
GeantV Geometry:

SIMD abstraction and interfacing with CUDA

Johannes de Fine Licht (johannes.definelicht@cern.ch)

Sandro Wenzel (sandro.wenzel@cern.ch)

Thanks to Michal Husejko and Romain Wartel at TechLab for providing hardware.

Abstracted algorithms

- Write **common code** for multiple types of SIMD “backends” (Vc/Cilk/CUDA).
- Eases code maintenance; requires only one kernel to be written for a given functionality
- **Must retain performance of a low level implementation**

Illustrating scalar/SIMD abstraction and kernels

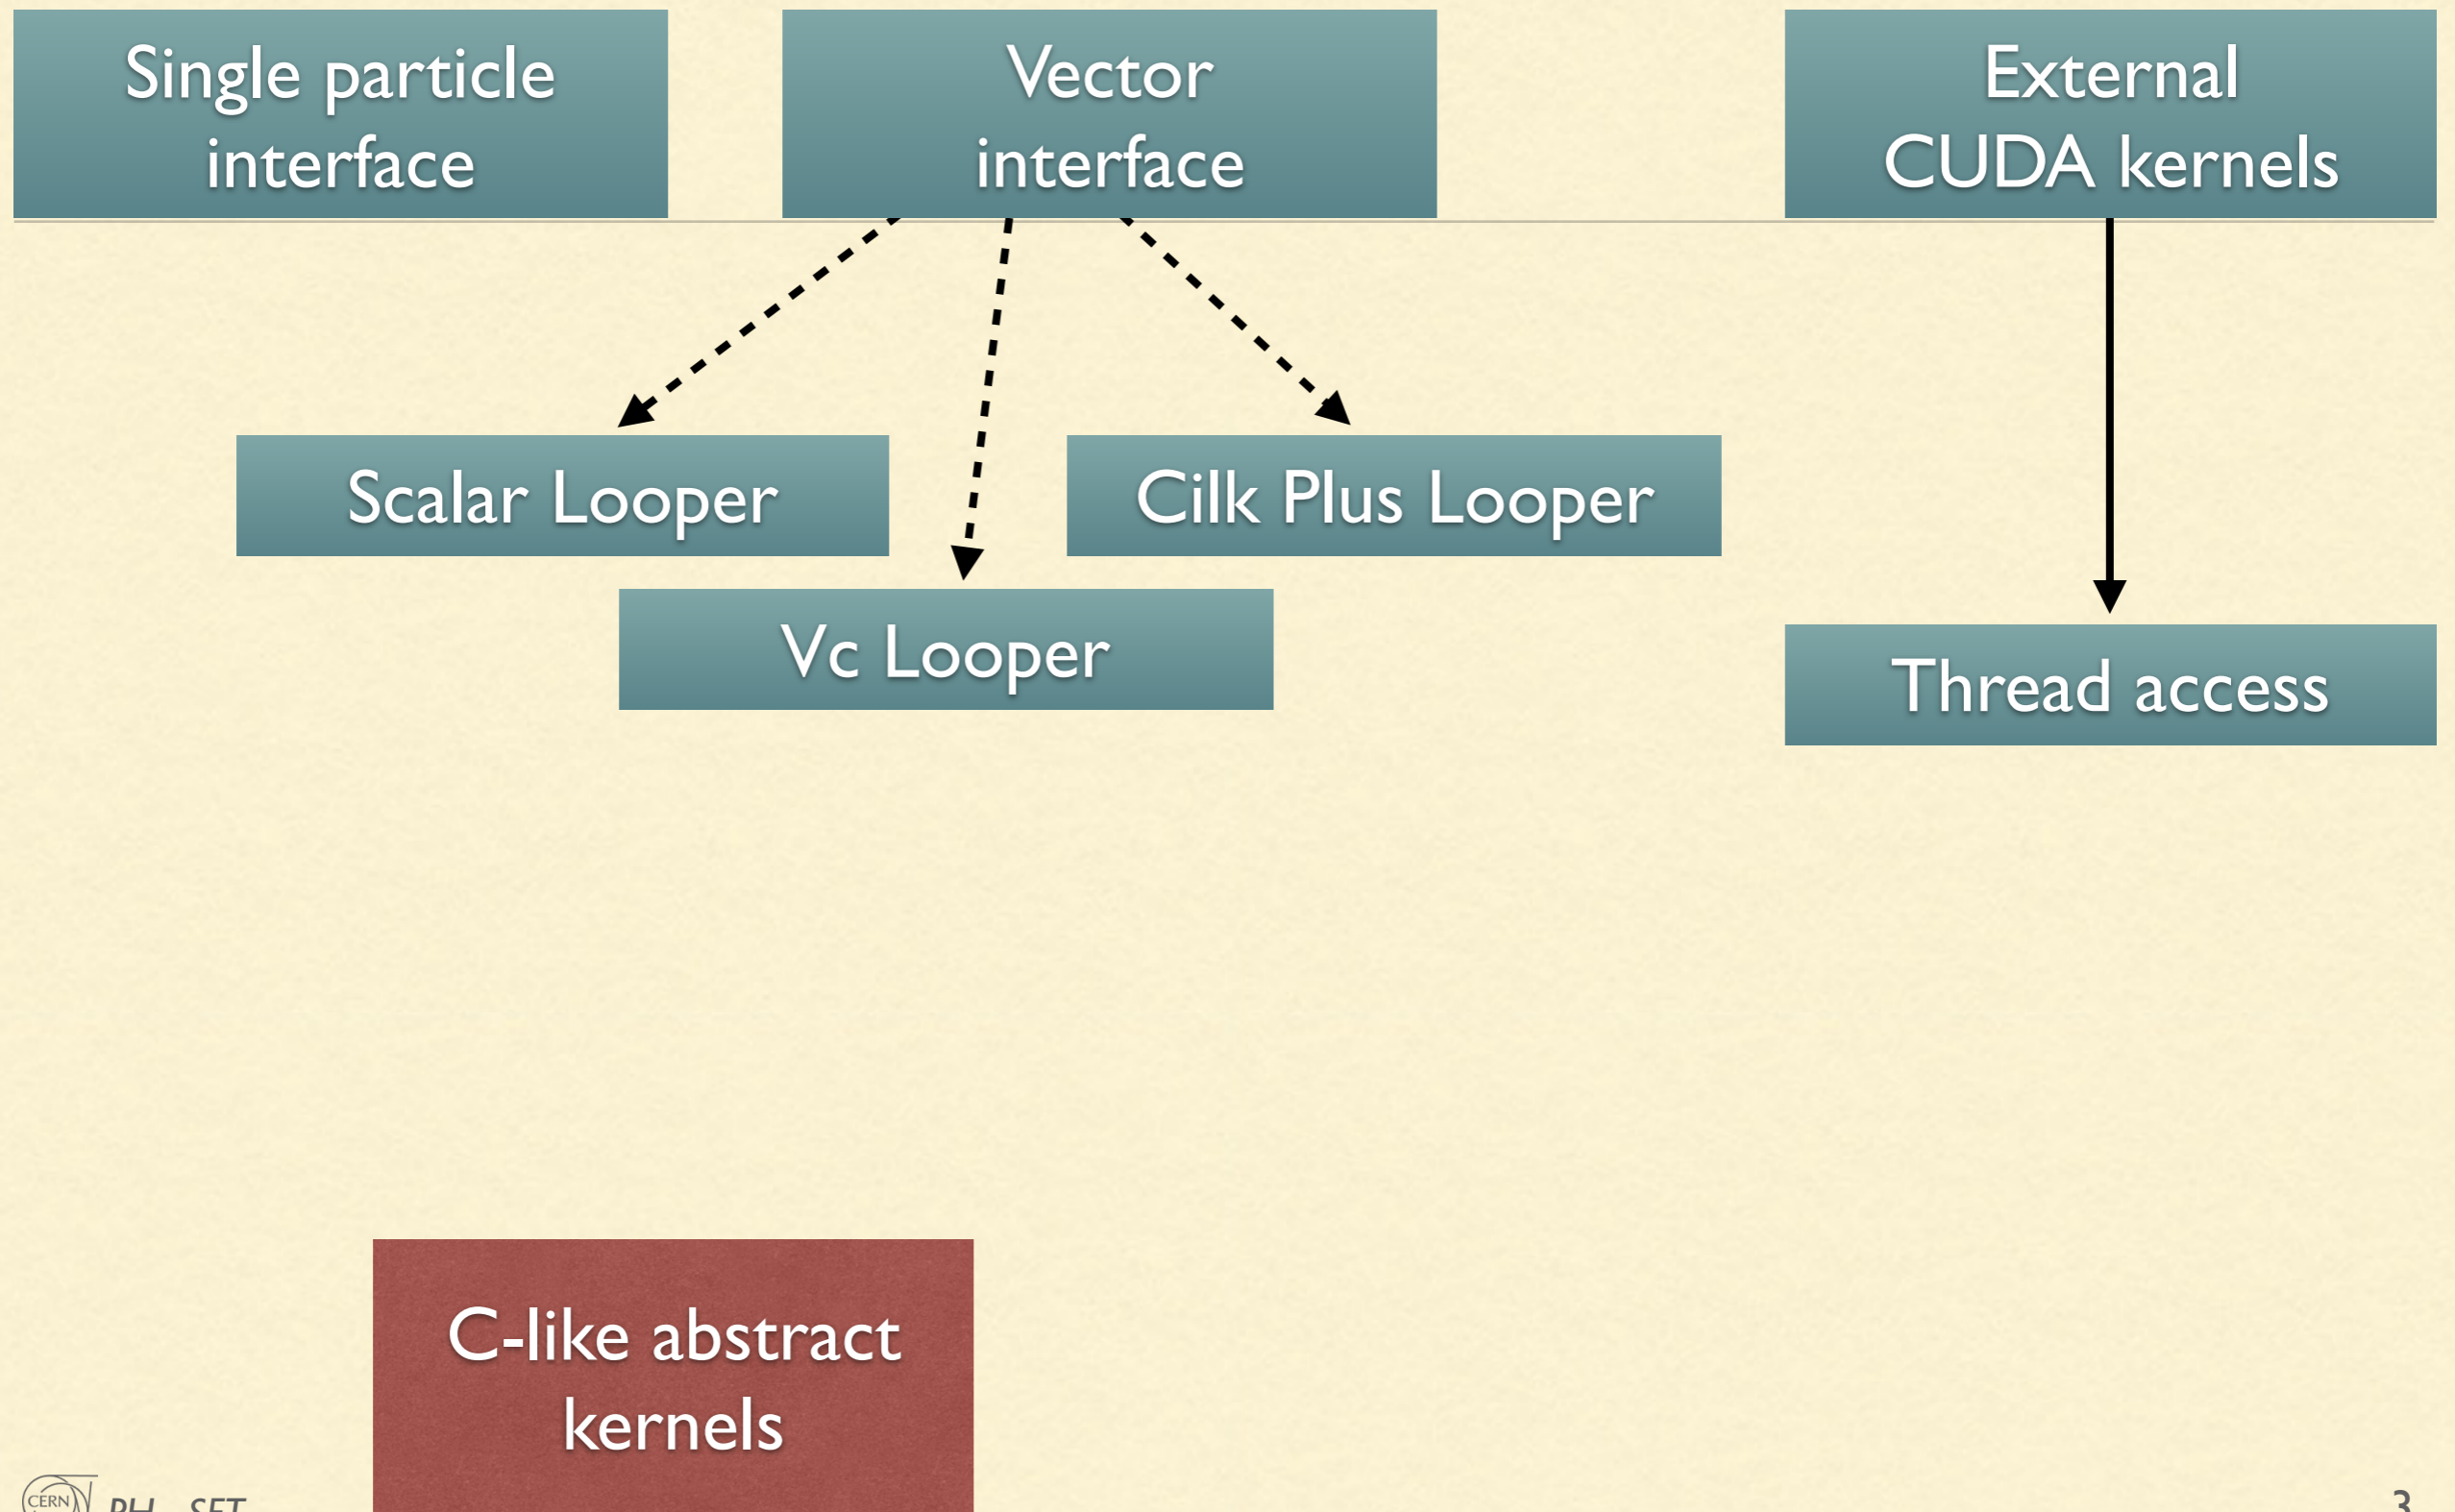
Single particle
interface

Vector
interface

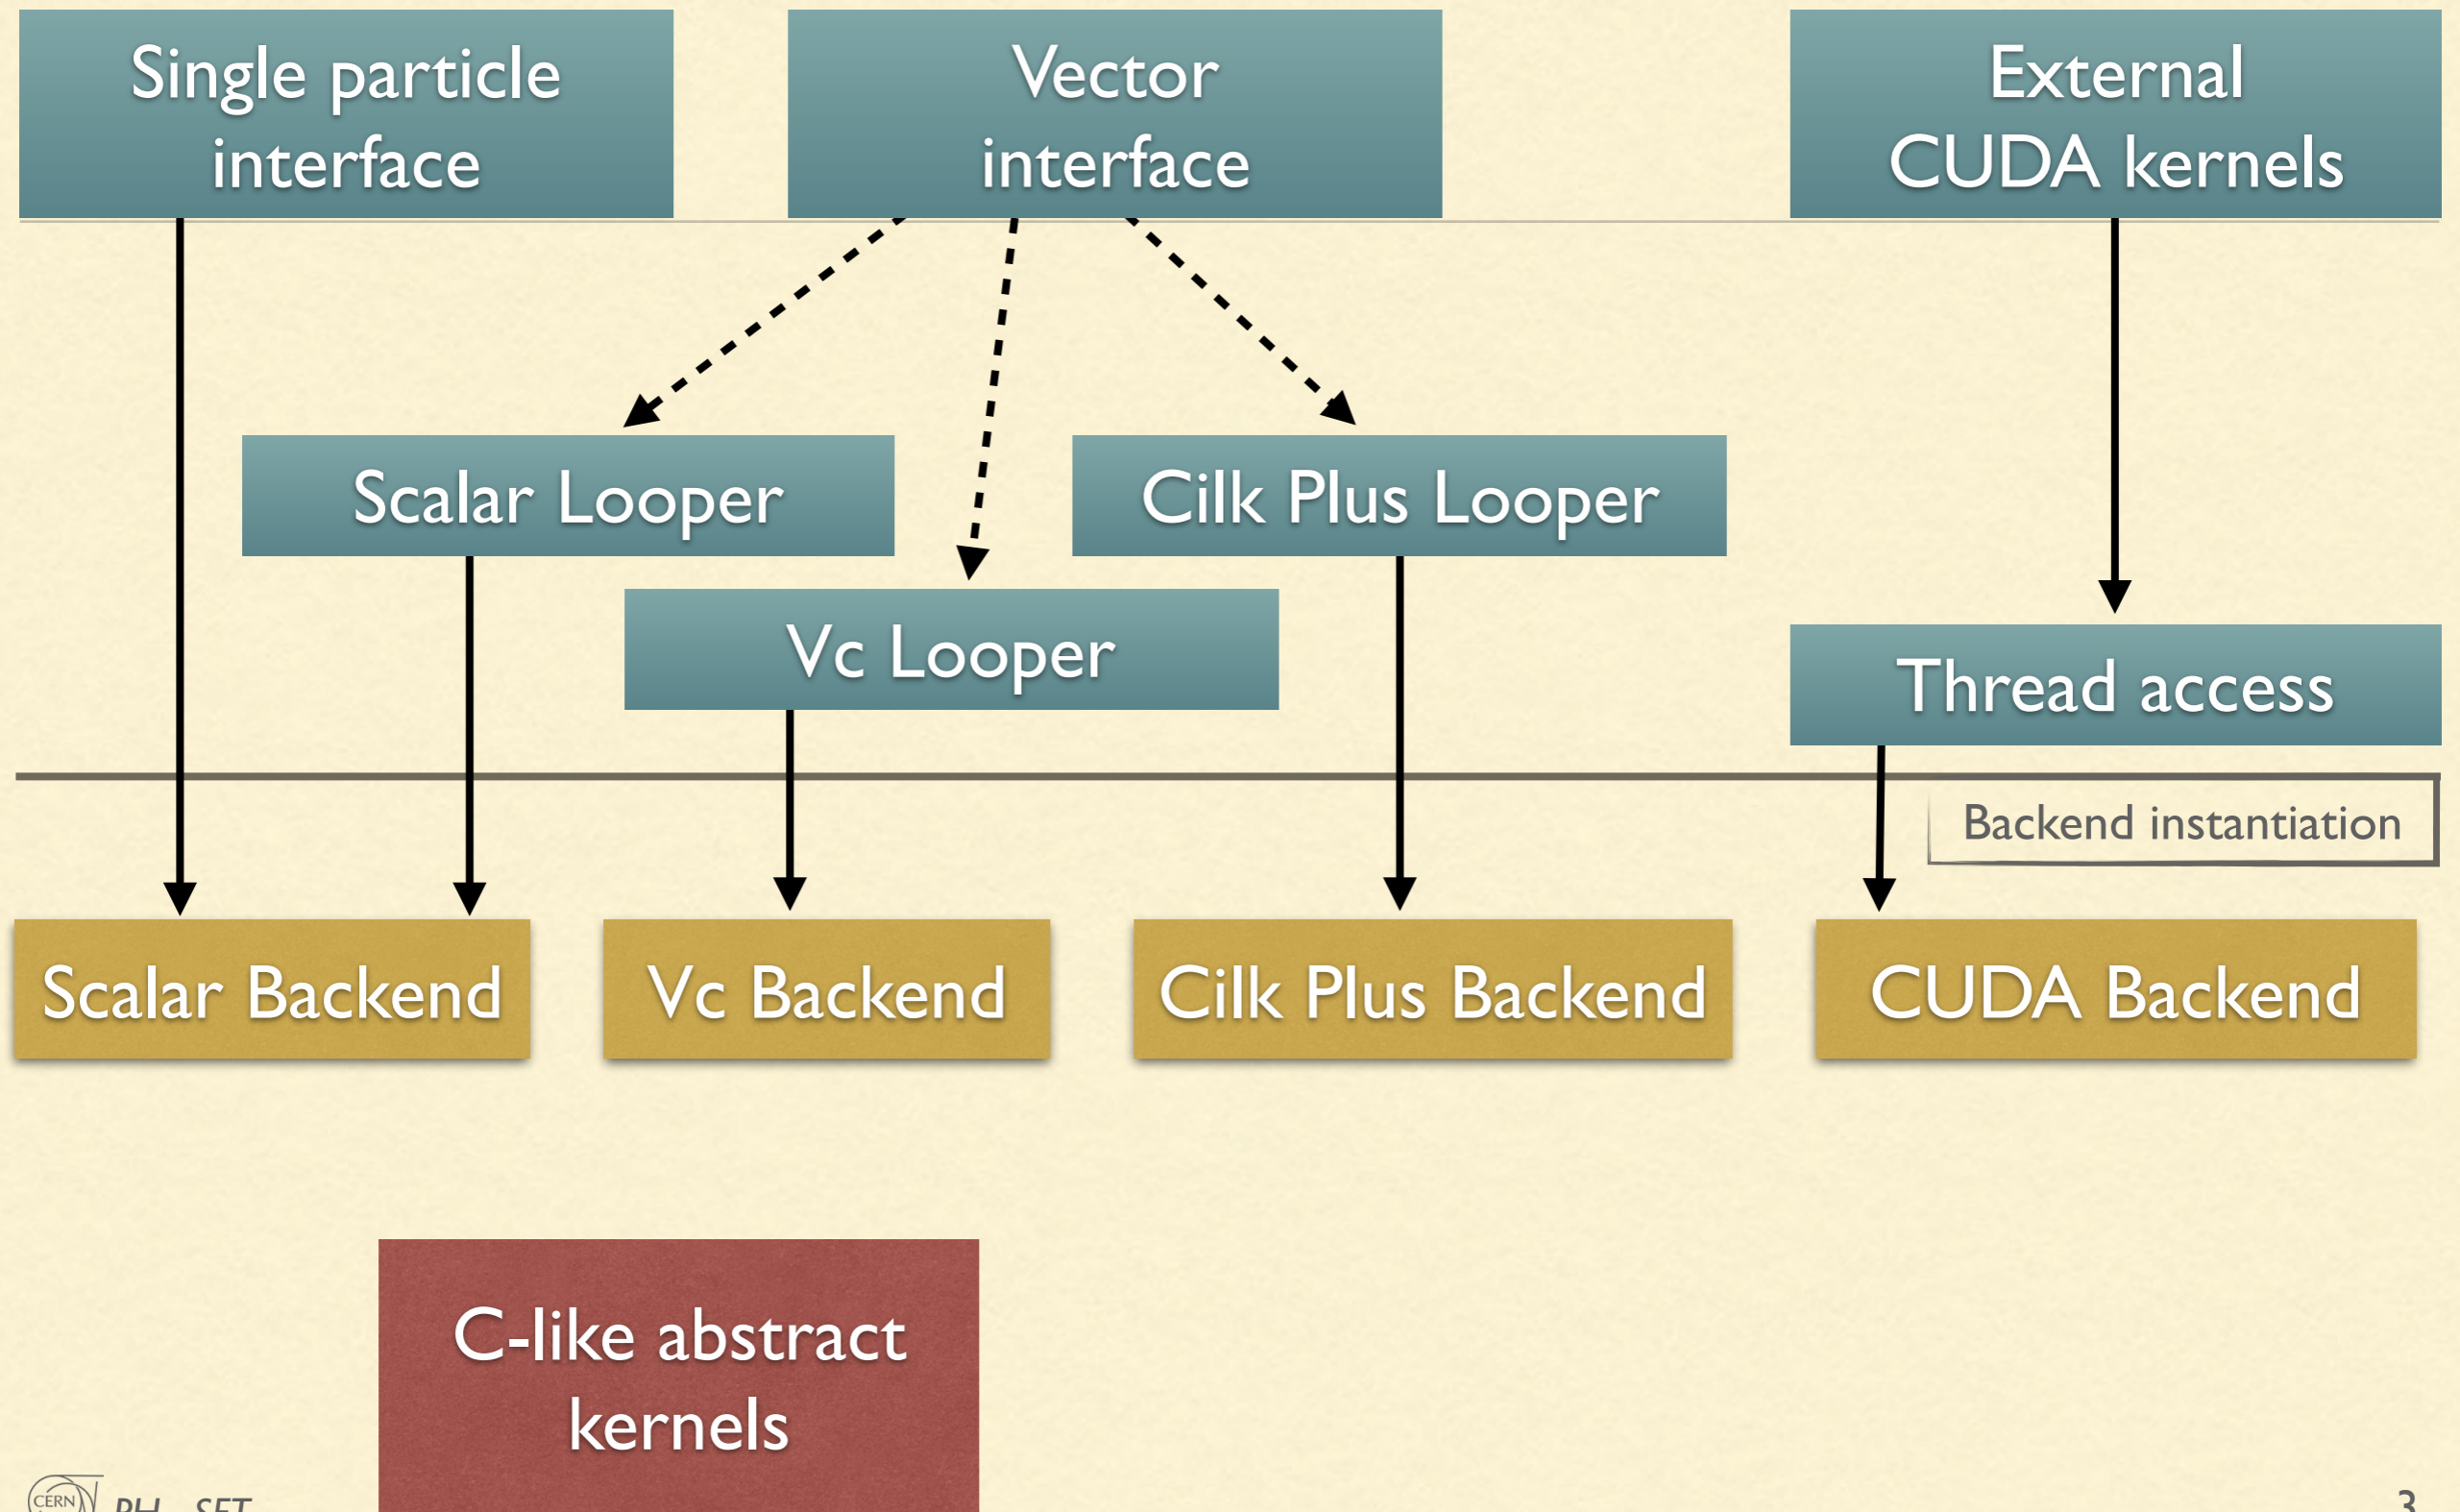
External
CUDA kernels

C-like abstract
kernels

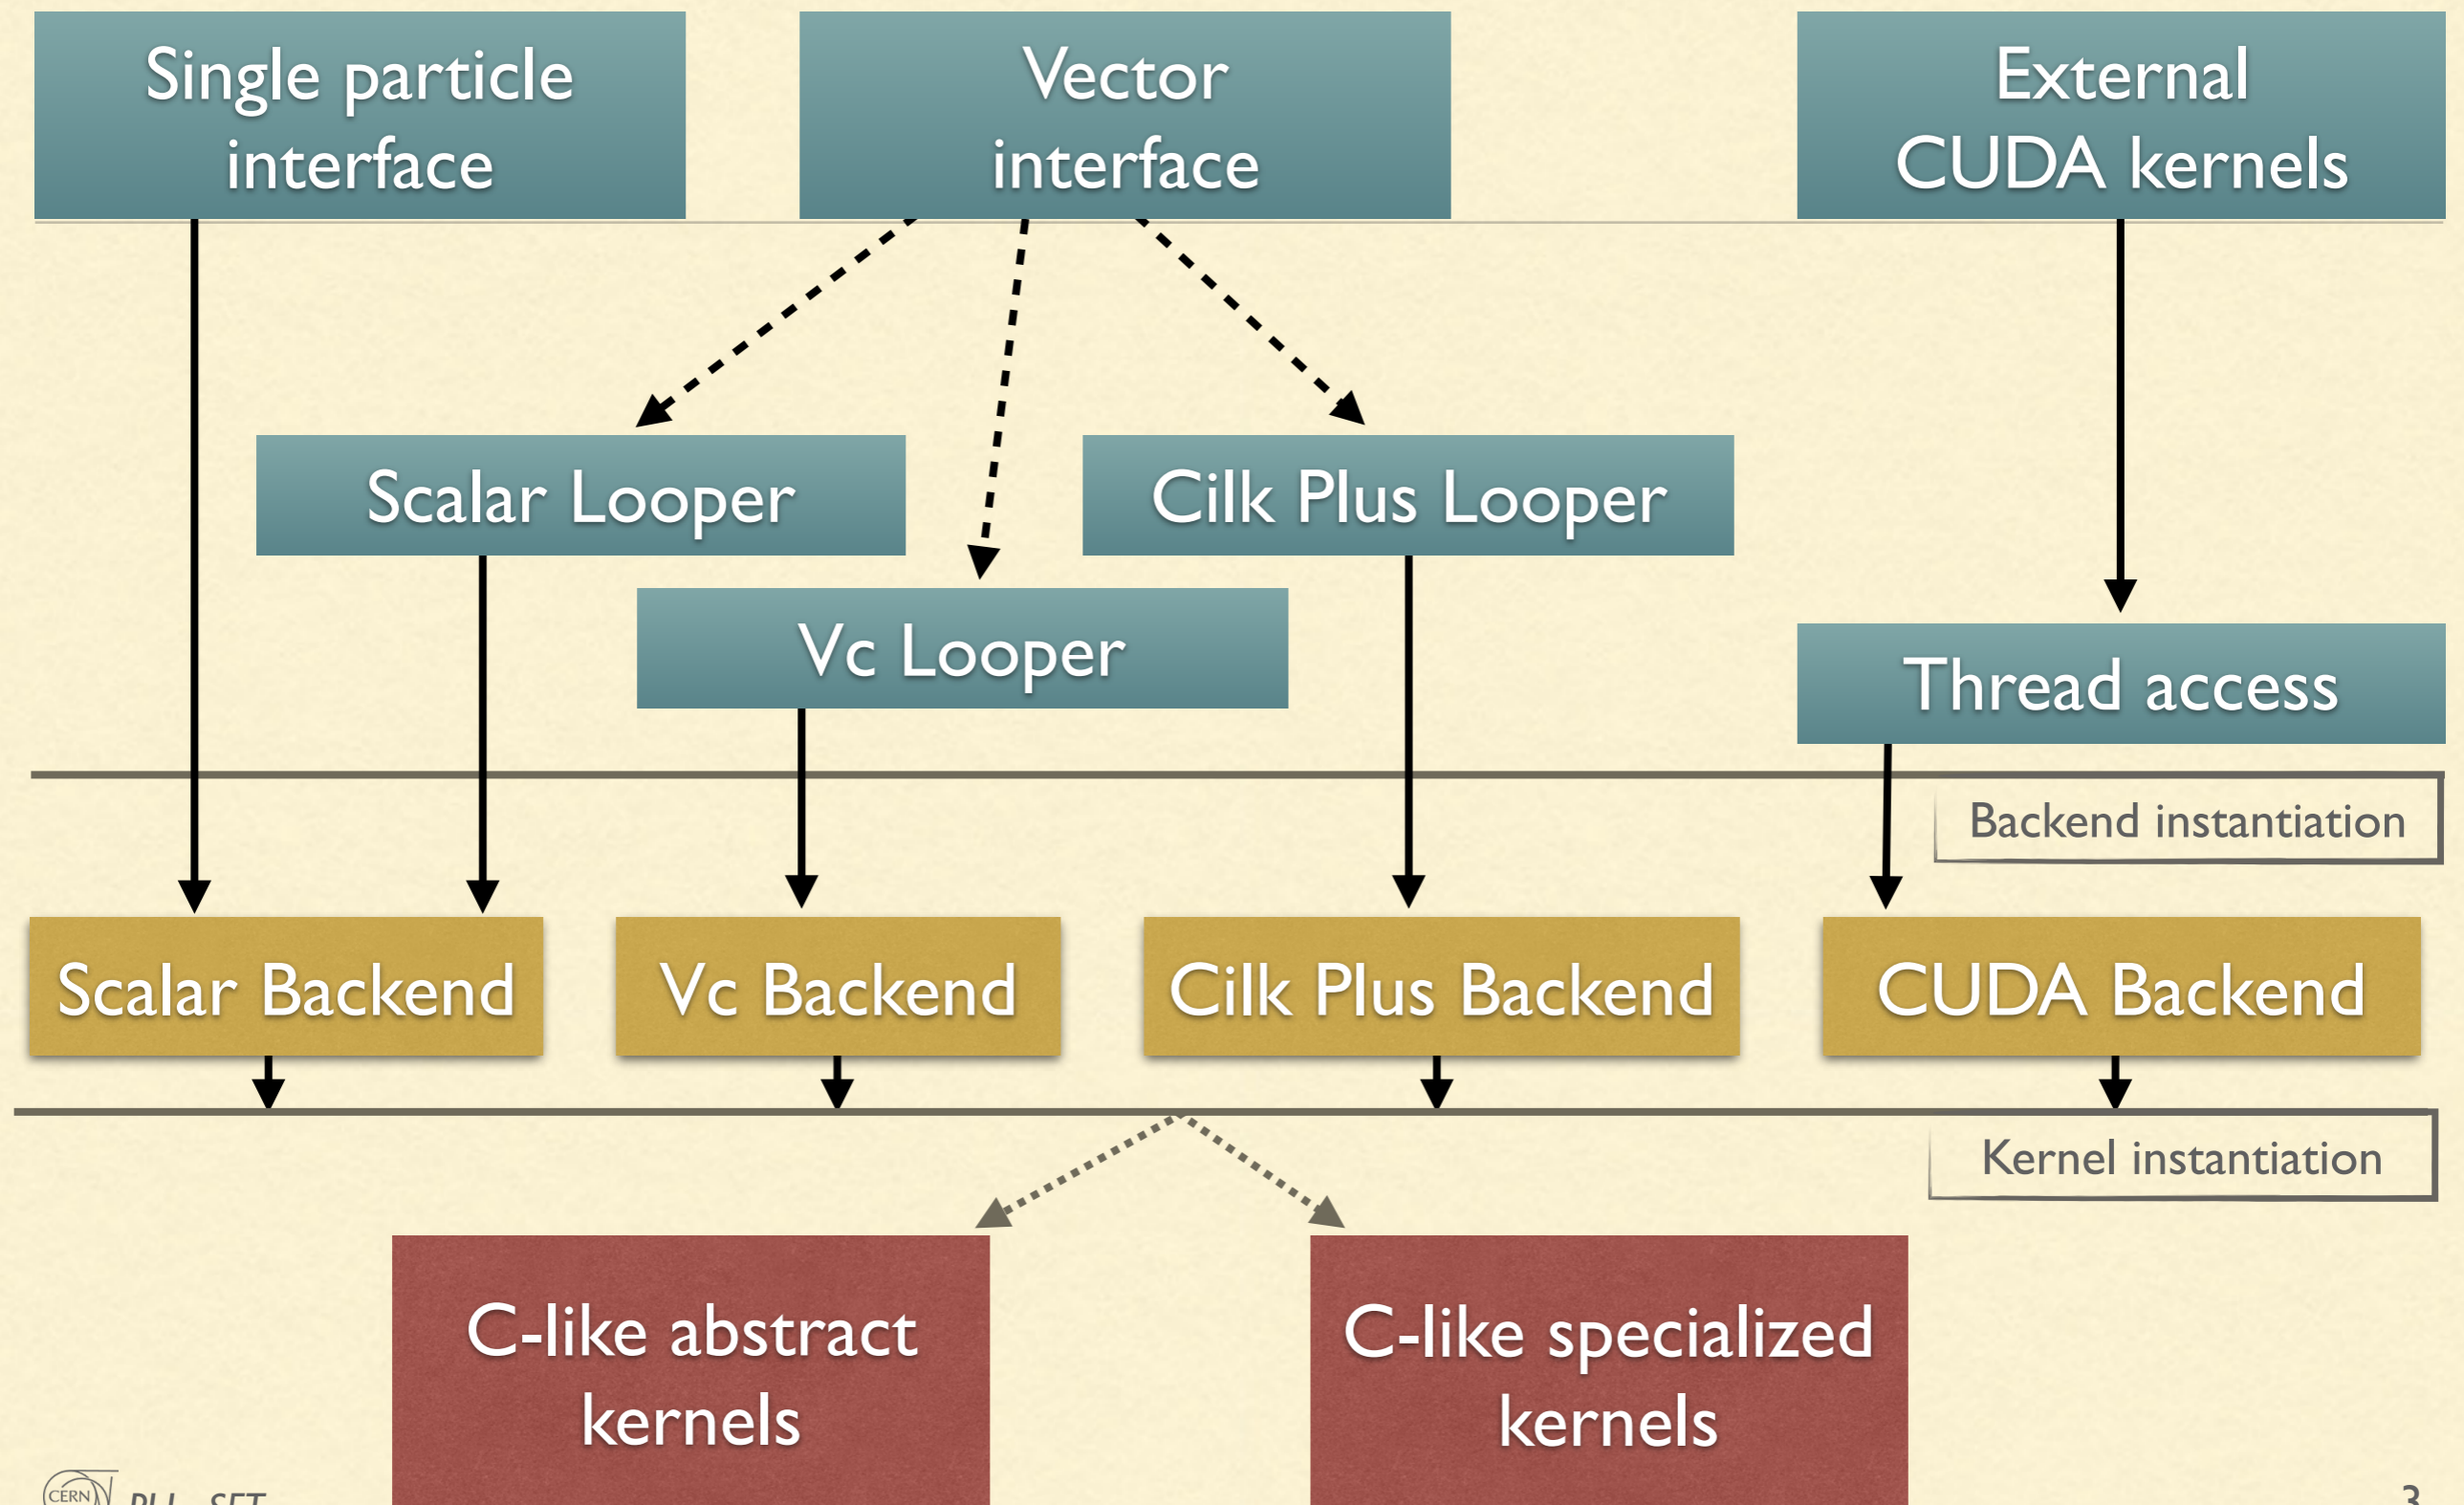
Illustrating scalar/SIMD abstraction and kernels



Illustrating scalar/SIMD abstraction and kernels



Illustrating scalar/SIMD abstraction and kernels



Generic kernels on abstract types

```
77 template <TranslationCode trans_code, RotationCode rot_code, typename Backend>
78 void BoxDistanceToIn(
```

- Generic programming; write algorithms in a **generic** way on **abstract** types
- Implementation depends on the backend
- Backend determines the types on which the high level operations are performed

```
86 typedef typename Backend::precision_v Float;
87 typedef typename Backend::bool_v Bool;
```

```
178 Float dist_x = (dimensions[0] - pos[0]) * dir_x_inv;
```

```
15 struct kVc {
16     typedef Vc::int_v int_v;
17     typedef Vc::Vector<Precision> precision_v;
18     typedef Vc::Vector<Precision>::Mask bool_v;
19     constexpr static bool early_returns = false;
20     const static precision_v kOne;
21     const static precision_v kZero;
22     const static bool_v kTrue;
23     const static bool_v kFalse;
24 };
```

Generic kernels on abstract types

```
77 template <TranslationCode trans_code, RotationCode rot_code, typename Backend>
78 void BoxDistanceToIn(
```

- Generic programming; write algorithms in a **generic** way on **abstract** types
- Implementation depends on the backend
- Backend determines the types on which the high level operations are performed

```
86 typedef typename Backend::precision_v Float;
87 typedef typename Backend::bool_v Bool;
```

```
178 Float dist_x = (dimensions[0] - pos[0]) * dir_x_inv;
```

```
15 struct kVc {
16     typedef Vc::int_v int_v;
17     typedef Vc::Vector<Precision> precision_v;
18     typedef Vc::Vector<Precision>::Mask bool_v;
19     constexpr static bool early_returns = false;
20     const static precision_v kOne;
21     const static precision_v kZero;
22     const static bool_v kTrue;
23     const static bool_v kFalse;
24 };
```


Implementations wrapped in types

- Operations are abstracted
- Backends implement all operations
- Vc already supports arithmetics
- CUDA uses primitive types
- Wrapper structs are implemented if needed

```
101 safety[0] = Abs(pos_local[0]) - dimensions[0];
```

```
48 VcPrecision Abs(VcPrecision const &val) {
49     return Vc::abs(val);
50 }
```

```
225 CilkPrecision Abs(CilkPrecision const &val) {
226     CilkPrecision result;
227     result.Map(std::fabs);
228     return result;
229 }
```

```
49 Type Abs(const Type val) {
50     return fabs(val);
51 }
```

```
50 template <typename Type, int vec_size>
51 struct CilkVector {
52
53     Type __attribute__((align(64))) vec[vec_size];
```

Interfacing with CUDA

- Involving the GPU should be easy
- Principle of common code
- *Issue:* GPU code needs nvcc, but we want a different compiler for CPU

```
31 CreateRootGeometry();  
32  
33 RootManager::Instance().LoadRootGeometry();  
34 CudaManager::Instance().LoadGeometry();  
35 CudaManager::Instance().Synchronize();
```

GNU/Clang:
C++11
Vc
Cilk

vs.

nvcc:
CUDA

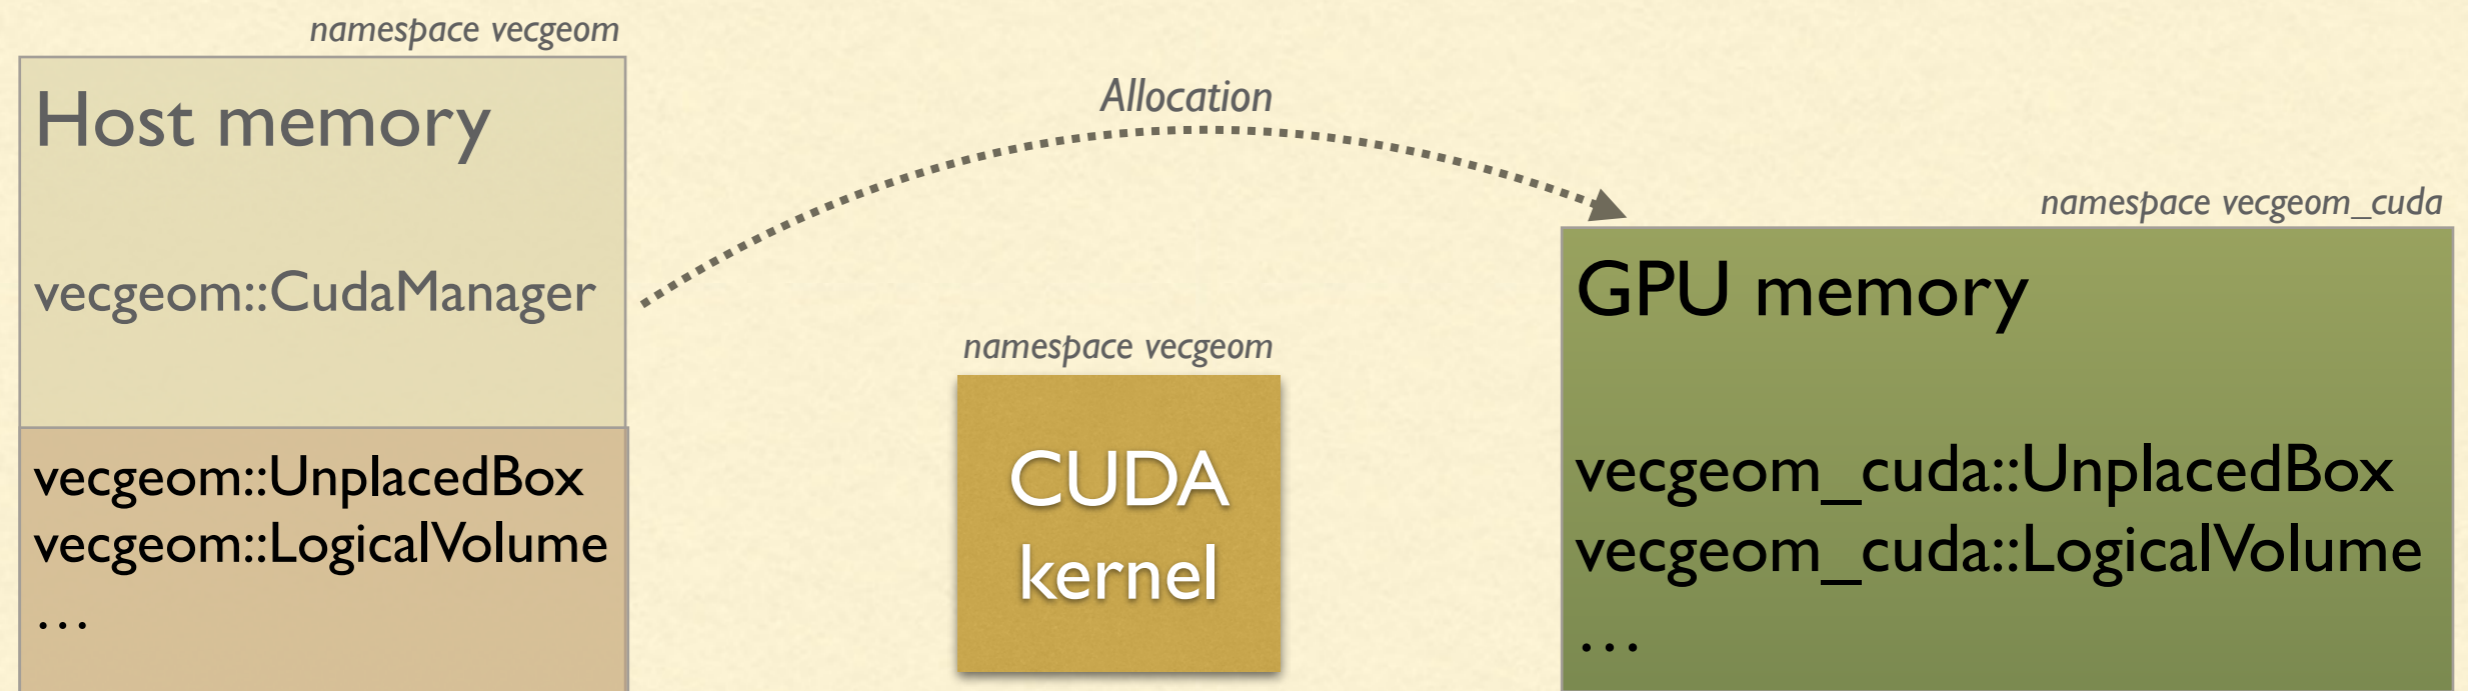
Dual namespaces

- Utilize distinct namespaces but **same code** for CPU and GPU
- Classes live in **one or both** environments
- Provide abstraction to interface between environments



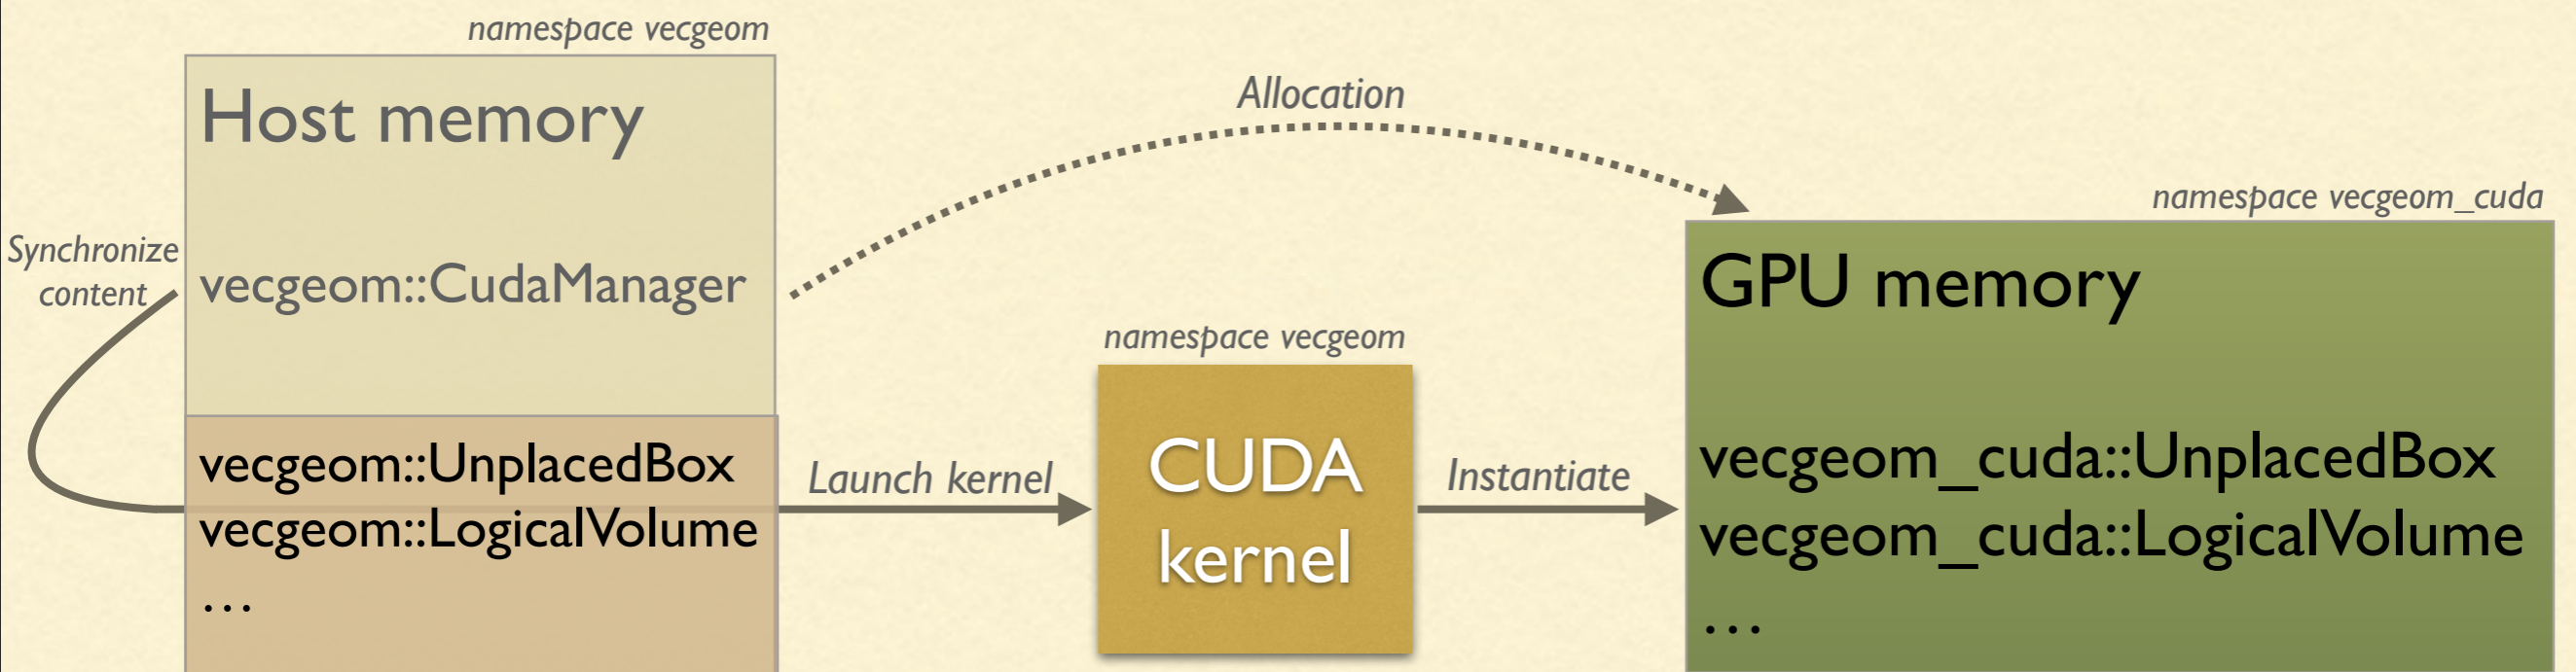
Dual namespaces

- Utilize distinct namespaces but **same code** for CPU and GPU
- Classes live in **one or both** environments
- Provide abstraction to interface between environments

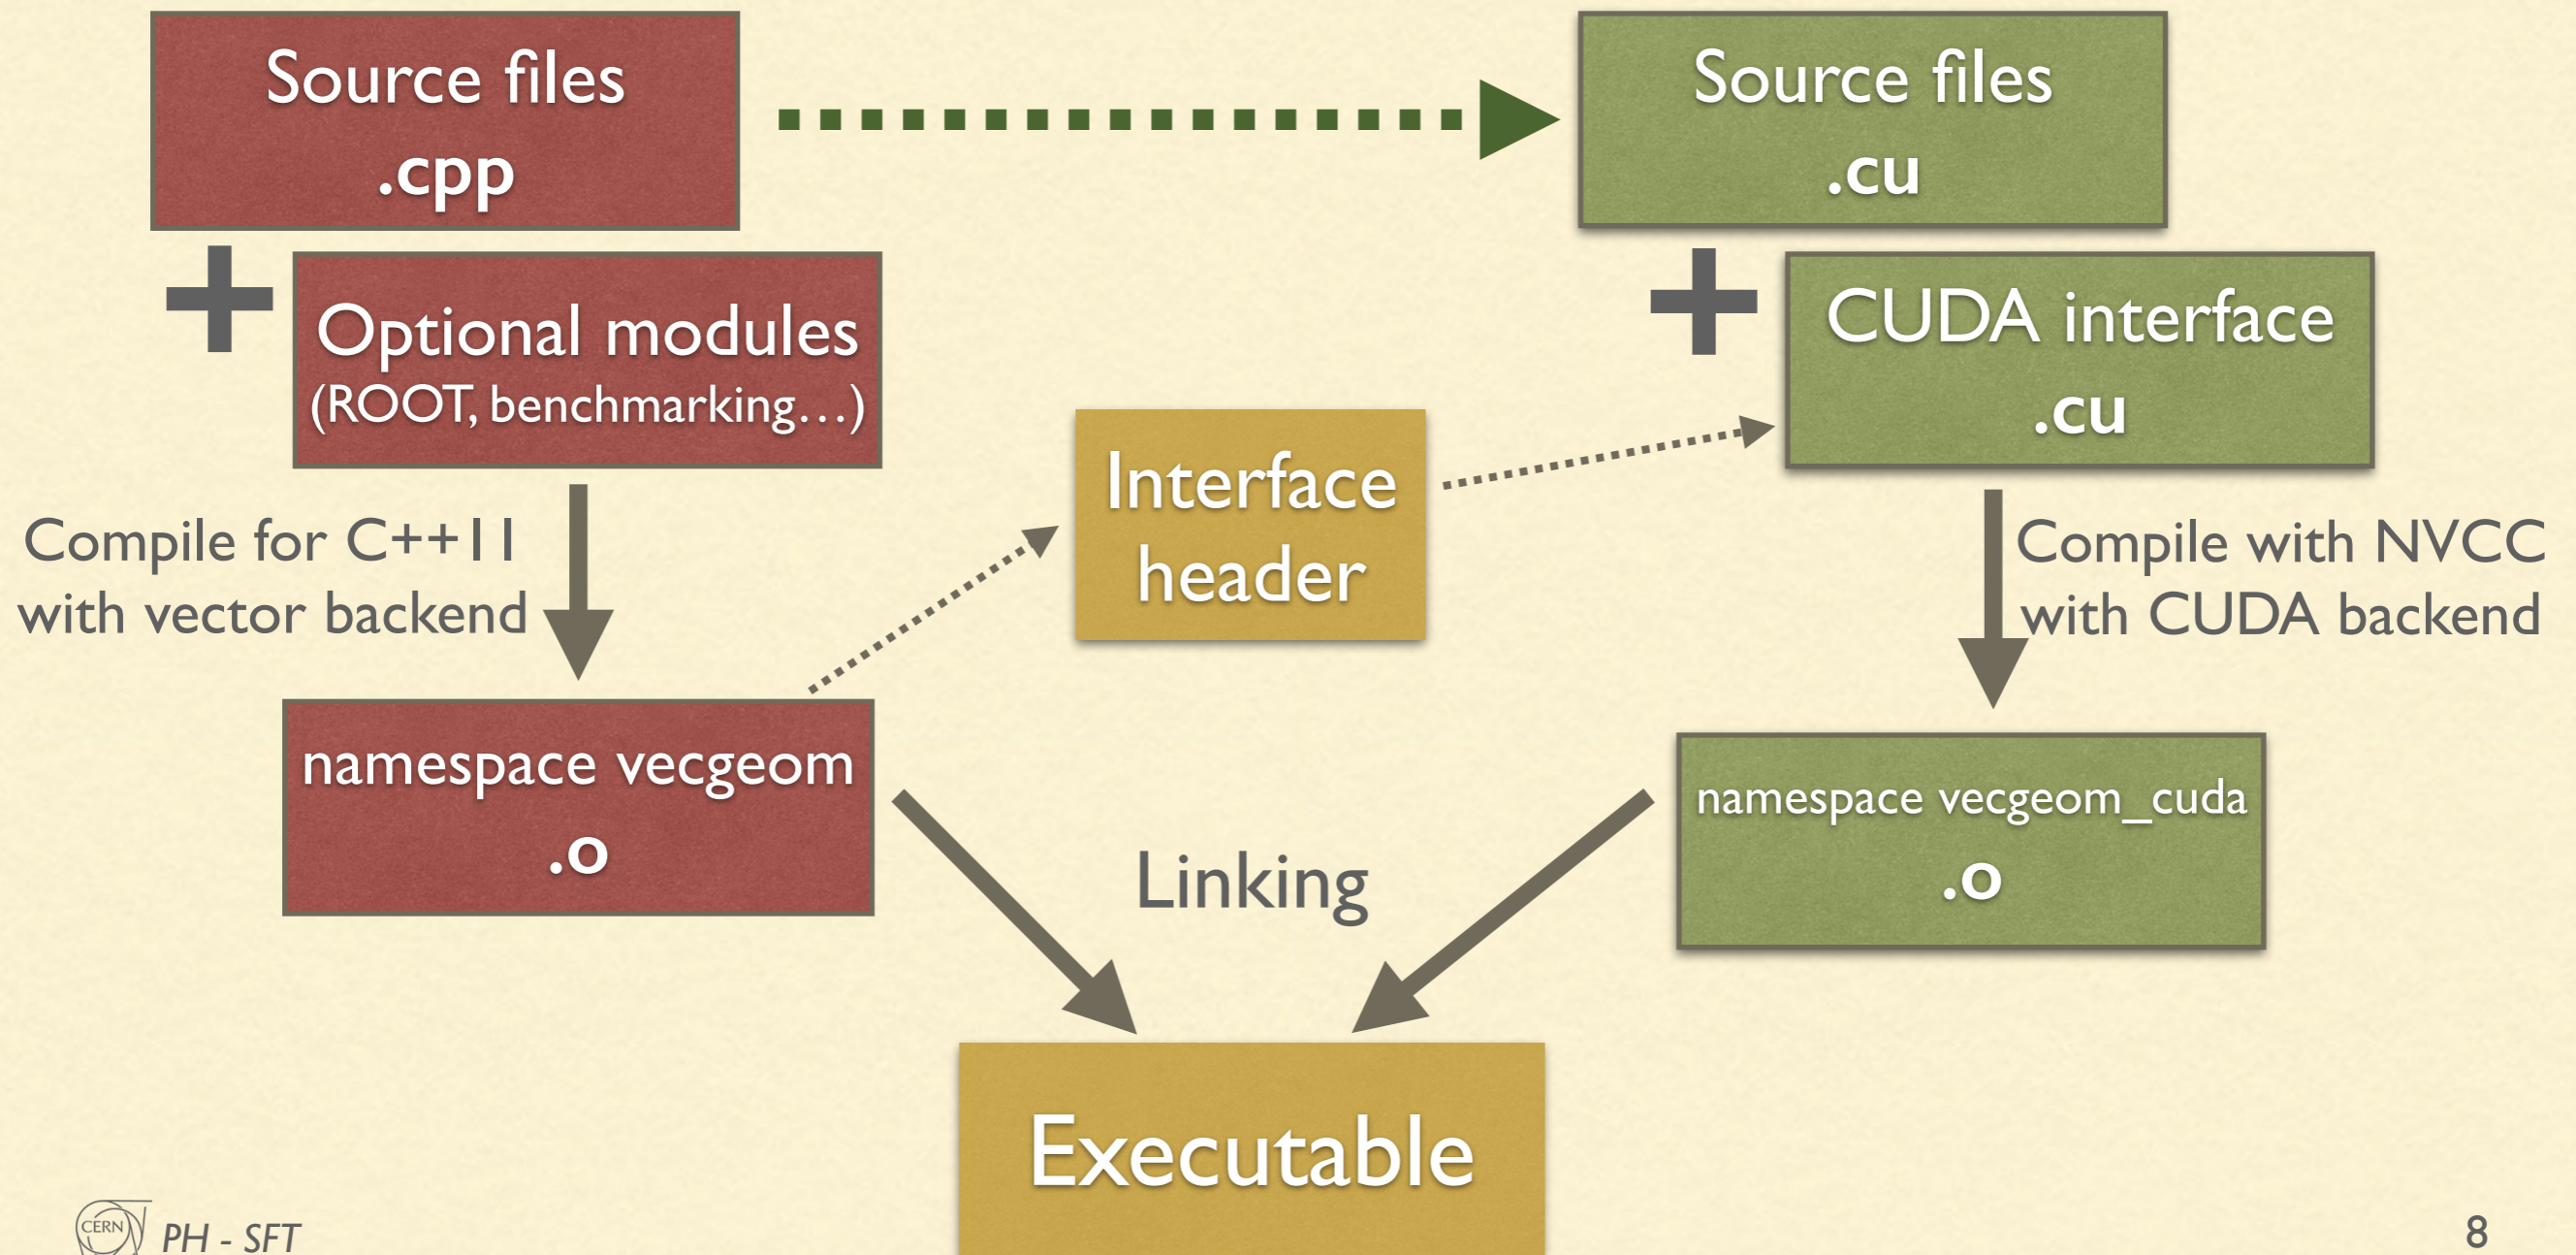


Dual namespaces

- Utilize distinct namespaces but **same code** for CPU and GPU
- Classes live in **one or both** environments
- Provide abstraction to interface between environments

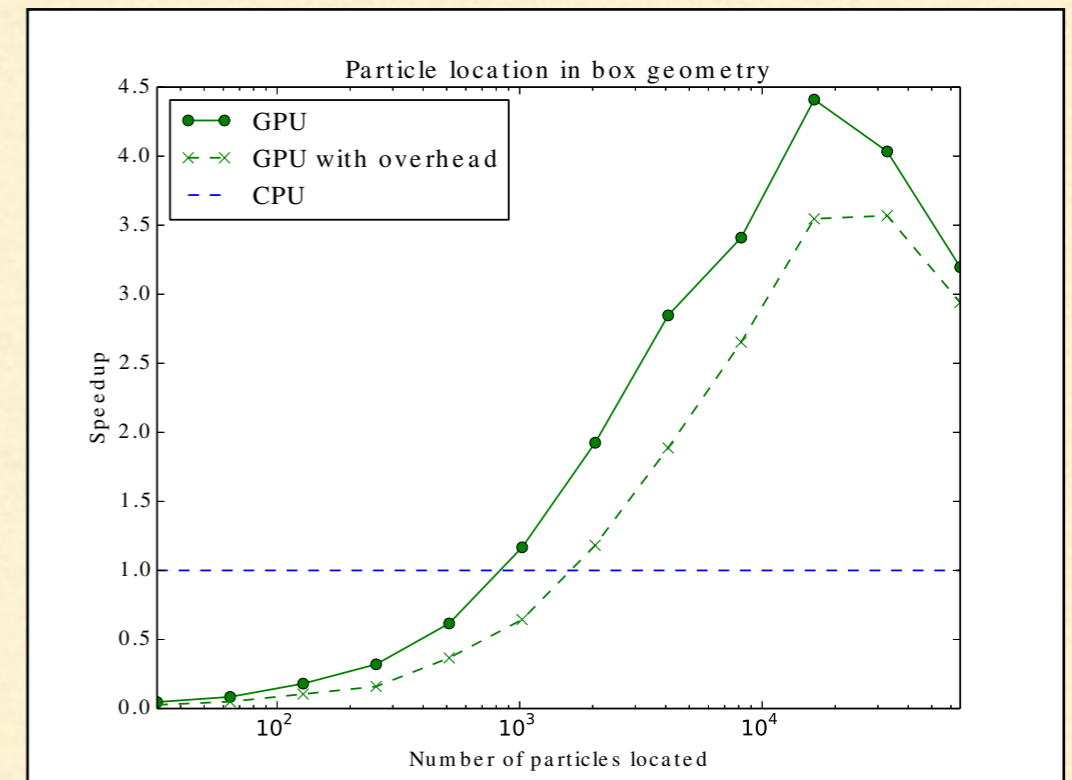


Compilation scheme



Status

- Common code for box methods runs for scalar, Vc, Cilk and CUDA backends
- Implemented GPU memory synchronization
- Dual namespaces allows dispatching work to CPU or GPU
- Location can run in either environment
- **Results are in agreement**



Preliminary benchmarks run in hybrid CPU/GPU environment for four-level box geometry. Location of particles is a tree algorithm, so speedups are only seen at very high particle multiplicity.

Appendix: CPU/GPU same interface

Main executable
compiled with
C++ compiler

```
50▼ for (int i = 0; i < n; ++i) {
51     NavigationState path(depth);
52▼     results[i] =
53         navigator.LocatePoint(GeoManager::Instance().world(), points[i], path,
54                               true)->id();
55 }
```

Same code used

```
75 template <typename TrackContainer>
76 __global__
77▼ void CudaManagerLocatePointsKernel(
78     vecgeom_cuda::VPlacedVolume const *const world,
79     vecgeom_cuda::SimpleNavigator const *const navigator,
80     vecgeom_cuda::NavigationState *const paths,
81     TrackContainer const *const points, const int n,
82     int *const output) {
83     const int i = vecgeom_cuda::ThreadIndex();
84     if (i >= n) return; // Out of range
85     output[i] =
86         navigator->LocatePoint(world, (*points)[i], paths[i], true)->id();
87 }
```

CUDA kernel
compiled with
NVCC

Appendix: Common code

```
17  template<typename Backend>
18  VECGEOM_INLINE
19  VECGEOM_CUDA_HEADER_BOTH
20  void BoxUnplacedInside( Vector3D<Precision> const & dimensions,
21                          Vector3D<typename Backend::precision_v> const &localpoint,
22                          typename Backend::bool_v *const inside )
23  {
24  Vector3D<typename Backend::bool_v> inside_dim(Backend::kFalse);
25  for (int i = 0; i < 3; ++i) {
26  inside_dim[i] = Abs(localpoint[i]) < dimensions[i];
27  if (Backend::early_returns) {
28  if (!inside_dim[i]) {
29  *inside = Backend::kFalse;
30  return;
31  }
32  }
33  }
34
35  if (Backend::early_returns) {
36  *inside = Backend::kTrue;
37  } else {
38  *inside = inside_dim[0] && inside_dim[1] && inside_dim[2];
39  }
40 }
```

Appendix: Template hierarchy

```

134 template <TranslationCode trans_code, RotationCode rot_code>
135 void SpecializedBox<trans_code, rot_code>::DistanceToIn(
136     SOA3D<Precision> const &positions,
137     SOA3D<Precision> const &directions,
138     Precision const *const step_max,
139     Precision *const output) const {
140     DistanceToIn_Looper<trans_code, rot_code>(*this, positions, directions,
141                                             step_max, output);
142 }

```

```

37 template <TranslationCode trans_code, RotationCode rot_code,
38           typename VolumeType, typename ContainerType>
39 VECGEOM_INLINE
40 void VPlacedVolume::DistanceToIn_Looper(VolumeType const &volume,
41                                         ContainerType const &positions,
42                                         ContainerType const &directions,
43                                         Precision const *const step_max,
44                                         Precision *const output) {
45     for (int i = 0; i < positions.fillsize(); i += kVectorSize) {
46         const VcPrecision result =
47             volume.template DistanceToInDispatch<trans_code, rot_code, kvC>(
48             Vector3D<VcPrecision>(VcPrecision(&positions.ContainerType::x(i)),
49                                 VcPrecision(&positions.ContainerType::y(i)),
50                                 VcPrecision(&positions.ContainerType::z(i))),
51             Vector3D<VcPrecision>(VcPrecision(&directions.ContainerType::x(i)),
52                                 VcPrecision(&directions.ContainerType::y(i)),
53                                 VcPrecision(&directions.ContainerType::z(i))),
54             VcPrecision(&step_max[i])
55             );
56         result.store(&output[i]);
57     }
58 }

```

```

15 struct kvC {
16     typedef Vc::int_v      int_v;
17     typedef Vc::Vector<Precision> precision_v;
18     typedef Vc::Vector<Precision>::Mask bool_v;
19     constexpr static bool early_returns = false;
20     const static precision_v kOne;
21     const static precision_v kZero;
22     const static bool_v kTrue;
23     const static bool_v kFalse;
24 };

```

```

77 template <TranslationCode trans_code, RotationCode rot_code, typename Backend>
78 void BoxDistanceToIn(
79     Vector3D<Precision> const &dimensions,
80     TransformationMatrix const &matrix,
81     Vector3D<typename Backend::precision_v> const &pos,
82     Vector3D<typename Backend::precision_v> const &dir,
83     typename Backend::precision_v const &step_max,
84     typename Backend::precision_v *const distance) {
85
86     typedef typename Backend::precision_v Float;
87     typedef typename Backend::bool_v Bool;
88
89     Vector3D<Float> safety;
90     Vector3D<Float> pos_local;
91     Vector3D<Float> dir_local;
92     Bool hit(false);
93     Bool done(false);
94     *distance = kInfinity;
95
96     matrix.Transform<trans_code, rot_code>(pos, pos_local);
97     matrix.TransformRotation<rot_code>(dir, dir_local);

```

Vector interface

```

134 template <TranslationCode trans_code, RotationCode rot_code>
135 void SpecializedBox<trans_code, rot_code>::DistanceToIn(
136     SOA3D<Precision> const &positions,
137     SOA3D<Precision> const &directions,
138     Precision const *const step_max,
139     Precision *const output) const {
140     DistanceToIn_Looper<trans_code, rot_code>(*this, positions, directions,
141     step_max, output);
142 }

```

```

33 template <TranslationCode trans_code, RotationCode rot_code,
34     typename VolumeType, typename ContainerType>
35 VECGEOM_INLINE
36 void VPlacedVolume::DistanceToIn_Looper(VolumeType const &volume,
37     ContainerType const &positions,
38     ContainerType const &directions,
39     Precision const *const step_max,
40     Precision *const output) {
41     for (int i = 0; i < positions.fillsize(); ++i) {
42         output[i] =
43             volume.template DistanceToInDispatch<trans_code, rot_code, kScalar>(
44             positions[i], directions[i], step_max[i]
45         );
46     }
47 }

```

```

37 template <TranslationCode trans_code, RotationCode rot_code,
38     typename VolumeType, typename ContainerType>
39 VECGEOM_INLINE
40 void VPlacedVolume::DistanceToIn_Looper(VolumeType const &volume,
41     ContainerType const &positions,
42     ContainerType const &directions,
43     Precision const *const step_max,
44     Precision *const output) {
45     for (int i = 0; i < positions.fillsize(); i += kVectorSize) {
46         const VcPrecision result =
47             volume.template DistanceToInDispatch<trans_code, rot_code, kVc>(
48             Vector3D<VcPrecision>(VcPrecision(&positions.ContainerType::x(i)),
49             VcPrecision(&positions.ContainerType::y(i)),
50             VcPrecision(&positions.ContainerType::z(i))),
51             Vector3D<VcPrecision>(VcPrecision(&directions.ContainerType::x(i)),
52             VcPrecision(&directions.ContainerType::y(i)),
53             VcPrecision(&directions.ContainerType::z(i))),
54             VcPrecision(&step_max[i]
55         );
56         result.store(&output[i]);
57     }
58 }

```

```

36 template <TranslationCode trans_code, RotationCode rot_code,
37     typename VolumeType, typename ContainerType>
38 VECGEOM_INLINE
39 void VPlacedVolume::DistanceToIn_Looper(VolumeType const &volume,
40     ContainerType const &positions,
41     ContainerType const &directions,
42     Precision const *const step_max,
43     Precision *const output) {
44     for (int i = 0; i < positions.fillsize(); i += CilkPrecision::size()) {
45         const CilkPrecision result =
46             volume.template DistanceToInDispatch<trans_code, rot_code,
47             kCilk>(
48             Vector3D<CilkPrecision>(CilkPrecision(&positions.x(i)),
49             CilkPrecision(&positions.y(i)),
50             CilkPrecision(&positions.z(i))),
51             Vector3D<CilkPrecision>(CilkPrecision(&directions.x(i)),
52             CilkPrecision(&directions.y(i)),
53             CilkPrecision(&directions.z(i))),
54             CilkPrecision(&step_max[i]
55         );
56         result.store(&output[i]);
57     }
58 }

```

```

15 struct kScalar {
16     typedef int int_v;
17     typedef Precision precision_v;
18     typedef bool bool_v;
19     const static bool early_returns = true;
20     const static precision_v kOne;
21     const static precision_v kZero;
22     const static bool_v kTrue = true;
23     const static bool_v kFalse = false;
24 };

```

```

15 struct kVc {
16     typedef Vc::int_v int_v;
17     typedef Vc::Vector<Precision> precision_v;
18     typedef Vc::Vector<Precision>::Mask bool_v;
19     const static bool early_returns = false;
20     const static precision_v kOne;
21     const static precision_v kZero;
22     const static bool_v kTrue;
23     const static bool_v kFalse;
24 };

```

```

31 struct kCilk {
32     typedef CilkVector<int> int_v;
33     typedef CilkVector<Precision> precision_v;
34     typedef CilkVector<bool> bool_v;
35     const static bool early_returns = false;
36     const static precision_v kOne;
37     const static precision_v kZero;
38     const static bool_v kTrue;
39     const static bool_v kFalse;
40 };

```

```

227 template <TranslationCode trans_code, RotationCode rot_code, typename Backend>
228 VECGEOM_CUDA_HEADER_BOTH
229 VECGEOM_INLINE
230 typename Backend::precision_v PlacedBox::DistanceToInDispatch(
231     Vector3D<typename Backend::precision_v> const &position,
232     Vector3D<typename Backend::precision_v> const &direction,
233     const typename Backend::precision_v step_max) const {
234
235     typename Backend::precision_v output;
236     BoxDistanceToIn<trans_code, rot_code, Backend>(
237         unplaced_box()->dimensions(),
238         *this->matrix(),
239         position,
240         direction,
241         step_max,
242         &output
243     );
244     return output;
245 }

```

```

17 struct kCuda {
18     typedef int int_v;
19     typedef Precision precision_v;
20     typedef bool bool_v;
21     const static bool early_returns = false;
22     const static precision_v kOne = 1.0;
23     const static precision_v kZero = 0.0;
24     const static bool_v kTrue = true;
25     const static bool_v kFalse = false;
26 };

```

External CUDA kernels

```

77 template <TranslationCode trans_code, RotationCode rot_code, typename Backend>
78 void BoxDistanceToIn(
79     Vector3D<Precision> const &dimensions,
80     TransformationMatrix const &matrix,
81     Vector3D<typename Backend::precision_v> const &pos,
82     Vector3D<typename Backend::precision_v> const &dir,
83     typename Backend::precision_v const &step_max,
84     typename Backend::precision_v *const distance) {
85
86     typedef typename Backend::precision_v Float;
87     typedef typename Backend::bool_v Bool;
88
89     Vector3D<Float> safety;
90     Vector3D<Float> pos_local;
91     Vector3D<Float> dir_local;
92     Bool hit(false);
93     Bool done(false);
94     *distance = kInfinity;
95
96     matrix.Transform<trans_code, rot_code>(pos, pos_local);
97     matrix.TransformRotation<rot_code>(dir, dir_local);

```

C-like abstract kernels

Vector interface

```

134 template <TranslationCode trans_code, RotationCode rot_code>
135 void SpecializedBox<trans_code, rot_code>::DistanceToIn(
136     SOA3D<Precision> const &positions,
137     SOA3D<Precision> const &directions,
138     Precision const *const step_max,
139     Precision *const output) const {
140     DistanceToIn_Looper<trans_code, rot_code>(*this, positions, directions,
141     step_max, output);
142 }
    
```

```

33 template <TranslationCode trans_code, RotationCode rot_code,
34     typename VolumeType, typename ContainerType>
35 VECGEOM_INLINE
36 void VPlacedVolume::DistanceToIn_Looper(VolumeType const &volume,
37     ContainerType const &positions,
38     ContainerType const &directions,
39     Precision const *const step_max,
40     Precision *const output) {
41     for (int i = 0; i < positions.fillsize(); ++i) {
42         output[i] =
43             volume.template DistanceToInDispatch<trans_code, rot_code, kScalar>(
44             positions[i], directions[i], step_max[i]
45             );
46     }
47 }
    
```

```

37 template <TranslationCode trans_code, RotationCode rot_code,
38     typename VolumeType, typename ContainerType>
39 VECGEOM_INLINE
40 void VPlacedVolume::DistanceToIn_Looper(VolumeType const &volume,
41     ContainerType const &positions,
42     ContainerType const &directions,
43     Precision const *const step_max,
44     Precision *const output) {
45     for (int i = 0; i < positions.fillsize(); i += kVectorSize) {
46         const VcPrecision result =
47             volume.template DistanceToInDispatch<trans_code, rot_code, kVc>(
48             Vector3D<VcPrecision>(VcPrecision(&positions.ContainerType::x(i)),
49             VcPrecision(&positions.ContainerType::y(i)),
50             VcPrecision(&positions.ContainerType::z(i))),
51             Vector3D<VcPrecision>(VcPrecision(&directions.ContainerType::x(i)),
52             VcPrecision(&directions.ContainerType::y(i)),
53             VcPrecision(&directions.ContainerType::z(i))),
54             VcPrecision(&step_max[i]
55             );
56         result.store(&output[i]);
57     }
58 }
    
```

```

36 template <TranslationCode trans_code, RotationCode rot_code,
37     typename VolumeType, typename ContainerType>
38 VECGEOM_INLINE
39 void VPlacedVolume::DistanceToIn_Looper(VolumeType const &volume,
40     ContainerType const &positions,
41     ContainerType const &directions,
42     Precision const *const step_max,
43     Precision *const output) {
44     for (int i = 0; i < positions.fillsize(); i += CilkPrecision::size()) {
45         const CilkPrecision result =
46             volume.template DistanceToInDispatch<trans_code, rot_code,
47             kCilk>(
48             Vector3D<CilkPrecision>(CilkPrecision(&positions.x(i)),
49             CilkPrecision(&positions.y(i)),
50             CilkPrecision(&positions.z(i))),
51             Vector3D<CilkPrecision>(CilkPrecision(&directions.x(i)),
52             CilkPrecision(&directions.y(i)),
53             CilkPrecision(&directions.z(i))),
54             CilkPrecision(&step_max[i]
55             );
56         result.store(&output[i]);
57     }
58 }
    
```

Scalar Backend

```

15 struct kScalar {
16     typedef int int_v;
17     typedef Precision precision_v;
18     typedef bool bool_v;
19     const static bool early_returns = true;
20     const static precision_v kOne;
21     const static precision_v kZero;
22     const static bool_v kTrue = true;
23     const static bool_v kFalse = false;
24 };
    
```

Vc Backend

```

15 struct kVc {
16     typedef Vc::int_v int_v;
17     typedef Vc::precision precision_v;
18     typedef Vc::Vector<Precision>::Mask bool_v;
19     const static bool early_returns = false;
20     const static precision_v kOne;
21     const static precision_v kZero;
22     const static bool_v kTrue;
23     const static bool_v kFalse;
24 };
    
```

Cilk Plus Backend

```

31 struct kCilk {
32     typedef CilkVector<int> int_v;
33     typedef CilkVector<Precision> precision_v;
34     typedef CilkVector<bool> bool_v;
35     const static bool early_returns = false;
36     const static precision_v kOne;
37     const static precision_v kZero;
38     const static bool_v kTrue;
39     const static bool_v kFalse;
40 };
    
```

CUDA Backend

```

17 struct kCuda {
18     typedef int int_v;
19     typedef Precision precision_v;
20     typedef bool bool_v;
21     const static bool early_returns = false;
22     const static precision_v kOne = 1.0;
23     const static precision_v kZero = 0.0;
24     const static bool_v kTrue = true;
25     const static bool_v kFalse = false;
26 };
    
```

External CUDA kernels

```

227 template <TranslationCode trans_code, RotationCode rot_code, typename Backend>
228 VECGEOM_CUDA_HEADER_BOTH
229 VECGEOM_INLINE
230 typename Backend::precision_v PlacedBox::DistanceToInDispatch(
231     Vector3D<typename Backend::precision_v> const &position,
232     Vector3D<typename Backend::precision_v> const &direction,
233     const typename Backend::precision_v step_max) const {
234
235     typename Backend::precision_v output;
236     BoxDistanceToIn<trans_code, rot_code, Backend>(
237         unplaced_box()->dimensions(),
238         *this->matrix(),
239         position,
240         direction,
241         step_max,
242         &output
243     );
244     return output;
245 }
    
```

C-like abstract kernels

```

77 template <TranslationCode trans_code, RotationCode rot_code, typename Backend>
78 void BoxDistanceToIn(
79     Vector3D<Precision> const &dimensions,
80     TransformationMatrix const &matrix,
81     Vector3D<typename Backend::precision_v> const &pos,
82     Vector3D<typename Backend::precision_v> const &dir,
83     typename Backend::precision_v const &step_max,
84     typename Backend::precision_v *const distance) {
85
86     typedef typename Backend::precision_v Float;
87     typedef typename Backend::bool_v Bool;
88
89     Vector3D<Float> safety;
90     Vector3D<Float> pos_local;
91     Vector3D<Float> dir_local;
92     Bool hit(false);
93     Bool done(false);
94     *distance = kInfinity;
95
96     matrix.Transform<trans_code, rot_code>(pos, pos_local);
97     matrix.TransformRotation<rot_code>(dir, dir_local);
    
```