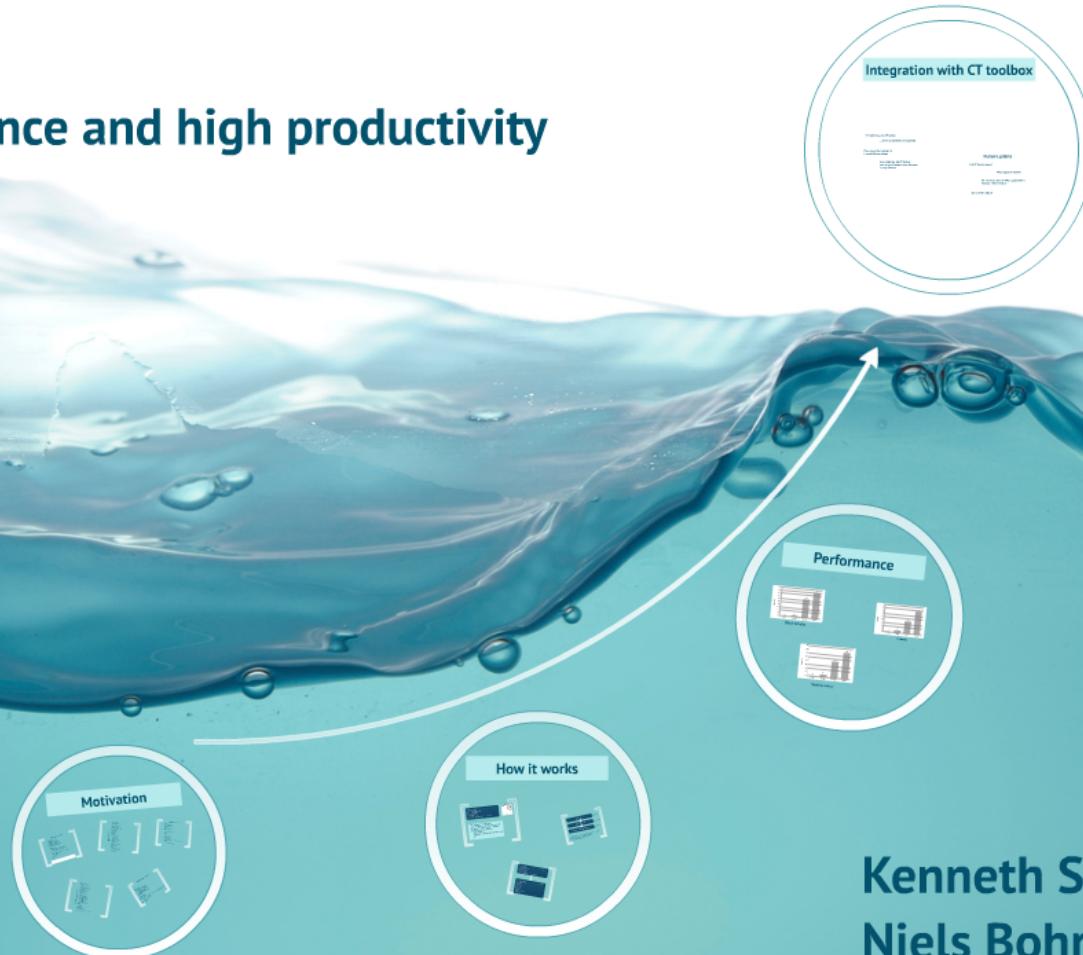


# Bohrium

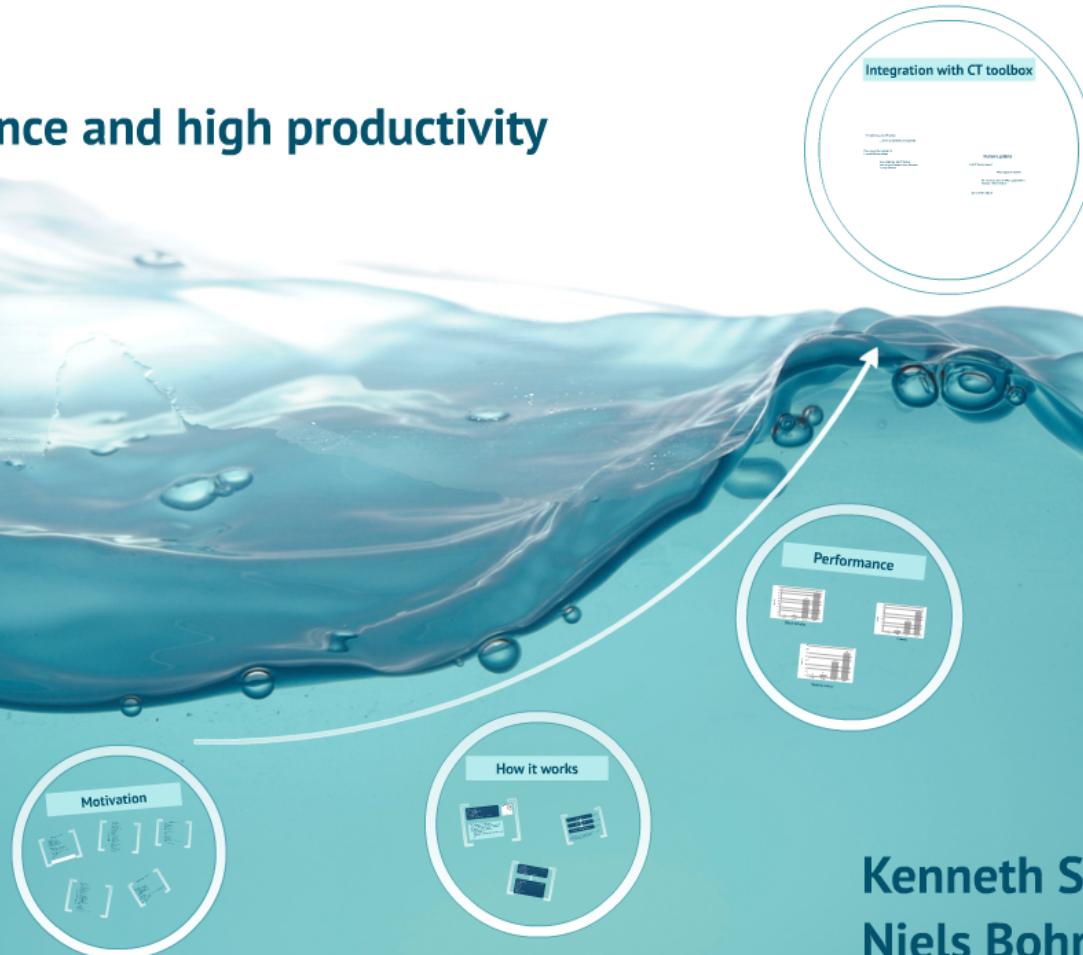
Bridging high performance and high productivity



Kenneth Skovhede  
Niels Bohr Institute  
Workshop 2014-02-13

# Bohrium

Bridging high performance and high productivity



Kenneth Skovhede  
Niels Bohr Institute  
Workshop 2014-02-13

# Motivation

Stencil example in Matlab

```
#Parameters
1. #Number of iterations
2. Input & Output Matrix
3. Temporary array
4.25. Temporary Matrix Size
```

#Computation

```
A = 1.25*ones(3,3); %Input slice vertical
I = 1.25*ones(1,3); %Output slice horizontal
for n=1:N
    for i=2:3
        for j=2:3
            T(i,j) = A(i-1,j)+A(i+1,j)+A(i,j-1)-
                10*I(i,j);
            A(i,j) = T(i,j);
        end
    end
end
```



Matlab Operator

```
% Parameters
1. #Number of iterations
2. Input & Output Matrix
3. Temporary array
4.25. Temporary Matrix Size
```

#Computation

```
A = 1.25*ones(3,3); %Input slice vertical
I = 1.25*ones(1,3); %Output slice horizontal
for n=1:N
    for i=2:3
        for j=2:3
            T(i,j) = A(i-1,j)+A(i+1,j)+A(i,j-1)-
                10*I(i,j);
            A(i,j) = T(i,j);
        end
    end
end
```

Stencil example in C

```
#include <math.h>
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <sys/malloc.h>
#include <sys/time.h>
#include <sys/conf.h>
#include <sys/sysroutines.h>
#include <sys/syscall.h>
#include <sys/types.h>
#include <sys/conf.h>
#include <sys/conf.h>
#include <sys/conf.h>
#include <sys/conf.h>
#include <sys/conf.h>
#include <sys/conf.h>
```

#Computation

```
A = 1.25*ones(3,3); %Input slice vertical
I = 1.25*ones(1,3); %Output slice horizontal
for n=1:N
    for i=2:3
        for j=2:3
            T(i,j) = A(i-1,j)+A(i+1,j)+A(i,j-1)-
                10*I(i,j);
            A(i,j) = T(i,j);
        end
    end
end
```

Stencil example with OpenMP

```
#include <omp.h>
#include <math.h>
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <sys/malloc.h>
#include <sys/time.h>
#include <sys/conf.h>
#include <sys/sysroutines.h>
#include <sys/syscall.h>
#include <sys/types.h>
#include <sys/conf.h>
#include <sys/conf.h>
#include <sys/conf.h>
#include <sys/conf.h>
#include <sys/conf.h>
```

#Computation

```
A = 1.25*ones(3,3); %Input slice vertical
I = 1.25*ones(1,3); %Output slice horizontal
for n=1:N
    for i=2:3
        for j=2:3
            T(i,j) = A(i-1,j)+A(i+1,j)+A(i,j-1)-
                10*I(i,j);
            A(i,j) = T(i,j);
        end
    end
end
```

Stencil example in NumPy

```
#Parameters
1. #Number of iterations
2. Input & Output Matrix
3. Temporary array
4.25. Temporary Matrix Size
```

#Computation

```
A = 1.25*ones(3,3); %Input slice vertical
I = 1.25*ones(1,3); %Output slice horizontal
for n=1:N
    for i=2:3
        for j=2:3
            T(i,j) = A(i-1,j)+A(i+1,j)+A(i,j-1)-
                10*I(i,j);
            A(i,j) = T(i,j);
        end
    end
end
```

# Stencil example in Matlab

#Parameters

I %Number of iterations

A %Input & Output Matrix

T %Temporary array

SIZE %Symmetric Matrix Size

#Computation

i = 2:SIZE+1;%Center slice vertical

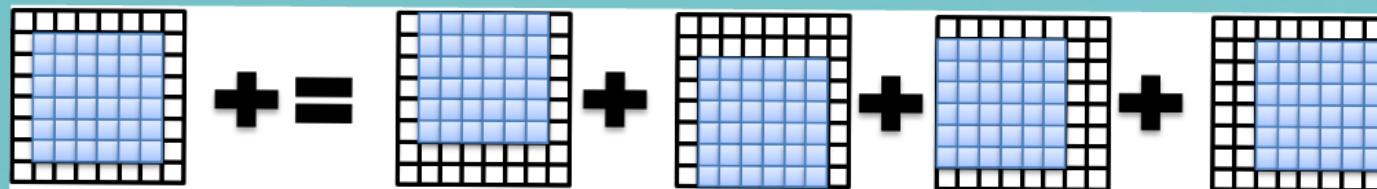
j = 2:SIZE+1;%Center slice horizontal

for n=1:I,

T(:) = (A(i,j) + A(i+1,j) + A(i-1,j) + A(i,j+1) ...  
+ A(i,j-1)) / 5.0;

A(i,j) = T;

end



# Stencil example in C

```
//Parameters
int l; //Number of iterations
double *A; //Input & Output Matrix
double *T; //Temporary array
int SIZE; //Symmetric Matrix Size

//Computation
int gsize = SIZE+2; //Size + borders.
for(n=0; n<l; n++)
{
    memcpy(T, A, gsize*gsize*sizeof(double));
    double *a = A;
    double *t = T;
    for(i=0; i<SIZE; ++i)
    {
        double *up    = a+1;
        double *left   = a+gsize;
        double *right  = a+gsize+2;
        double *down   = a+1+gsize*2;
        double *center = t+gsize+1;
        for(j=0; j<SIZE; ++j)
            *center++ = (*center + *up++ + *left++ + *right++ + *down++) / 5.0;
        a += gsize;
        t += gsize;
    }
    memcpy(A, T, gsize*gsize*sizeof(double));
}
```

## Stencil example with MPI

```
//Parameters
int l; //Number of iterations
double *A; //Input & Output Matrix (local)
double *T; //Temporary array (local)
int SIZE; //Symmetric Matrix Size (local)

//Computation
int gsize = SIZE+2; //Size + borders.
MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
MPI_Comm_size(MPI_COMM_WORLD, &worldsize);
MPI_Comm comm;
int periods[] = {0};
MPI_Cart_create(MPI_COMM_WORLD, 1, &worldsize,
    periods, 1, &comm);
int L_size = SIZE / worldsize;
if(myrank == worldsize-1)
    L_size += SIZE % worldsize;
int L_gsize = L_size + 2;//Size + borders.
for(n=0; n<l; n++)
{
    int p_src, p_dest;
    //Send/receive - neighbor above
    MPI_Cart_shift(comm,0,1,&p_src,&p_dest);
    MPI_Sendrecv(A+gsize,gsize,MPI_DOUBLE,
        p_dest,1,A,gsize,MPI_DOUBLE,
        p_src,1,comm,MPI_STATUS_IGNORE);
    //Send/receive - neighbor below
    MPI_Cart_shift(comm,0,-1,&p_src,&p_dest);
    MPI_Sendrecv(A+(L_gsize-2)*gsize,
        gsize,MPI_DOUBLE,
        p_dest,1,A+(L_gsize-1)*gsize,
        gsize,MPI_DOUBLE,
        p_src,1,comm,MPI_STATUS_IGNORE);
    memcpy(T, A, L_gsize*gsize*sizeof(double));
    double *a = A;
    double *t = T;
    for(i=0; i<SIZE; ++i)
    {
        int a = i * gsize;
        double *up   = &A[a+1];
        double *left  = &A[a+gsize];
        double *right = &A[a+gsize+2];
        double *down  = &A[a+1+gsize*2];
        double *center = &T[a+gsize+1];
        for(j=0; j<SIZE; ++j)
            *center++ = (*center + *up++ + *left++ + *right++ + *down++) / 5.0;
    }
    MPI_Barrier(MPI_COMM_WORLD);
```

# MPI with OpenMP

```
//Parameters
int l; //Number of iterations
double *A; //Input & Output Matrix (local)
double *T; //Temporary array (local)
int SIZE; //Symmetric Matrix Size (local)

//Computation
int gsize = SIZE+2; //Size + borders.
MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
MPI_Comm_size(MPI_COMM_WORLD, &worldsize);
MPI_Comm comm;
int periods[] = {0};
MPI_Cart_create(MPI_COMM_WORLD, 1, &worldsize,
    periods, 1, &comm);
int l_size = SIZE / worldsize;
if(myrank == worldsize-1)
    l_size += SIZE % worldsize;
int l_gsize = l_size + 2;//Size + borders.
for(n=0; n<l; n++)
{
    int p_src, p_dest;
    MPI_Request reqs[4];

    //Initiate send/receive - neighbor above
    MPI_Cart_shift(comm, 0, 1, &p_src, &p_dest);
    MPI_Isend(A+gsize, gsize, MPI_DOUBLE, p_dest,
        1, comm, &reqs[0]);
    MPI_Irecv(A, gsize, MPI_DOUBLE, p_src,
        1, comm, &reqs[1]);

    //Initiate send/receive - neighbor below
    MPI_Cart_shift(comm, 0, -1, &p_src, &p_dest);
    MPI_Isend(A+(l_gsize-2)*gsize, gsize,
        MPI_DOUBLE,
        p_dest, 1, comm, &reqs[2]);
    MPI_Irecv(A+(l_gsize-1)*gsize, gsize,
        MPI_DOUBLE,
        p_src, 1, comm, &reqs[3]);

    //Handle the non-border elements.
    memcpy(T+gsize, A+gsize, l_size*gsize*sizeof(double));
    #pragma omp parallel for shared(A,T)
    for(i=1; i<l_size-1; ++i)
        compute_row(i,A,T,SIZE,gsize);

    //Handle the upper and lower ghost line
    MPI_Waitall(4, reqs, MPI_STATUSES_IGNORE);
    compute_row(0,A,T,SIZE,gsize);
    compute_row(l_size-1,A,T,SIZE,gsize);

    memcpy(A+gsize, T+gsize, l_size*gsize*sizeof(double));
}
MPI_Barrier(MPI_COMM_WORLD);
```

# Stencil example in NumPy

#Parameters

I #Number of iterations  
A #Input & Output Matrix  
T #Temporary array  
SIZE #Symmetric Matrix Size

#Computation

```
for i in xrange(I):  
    T[:] = (A[1:-1,1:-1] + A[1:-1,:-2] + A[1:-1,2:] + A[:-2,1:-1] \  
            + A[2:,1:-1]) / 5.0  
    A[1:-1, 1:-1] = T
```

# How it works

```
#Monte Carlo Pi  
size = [100, 2000, 10]  
x = random(size)  
y = random(size)  
sum = add.aggregate((x*x + y*y) <= 1)*4
```



- Can be written as sequential code
- Can be implemented as library in any language
- Can use special source language constructs
- Bohrium currently has support for:
  - Python, NumPy
  - CIL languages (.NET): C#, F#, VB, IronPython, etc.
  - C++
  - C

An interesting thing we can immediately see from Bohrium's execution of the above code is that it gets converted into vectorized code. This is because Bohrium uses a JIT compiler to convert user code into efficient native code.

Source code  
↓  
Vector byte code  
↓  
Hardware optimized execution

Any language  
Abstraction glue  
Any hardware

```
#Monte Carlo Pi  
size = [100, 2000, 10]  
x = random(size)  
y = random(size)  
sum = add.aggregate((x*x + y*y) <= 1)*4
```

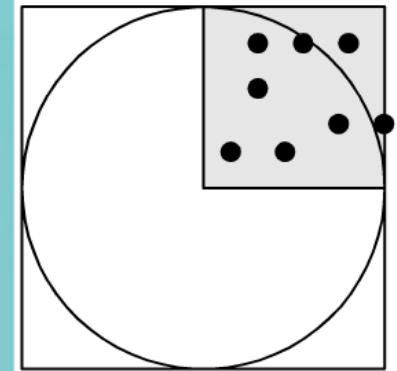
Native code

Metered code

Native code

Metered code

```
#Monte Carlo PI  
size = [100, 2000, 20]  
x = random(size)  
y = random(size)  
sum = add.aggregate((x^2 + y^2) <= 1)*4
```



- Can be written as sequential code
- Can be implemented as library in any language
- Can use special source language constructs
- Bohrium currently has support for:
  - Python/NumPy
  - CIL languages (.Net): C#, F#, VB, IronPython, etc.
  - C++
  - C

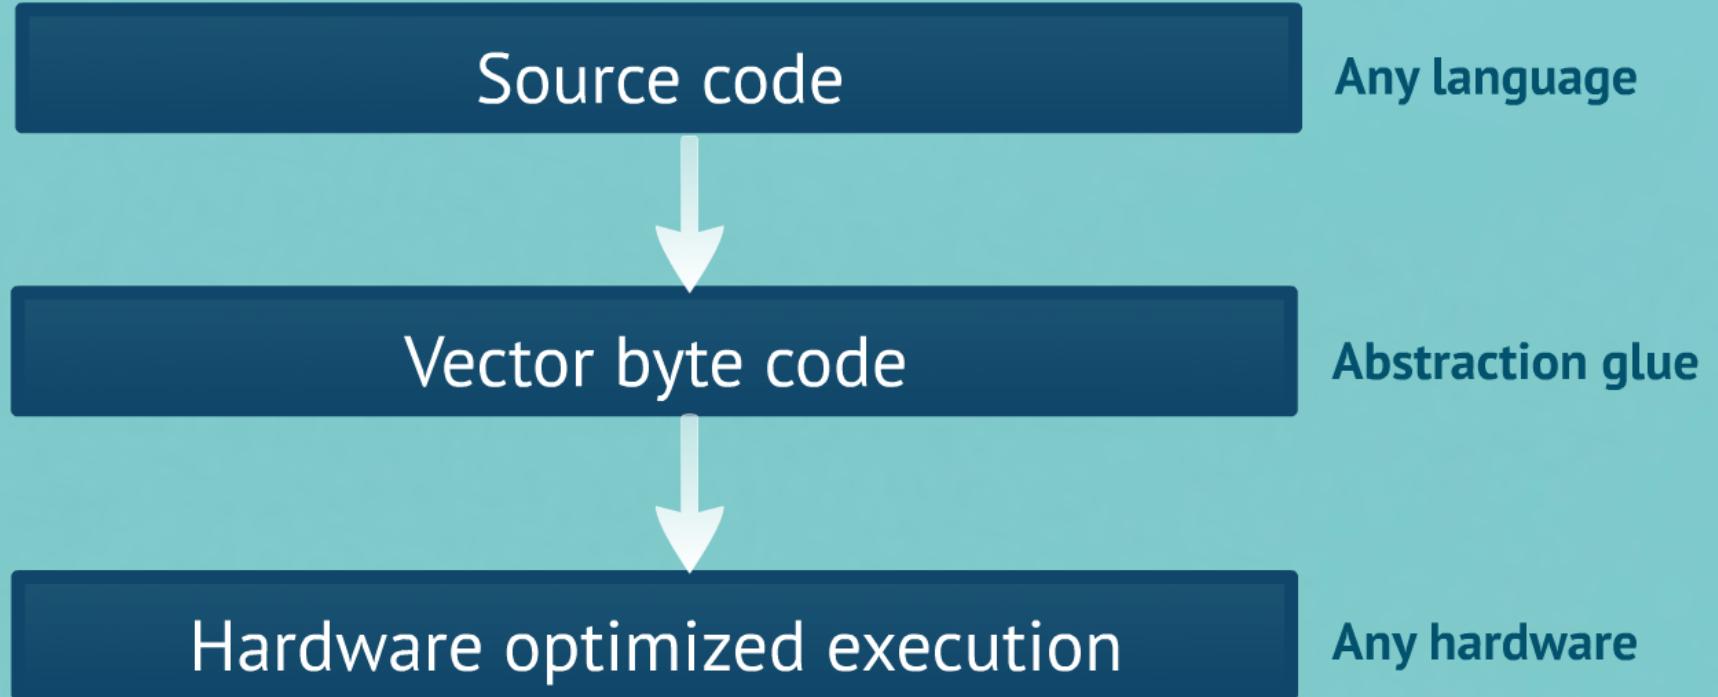
```
#Monte Carlo PI  
size = [100, 2000, 20]  
x = random(size)  
y = random(size)  
sum = add.aggregate((x^2 + y^2) <= 1)*4
```

Sequential code



```
X = NEW(100,2000,20)  
Y = NEW(100,2000,20)  
RND(X,X)  
RND(Y,Y)  
POW(T0,X,2)  
POW(T1,Y,2)  
ADD(T2,T1,T0)  
LTE(T3,T2,1)  
AGGREGATE(ADD,T4,T3)  
MUL(T5,T4,4)
```

Vector bytecode



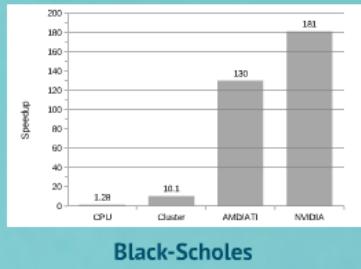
**Any existing program can immediately use new hardware because of the abstraction**

**Every language gets all hardware support without any effort**

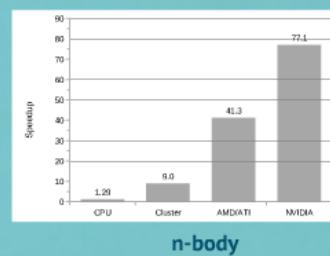
**Same basic idea as compiler intermediate files (aka object files)**

**but for vector operations and done runtime**

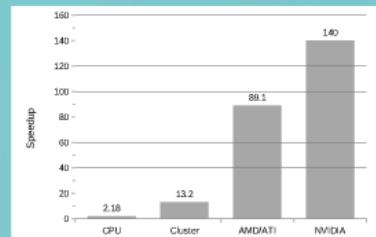
# Performance



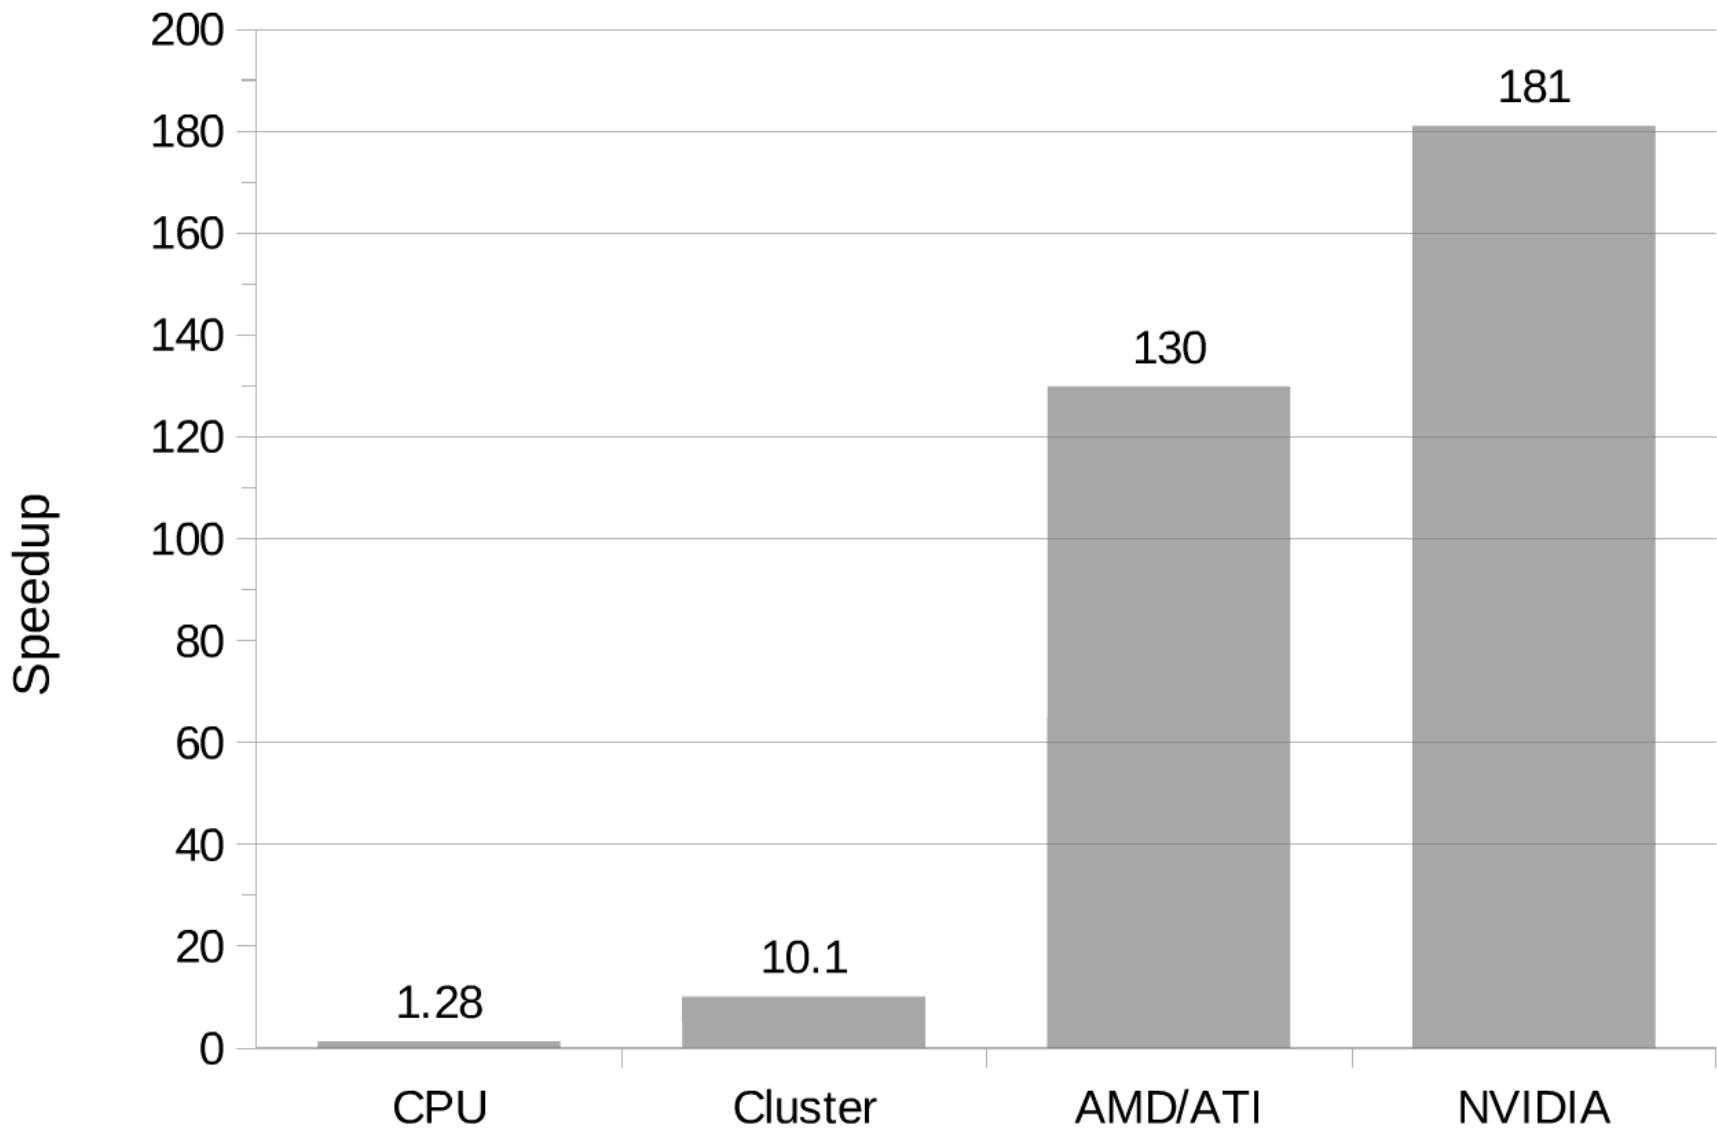
Black-Scholes



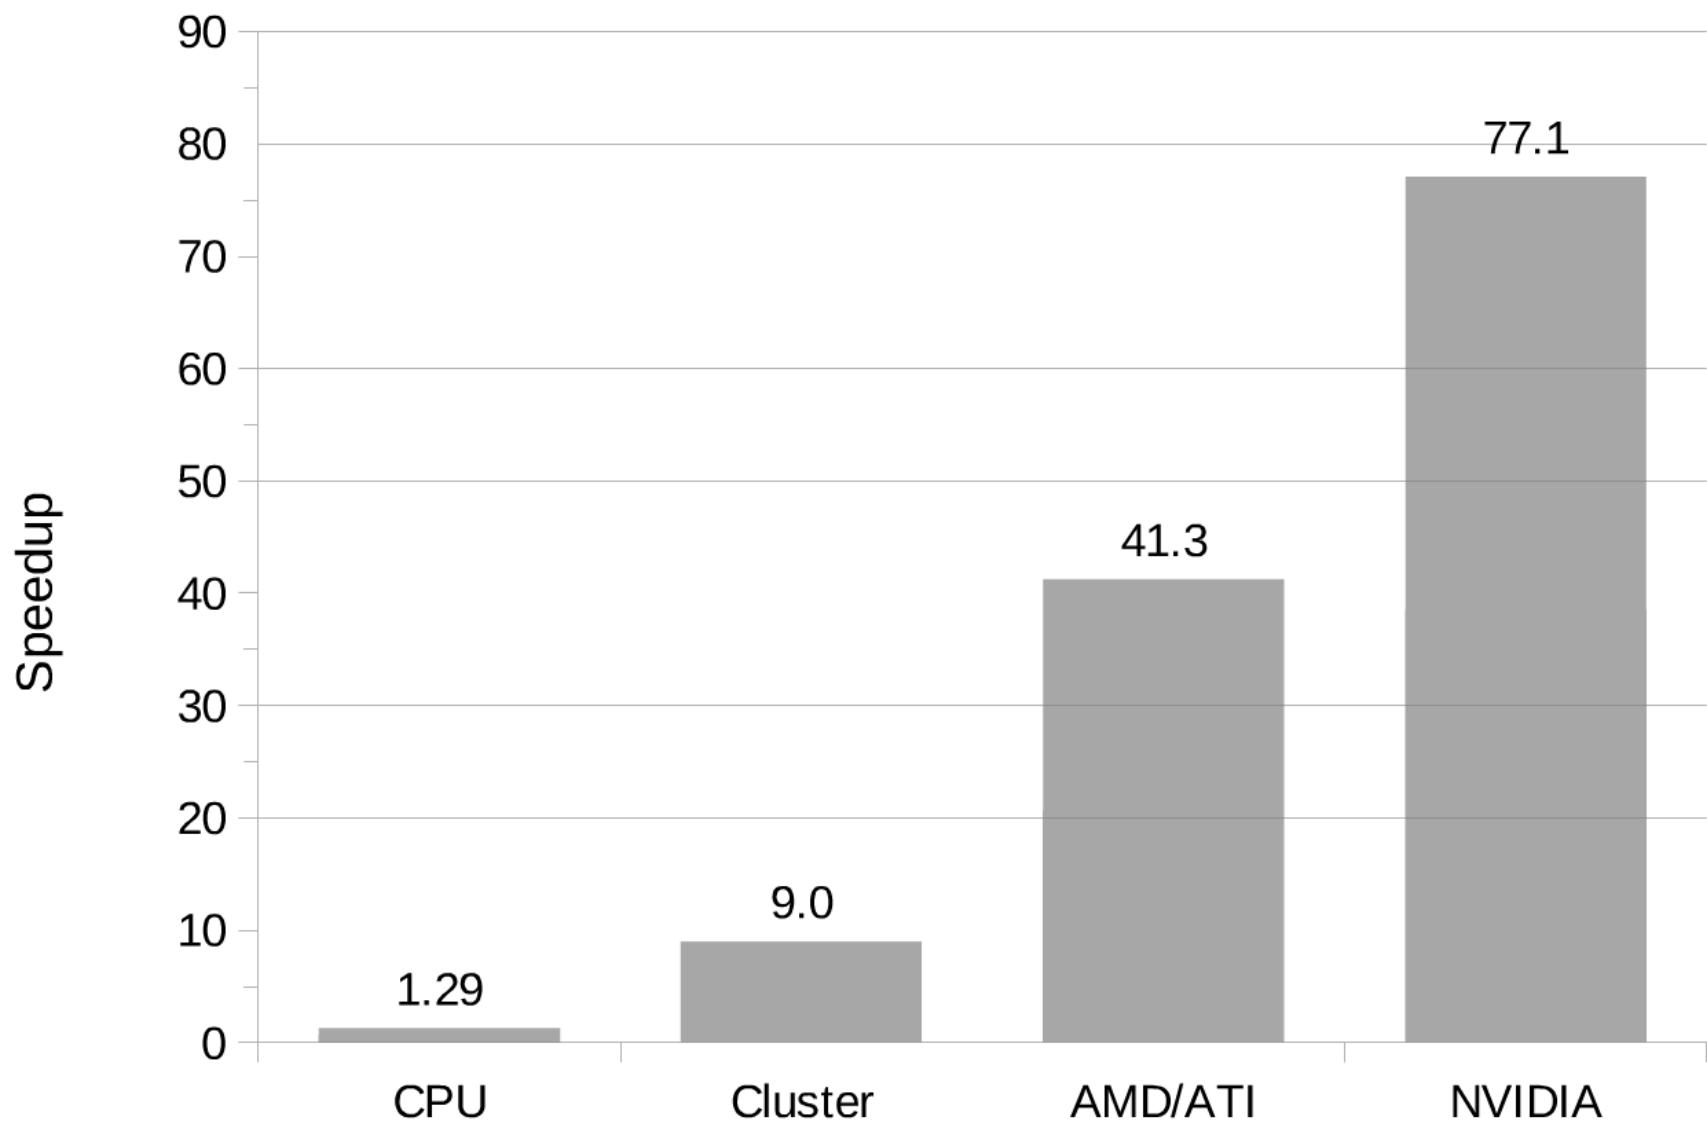
n-body



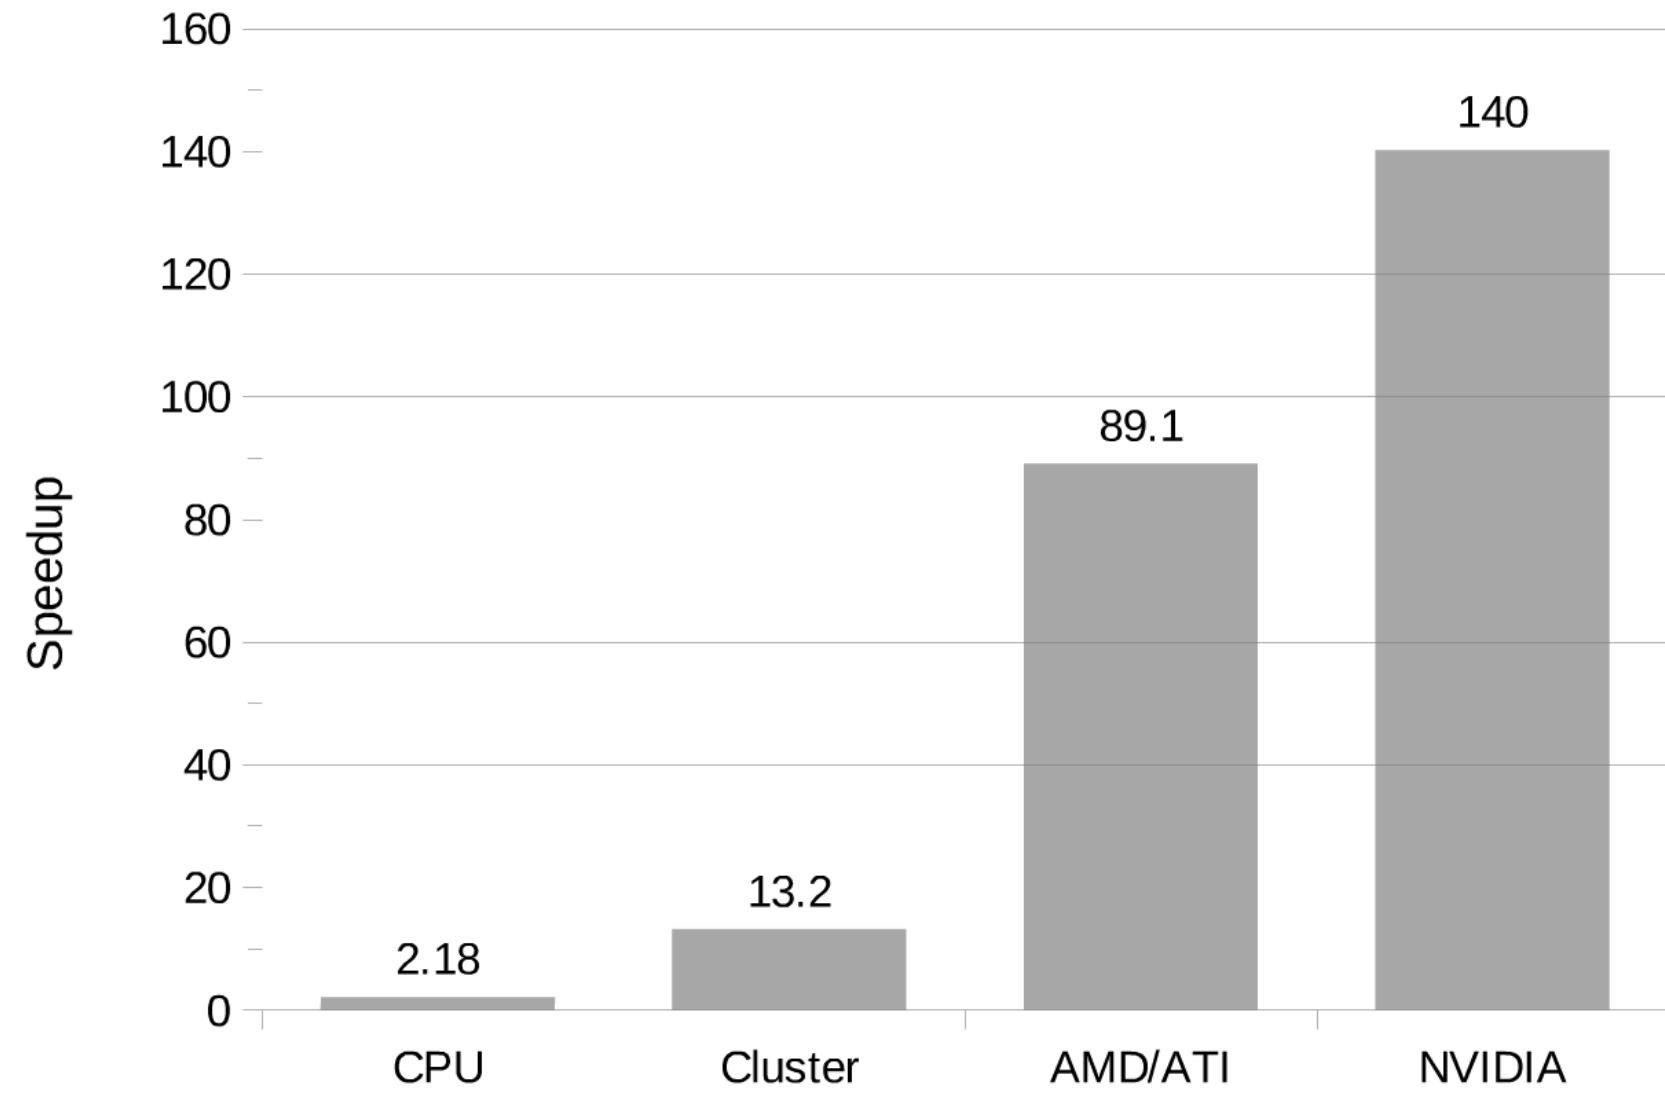
Shallow water



# Black-Scholes



# n-body



# Shallow water

# Integration with CT toolbox

CT Toolbox has a NumPy version

... but not all operations are supported

These are on the roadmap for  
the next Bohrium release

Once completed, any CT Toolbox  
application will support more hardware  
through Bohrium

## Future plans

Full CT Toolbox support

FPGA support in Bohrium

We have most parts of a FPGA based Bohrium  
Processor with a simulator

Sparse matrix support

**CT Toolbox has a NumPy version**

**... but not all operations are supported**

**These are on the roadmap for  
the next Bohrium release**

**Once completed, any CT Toolbox  
application will support more hardware  
through Bohrium**

# Future plans

**Full CT Toolbox support**

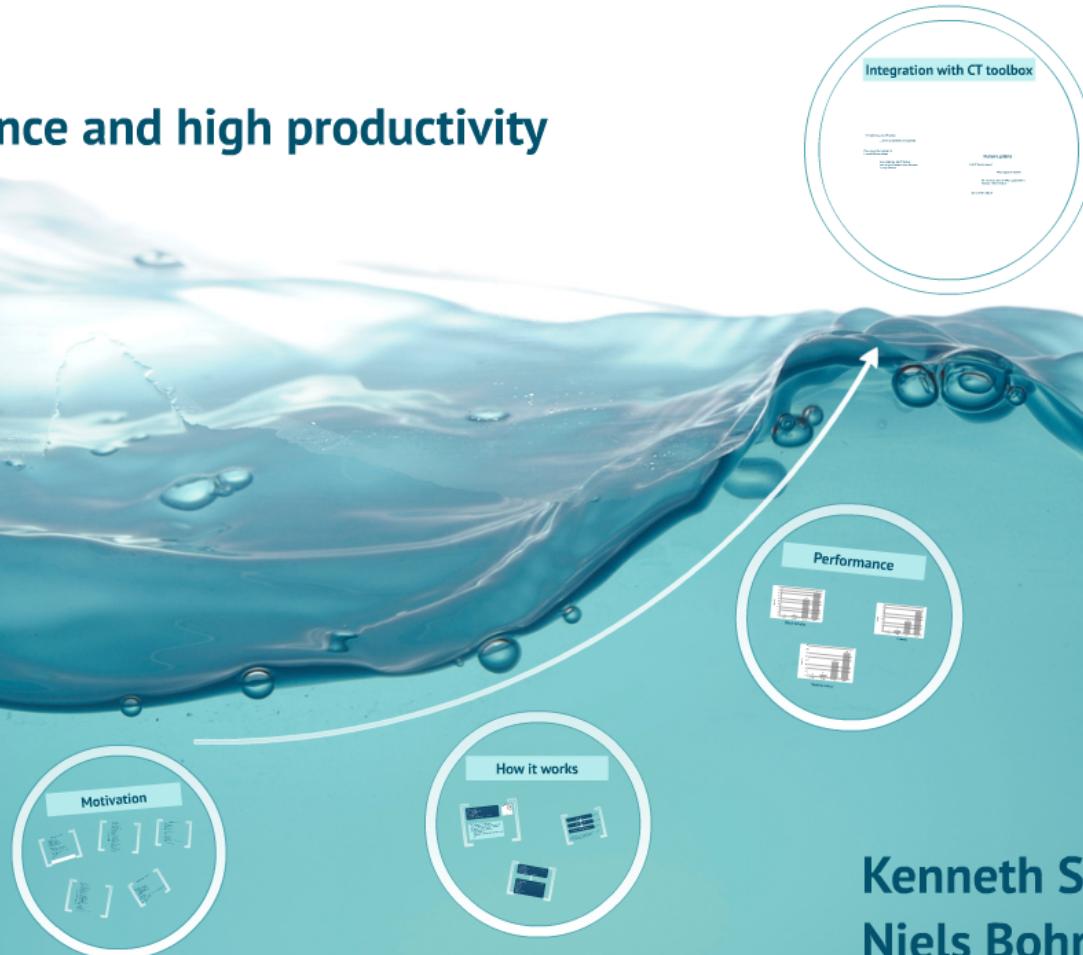
**FPGA support in Bohrium**

**We have most parts of a FPGA based Bohrium  
Processor with a simulator**

**Sparse matrix support**

# Bohrium

Bridging high performance and high productivity



Kenneth Skovhede  
Niels Bohr Institute  
Workshop 2014-02-13