to replicate the same output on any parallel backend, without the need of global synchronization or advance knowledge of parallelism degree. As in the stream splitting approach, data-parallel work items are deterministically mapped across available computational units (threads). This means that the number of tasks assigned to each thread is known ahead of time, and for simplicity without the loss of generalization we assume the work can be split evenly among threads.

Furthermore, we assume that the number of times a PRNG state is updated (i.e., the number of times a random value is generated) is known ahead of time for each work unit. Combining the knowledge of *work unit count* and *random calls per work unit*, we know exactly how many state-forwards each thread will generate in the respective data-parallel construct (i.e., skeleton invocation).

We can therefore, for each thread, *pre-forward* the state of the PRNG and store a copy of the forwarded state. These per-thread forwarded clones of the original PRNG can now act as the thread-private PRNG streams in the stream-splitting approach, with the additional property that when interleaved during the data-parallel execution, the aggregate observed stream now is equivalent to the sequential stream, which was the primary hurdle in the stream-splitting approach. Figure 2b illustrates the resulting pattern.

Still, the extra memory footprint of the thread-private PRNG states persists and will lead to additional overhead. The state-forwarding adds an additional computation step before the execution of the tasks, which can in the worst case be equally costly as the PRNG value extraction process itself (though it can also be parallelized). Properties of the PRNG state space have to be exploited to speed up the forwarding process and reduce the induced overhead.

The leapfrog resp. sequence splitting method for state forwarding, introduced by Celmaster and Moriarty [1] for use with vector computers, considers a special case that allows to parallelize the forwarding phase of the PRNG. A linearly congruential PRNG with factor a is partitioned into p linearly congruential PRNGs each to be used r times, which are defined based on the same linear factor a, by $seed(i) = (a^r \cdot seed(i-1)) \bmod m$ for $i = 1, ..., p$, $rand(i, 0) = seed(i)$ and $rand(i, j) = a \cdot rand(i, j-1) \bmod m$. Hence, the p PRNGs equally partition the period of the seed PRNG in contiguous sub-sequences of length r. First, the a^{ir} for $i = 0, ..., p-1$ and the seed sequence can be calculated in parallel by a `Scan` in $O(\log p)$ steps, using the property $a^{2k} \bmod m = ((a^k \bmod m)^2) \bmod m$. Then the *rand* calls are independent for each i. (For reasonably low numbers of p such as for a current multicore CPU, sequential computation of the seeds should be faster; this is done in the current implementation.) The leapfrog / sequence splitting method scales well but is known to have problems for lcg with power-of-2 values for modulus and p. Skipping can also be applied for counter-based PRNGs [27] with output functions based on block ciphers for better statistics at a higher cost.

3.4 Optimizing Long or Iterated Skeleton Chains by Pre-Forwarding

While some applications may consist of a single parallelized step (such as a parallel for loop or skeleton call; we will use the latter here), others, in particular larger applications, will have multiple phases which are individually parallelized. A common example is iterative applications where each iteration in turn consists of one or more skeleton calls. To achieve good efficiency, we need to ask the question: *when* is the PRNG state split and forwarded for the purposes of parallelization in a skeleton invocation scenario?

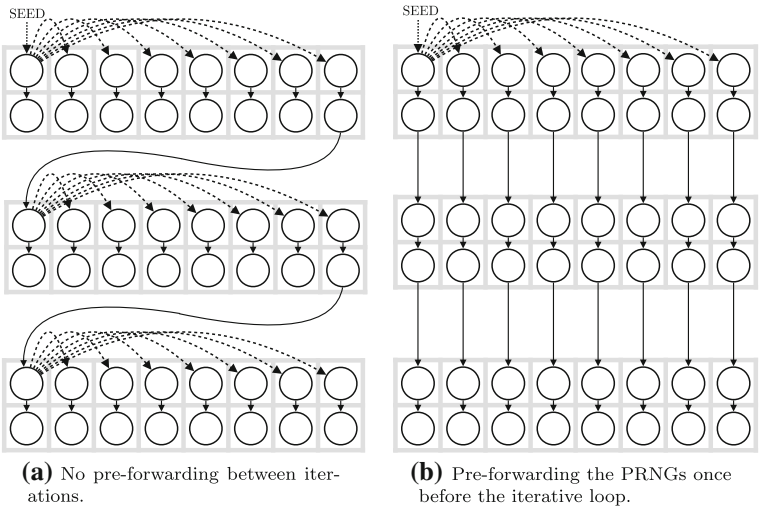

(a) No pre-forwarding between iterations.

(b) Pre-forwarding the PRNGs once before the iterative loop.

Fig. 3 Data container indexing and memory layout

In a naive implementation of the state-splitting approach, the state splitting and forwarding step (see Fig. 3a) is done right before each skeleton call. On the other hand, if we have a known number of skeleton calls (determinable by static analysis, lineage building [10], or program instrumentation), we only need to perform the splitting and forwarding of the PRNG states once per application. This is referred to as *pre-forwarding* and is illustrated in Fig. 3b. In practice, restrictions such as data-dependent control flow (e.g., branches or iteration bounds) may limit the degree to which pre-forwarding can be applied, and application programmers may benefit from awareness of the cost-reduction opportunities from pre-forwarding already during program design.

3.5 API Extension Design

We have implemented the state-forwarding approach in the skeleton programming framework SkePU 3. SkePU did not previously have a random number generation component, and as shown in Sect. 2.4, previous manual implementations of PRNG-like functionality in SkePU applications have been ad-hoc and substantially different from each other. A baseline contribution of a framework-level PRNG library in SkePU is the programmability gains from reducing the effort of designing probabilistic applications on top of SkePU, as well as readability benefits from having a unified system for random number generation across all SkePU programs.

3.5.1 Random Number Extraction in User Functions

As explained in Sect. 2.1, a SkePU skeleton is defined entirely by its type (e.g., `Map`), the signature of its instantiating user function, and state properties set on the resulting skeleton instance (such as `.setOverlap(...)` for `MapOverlap` instances).

PRNG extraction is made available in all skeletons with a fully data-parallel mapping stage, which is the entire skeleton set except for `Reduce` and `Scan`.[2]

As such, the user function signature ("header") itself should encode the use of random number extraction. This is analogous to the preexisting option for mapping user functions to request the index of the currently processed element (see Listing 2). Therefore, we encode PRNG reliance in the same way. At the start of the parameter list (after the index parameter, if any), a parameter of type `skepu::Random <N>&` is added. `N` is a compile-time constant used in SkePU's *template metaprogramming*-based implementation to deduce the number of random values extracted by the user function in the *dynamic extent* of its evaluation. N is required to be known ahead-of-time for the state forwarding to work and determinism to be preserved.[3] A compilation option allows for run-time verification that the extraction count is obeyed.

Value extraction is carried out by a call to one out of two member functions of the `skepu::Random <N>` object. `random.get()` produces integers in $[0, \text{SKEPURANDMAX})$ while `random.getNormalized()` returns real numbers in $[0, 1)$. Each call corresponds to one extraction and state update of the PRNG stream. A basic example of a user function with 5 random number extractions is shown in Listing 4.

Listing 4: User function with calls to the SkePU random number generator.

```
float uf( skepu::Random<5> &prng, skepu::Region1D<float> region )
{
    float res = 0;
    for (int i = -2; i <= 2; ++i)  // 1D stencil with random weights
        res += region(i) * prng.getNormalized();
    return res;
}
```

SkePU user functions are allowed to call other functions, subject to some but not all restrictions of skeleton-instantiating user functions. As the extraction count N is only required for instantiation, passing a PRNG stream object to indirect user functions is instead done with a `skepu::Random <>*` parameter with no positional requirement.

3.5.2 PRNG Streams and Skeleton Invocations

Once a `skepu::Random <N>&`-enabled user function is present, a skeleton can be instantiated as usual. In addition to the skeleton instance, a PRNG stream object needs to be defined in the program: an object of type `skepu::PRNG`. Initialization of the PRNG stream takes an optional *seed* integer argument. The seed changes the deterministic sequence generated in the stream and can be assigned from an external entropy source (e.g., a timestamp) if non-determinism across program runs is preferred.

[2] `Reduce` and `Scan` are parallelized through tree reductions reliant on the associativity property of their user functions.

[3] If determinism is not required by the application, `N` can be treated as an upper bound, which instead guarantees that no sub-sequences of random numbers are overlapping.

The stream object is a state machine which registers skeletons ahead of invocation time. Also in this way PRNG streams work like SkePU's index parameters: the stream is not part of a skeleton call's argument list. Instead they are registered as `skeleton.setPRNG(prng)`, if the programmer wants explicit control over stream objects and the way they map to skeleton instances; if no stream is registered, the skeleton picks a global PRNG stream by default at invocation time. The full flow chart of the registration and evaluation process is shown in Fig. 4. In short, several skeleton instances may be *registered* before reaching an *evaluation* event. Only at this point is the PRNG sequence split across computational units and forwarded to the appropriate state. The input size (i.e., the maximum degree of parallelism) has direct impact on the forwarding leaps and is only known at the evaluation point from the input arguments to the skeleton call.[4] In subsequent skeleton invocations, the PRNG object checks for existing forwarded state and skips directly to evaluation (refer to Fig. 3b).

Listing 5: Pi approximation using the new SkePU PRNG API.

```
#include <iostream>
#include <skepu>

int monte_carlo_sample(skepu::Random<2> &random)
{
    float x = random.getNormalized();
    float y = random.getNormalized();
    // check if (x,y) is inside region:
    return ((x*x + y*y) < 1) ? 1 : 0;
}

int add(int lhs, int rhs) { return lhs + rhs; }

int main(int argc, char *argv[])
{
    auto montecarlo = skepu::MapReduce<0>(monte_carlo_sample, add);

    skepu::PRNG prng;          // optional
    montecarlo.setPRNG(prng); // optional

    const size_t samples = atoi(argv[1]);
    montecarlo.setDefaultSize(samples);

    double pi = (double)montecarlo() / samples * 4;
    std::cout << pi << "\n";
}
```

Listing 5 shows a variant of the Monte-Carlo Pi calculation algorithm using the new SkePU API. Implementation with a `MapReduce <0>` skeleton enables a data-parallel computation without explicit data container allocation, as the algorithm needs no element-wise input data to the user function; all input is derived from the PRNG stream. Internally, SkePU will use two data containers: one input data set for the split PRNG sub-sequence states, and one output data set for the results of the user function invocations. Note, however, that SkePU will optimize the size of these intermediate data sets; they grow by $O(p)$, the number of computational units, and not $O(n)$, problem size (here the sample count).

[4] The input size is assumed to be uniform over a sequence of skeleton calls.

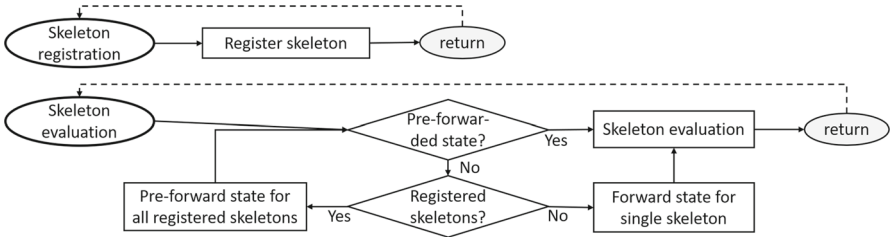

Fig. 4 Flow-chart of the deterministic PRNG implementation. Here ellipses are events and boxes correspond to processes

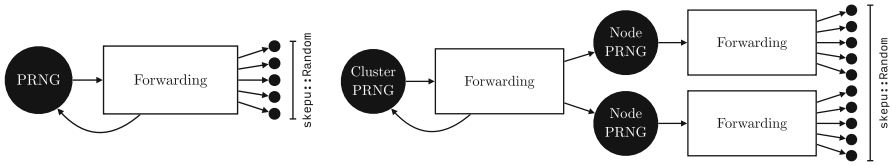

(a) Forwarding on node-level backends.

(b) Forwarding on cluster-level backends.

Fig. 5 Differences in state-forwarding approach in cluster backend

Our prototype implementation handles multiple PRNG stream objects across different skeleton calls, but a single skeleton call (and thus its user function) can only receive values from one PRNG stream per invocation. `skepu::Random` usage can be combined with most other SkePU features, with a notable exception being dynamic scheduling for multi-core execution introduced [9] in SkePU 3.

The SkePU implementation of deterministic PRNG streams cover a wide set of backend targets. For *OpenMP*, *OpenCL*, and *CUDA*, the forwarding is straightforward as the worker threads are homogeneous, running on equal hardware resources. In SkePU's *hybrid backend* which simultaneously targets multi-core CPU and GPU resources, the normal forwarding process is done twice in sequence: first once for the GPU, then the same stream is forwarded again for the CPU threads. The end result is a single set of forwarded thread-specific state objects (Fig. 5a), but with non-equal step count. Finally, we have an early prototype implementation targeting the *cluster* backend. In this case, parallelization is done in two steps: once among nodes through StarPU and once among CPU cores by means of OpenMP. The forwarding process for PRNG streams is therefore hierarchical, as illustrated in Fig. 5b.

4 Experimental Evaluation

For the performance evaluations in this section, we use a server with two six-core Xeon E5-2630L CPUs with two-way hardware multi-threading, a Nvidia K20c GPU, and 64 GiB of main memory. The system runs Ubuntu 18.04.5 LTS and GCC 10.3.0 is used as backend compiler with -O3 optimization level. Results are presented as the median of several measurement runs (varying between the programs).

4.1 Monte-Carlo Pi Approximation

We begin with the probabilistic Pi calculation from Sect. 2.4.1. SkePU code using the new `skepu::Random` API is shown in Listing 5.

Figure 6 contains the performance results from executing the SkePUized program on various backends. The Monte-Carlo Pi calculation algorithm is an interesting stress test due to the random number generation dominating the total work. The application scales well on the GPU using the OpenCL backend (up to 180x speedup compared to sequential in the presented results), even though the work done in the user function is very lightweight.

4.2 LQCD Mini-Application

For the LQCD mini-application introduced in Sect. 2.4.2, SkePU code using the new PRNG API is shown in Listing 6.

Figure 7 shows the times of 10 iterations of LQCD with the OpenMP backend, comparing the manual workaround of Listing 3 to the new version using `skepu::Random` of Listing 6. We can see that no new overheads are introduced while code complexity decreases (see Sect. 4.5).

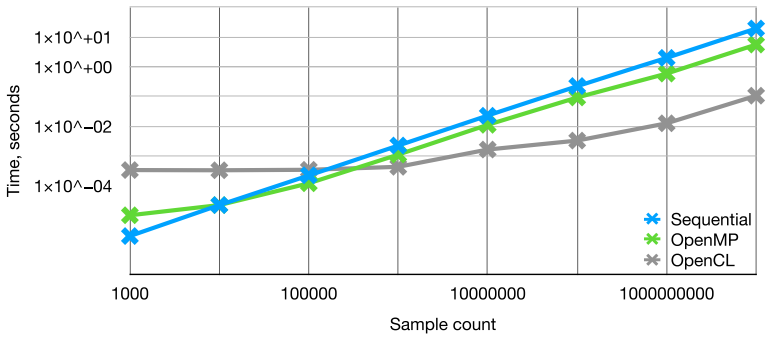

Fig. 6 Monte-Carlo Pi calculation with varying sample count on different backends

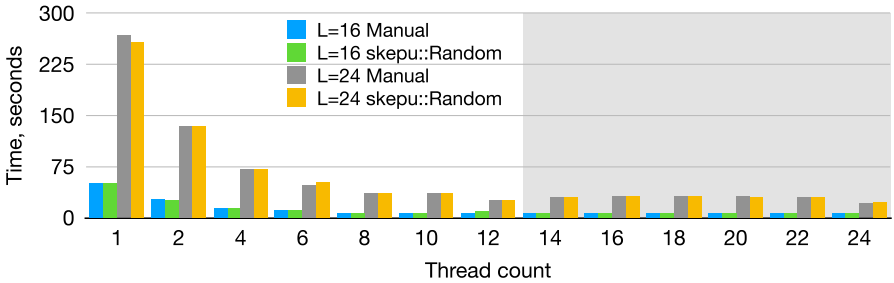

Fig. 7 Time (seconds) for 10 iterations of LQCD with lattice sizes $L = 16$ and $L = 24$ for varying number of hardware threads in the OpenMP backend

Listing 6: Markov-chain-based LQCD application with new SkePU PRNG API

```
// Data management:
struct localGauge; // 36 double-precision complex numbers
skepu::Tensor4<localGauge> gaugeField( L, L, L, L);

// Gauge randomization:
localGauge randomizeLocalGauge( skepu::Random<> *prng)
{
    localGauge gaugeNew;
    for (int idx = 0; idx < 36; idx++) {
        gaugeNew.at(idx).re = prng->getNormalized();
        gaugeNew.at(idx).im = prng->getNormalized();
    }
    return gaugeNew;
}

// Metropolis step:
localGauge localUpdate( skepu::Random<73>& prng,
                        skepu::Region4D<localGauge> stencil )
{
    localGaugeAndPRNG update = randomizeLocalGauge( prng);
    double limen = someMatrixArithmetic( stencil, update);
    if (prng.getNormalized() >= limen) {
        stencil( 0,0,0,0) = update;
    }
    return stencil( 0,0,0,0);
}

...
skepu::PRNG prng;
auto metropolisUpdate = skepu::MapOverlap( localUpdate);
metropolisUpdate.setPRNG( prng);
for (int iter = 0; iter < Niter; iter++) {
    metropolisUpdate( gaugeField, gaugeField);
}
```

4.3 Miller-Rabin Primality Testing

Listing 7: Pseudocode of Miller-Rabin probabilistic primality testing

```
int test( int n )
{
   result = true;
   for (int i = 1; i <= t; i++) {
      int a = rand();
      result = result & millerrabin( n, a );
      // if (result == false) break;
   }
   return result;
}
```

The Miller-Rabin primality test [25] is a probabilistic algorithm to determine for a given number if it is likely prime or not. The actual test gets two inputs: n, the number to be tested for primality, and a value a in the range $\{2...n-2\}$. The test performs a computation on a and n, and depending on the result it outputs "n is prime" or "n is composite". While the latter answer is always true, there is a certain probability that the former answer is wrong, and this probability can be reduced by doing the computation repeatedly with randomly chosen a, see Listing 7. This can be easily parallelized, as the t iterations are independent (except for calls to the PRNG), but for comparability it is helpful that the random choices are similar to the sequential version.

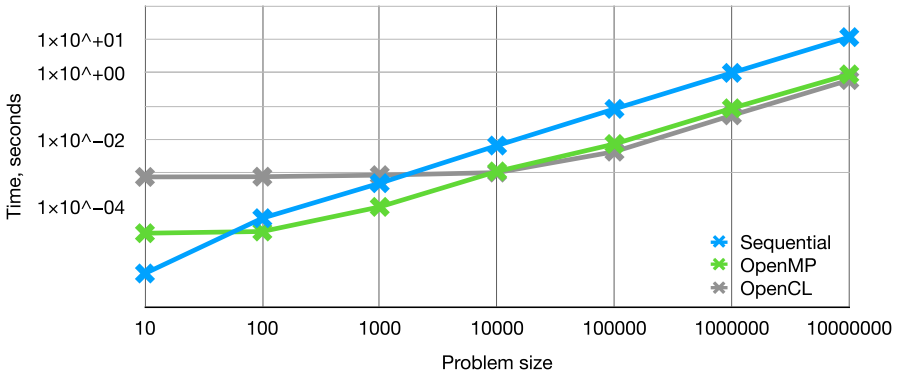

Fig. 8 Miller-Rabin primality test with varying sample count on different backends

Our SkePU implementation of the Miller-Rabin algorithm is largely based on an open-source implementation in C++ by Larsen[5] where the main Monte-Carlo parallelism is expressed by a `MapReduce <0>` skeleton instance.

Parallel performance of the SkePUized Miller-Rabin application can be seen in Fig. 8. Instruction flow is highly divergent throughout the algorithm due to data-dependent control flow, which is challenging for the GPU backend: it is just barely faster than multi-core CPU computation. This property distinguishes the program from the Monte-Carlo sampling algorithm wherein the PRNG values have no effect on control flow. For the multi-core OpenMP backend we observe speedup up to 13x.

4.4 Natural Noise Generation

The PRNG construct now built-in to SkePU generates uniformly distributed real or integer values ("white" noise; Fig. 9a). When other distributions are desired, post-processing of the generated data in application space can be a solution. One such scenario is for the generation of natural-looking noise patterns where the value distribution is dependent on factors such as signal frequency.[6] One way of generating such "colored" noise is with a gradient noise algorithm, also known as *Perlin noise* [24]. The algorithm (Listing 8) first generates n-dimensional grids of uniformly distributed values at different grid densities. The values are interpreted as gradients for the resulting n-dimensional noise pattern, and for each sample point, the neighbor gradients are interpolated to produce a noise value. The sampling process is repeated for each grid density (frequency) level, taking a weighted sum of the individual samples as the resulting output value. A typical output is found in Fig. 9b.

[5] C.S. Larsen: The Miller-Rabin primality test in C++. https://github.com/cslarsen/miller-rabin

[6] Used for example in computer-generated image production or as initial values in simulations.

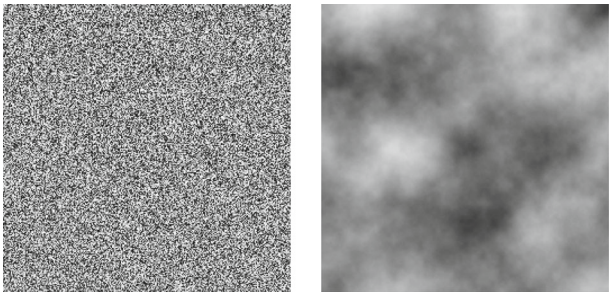

(a) Uniformly distributed noise.

(b) Colored (approximately Perlin) noise.

Fig. 9 Two-dimensional noise variants generated by a SkePU application

Listing 8: Skeleton call site in the 3D natural noise generator program.

```
skepu::Tensor3<float> noise(outDepth, outHeight, outWidth, 0);
auto grid_gen  = skepu::Map(generate_grid_point_value); // uses PRNG
auto noise_gen = skepu::Map(render_value_noise_3d);
for (int i = 0; i < octaves; ++i)
{
  // ... (computation of grid sizes and amplitude omitted)
  skepu::Tensor3<float> grid(gridWidth, gridHeight, gridDepth);
  grid_gen(grid); // consumes PRNG values
  noise_gen(noise, noise, grid, amplitude, outDepth, outHeight, outWidth);
}
```

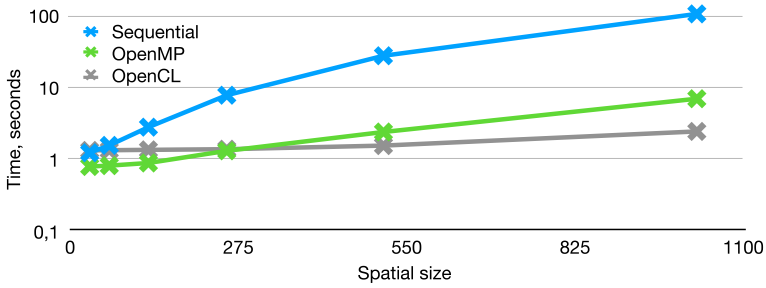

Fig. 10 3D natural noise generation results, varying the two spatial dimension side lengths. Vertical time axis is in logarithmic scale

Performance evaluation of the noise generator is done on a 3D domain. It produces 256 time-linked textures, with the spatial domain as a square with side length sampled at powers of two. The program generates the entire 3D domain in one sweep with the SkePU tensor data-container, with 10 superimposed octaves of noise. The results in Fig. 10 indicate very good scalability with just over 15x speedup on OpenMP and 43x on OpenCL.

4.5 Programmability Evaluation

In all[7] of Pi calculation, Lattice QCD and Miller-Rabin primality test, the implementations see a reduction in lines-of-code count after applying the SkePU PRNG API. This effect primarily comes from abstracting the implementation details of the PRNG engine itself. The entire code base of the SkePUized LQCD application is reported by `sloccount` to be 1,212 lines of code before applying the new `skepu::Random` API, and 1,137 afterwards. This amounts to a reduction by 6.2%. In addition, the change simplifies the data structure hierarchy, and fewer skeleton calls and user function declarations are necessary.

5 Conclusion and Future Work

We have proposed a method for realizing a deterministic parallel PRNG for use with skeleton-based high-level programming of heterogeneous parallel systems where the type and number of parallel execution resources of skeleton calls can be selected dynamically at run-time. We provided an extension of the SkePU API and its prototype implementation based on state forwarding. We evaluated it with four probabilistic applications on multi-core CPU and GPU backends, and found that the proposed API and parallelization approach performs at least equally well as manual workarounds while code complexity is reduced. The SkePU PRNG implementation, with CPU, GPU, and Hybrid backends, is publicly available as part of the open-source distribution.

Future work will extend the current implementation of the PRNG functionality in SkePU, in particular the cluster backend. Another possible area of investigation is different types of PRNG generators, which may need specially-adapted forwarding schemes.

Acknowledgements This work was partly funded by EU H2020 project 801015 (EXA2PRO), http://www.exa2pro.eu and by CUGS, Linköping University.

Funding Open access funding provided by Linköping University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

[7] The natural noise generator was written without a prior reference implementation.

References

1. Celmaster W., Kevin J.M.: Moriarty A method for vectorized random number generators. J. Comput. Phys. **64**, 271–275 (1986)
2. Ciglarič T., Štrumbelj E., et al.: An OpenCL library for parallel random number generators. J. Supercomput. **75**(7), 3866–3881 (2019)
3. Cole M.: Bringing skeletons out of the closet: a pragmatic manifesto for skeletal parallel programming. Parallel comput. **30**(3), 389–406 (2004)
4. Cole M.: Algorithmic Skeletons: Structured Management of Parallel Computation. Pitman and MIT Press, Cambridge, Mass (1989)
5. Danelutto M., Torquati M.: Structured parallel programming with "core" FastFlow. In: Central European Functional Programming School, volume 8606 of LNCS, pages 29–75. Springer (2015)
6. Dastgeer U., Li L., Kessler C.: Adaptive implementation selection in the SkePU skeleton programming library. In: Advanced Parallel Processing Technologies, pages 170–183. Springer (2013)
7. del Rio Astorga D., Dolz M.F., Fernández J, García J.D.: A generic parallel pattern interface for stream and data processing. Concurrency Comput. Pract. Exp. **29**(24), 4175 (2017)
8. Enmyren J., Kessler C.: SkePU: A multi-backend skeleton programming library for multi-GPU systems. In: Proceedings of the fourth international workshop on High-level parallel programming and applications, pp. 5–14. ACM (2010)
9. Ernstsson A., Ahlqvist J., Zouzoula S., Kessler C.: SkePU 3: Portable high-level programming of heterogeneous systems and HPC clusters. Int. J. Parallel Programm., May (2021) https://doi.org/10.1007/s10766-021-00704-3
10. Ernstsson A., Kessler C.: Extending smart containers for data locality-aware skeleton programming. Concurrency Comput. Pract. Exp. **31**(5), e5003 (2019)
11. Flajolet P., Odlyzko A.M.: Random mapping statistics. In: Jean-Jacques Quisquater and Joos Vandewalle, editors, Advances in Cryptology—Proc. EUROCRYPT '89 Workshop on the Theory and Application of Cryptographic Techniques, Houthalen, Belgium, volume 434 of LNCS, pages 329–354. Springer, (1989)
12. Fog A.: Pseudo-random number generators for vector processors and multicore processors. J. Modern Appl. Stat. Meth. **14**, 308–334 (2015)
13. Frederickson P., Hiromoto R., Jordan T.L., Smith B., Warnock T.: Pseudo-random trees in Monte Carlo. Parallel Comput. **1**(2), 175–180 (1984)
14. Gao S., Peterson G.D.: GASPRNG: GPU accelerated scalable parallel random number generator library. Computer Phys. Comm. **184**(4), 1241–1249 (2013)
15. Hanada M.: Markov chain Monte Carlo for dummies. arXiv:1808.08490 (2018)
16. Hennessy J.L., Patterson D.A.: A new golden age for computer architecture. Commun. ACM **62**(2), 48–60 (2019)
17. Kneusel R.T.: Random Numbers and Computers. Springer, Cham (CH) (2018)
18. Knuth D.E.: The Art of Computer Programming. Seminumerical Algorithms, vol. 2, 3rd edn. Addison-Wesley Longman, Boston, MA (1997)
19. L'Ecuyer P., Munger D., Oreshkin B.N., Simard R.J.: Random numbers for parallel computers: requirements and methods, with emphasis on GPUs. Math. Comput. Simul. **135**, 3–17 (2017)
20. Leiserson C.E., Schardl T.B., Sukha J.: Deterministic parallel random-number generation for dynamic-multithreading platforms. In: Proc. 17th Symposium on Principles and Practice of Parallel Programming, page 193–204. ACM, (2012)
21. Marques R., Paulino H., Alexandre F., Medeiros P.D.: Algorithmic skeleton framework for the orchestration of GPU computations. In: Euro-Par 2013 Parallel Processing, volume LNCS 8097, pages 874–885. Springer, (2013)
22. Mascagni M., Srinivasan A.: Algorithm 806: SPRNG: a scalable library for pseudorandom number generation. ACM Trans. Math. Softw. **26**(3), 436–461 (2000)
23. Passerat-Palmbach J., Mazel C., Hill D.R.C.: Pseudo-random number generation on GP-GPU. In: Proc. IEEE/ACM/SCS Workshop on Principles of Advanced and Distributed Simulation, pages 146–153, June (2011)
24. Perlin, K.: Improving noise. In: Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '02, pp. 681–682, New York, NY, USA, Association for Computing Machinery (2002)

25. Rabin M.O.: Probabilistic algorithms. In: Traub, J.F. (ed.) Algorithms and complexity, pp. 21–39. Academic Press, Cambridge (1976)
26. Rieger C., Wrede F., Kuchen H.: Musket: a domain-specific language for high-level parallel programming with algorithmic skeletons. In: Proc. Symposium on Applied Computing (SAC'19), pp. 1534–1543. ACM (2019)
27. Salmon J.K., Moraes M.A., Dror R.O., Shaw D.E.: Parallel random numbers: As easy as 1, 2, 3. In: Proc. Int. Conf. for High Performance Computing, Networking, Storage and Analysis, SC '11, pp. 16:1–16:12. ACM (2011)
28. Steuwer M., Kegel P., Gorlatch S.: SkelCL–a portable skeleton library for high-level GPU programming. In: 16th Int. Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS'11), pp. 1176–1182 (2011)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

August Ernstsson[1] · Nicolas Vandenbergen[2] · Jörg Keller[3] · Christoph Kessler[1]

✉ August Ernstsson
august.ernstsson@liu.se

Nicolas Vandenbergen
n.vandenbergen@fz-juelich.de

Jörg Keller
joerg.keller@fernuni-hagen.de

Christoph Kessler
christoph.kessler@liu.se

1 PELAB, Department of Computer and Information Science, Linköping University, Linköping, Sweden

2 Inst. f. Adv. Simulation, Jülich Supercomputing Center, FZ, Jülich, Germany

3 Faculty of Mathematics and Computer Science, FernUniversität in Hagen, Hagen, Germany