

Implementing Scalable Matrix-Vector Products for the Exact Diagonalization Methods in Quantum Many-Body Physics

Tom Westerhout and Bradford L. Chamberlain

Radboud Universiteit



Matrix-vector products



Matrix-vector products

$$y = A \cdot x$$

Matrix-vector products

$$y = A \cdot x$$


$$y_i = \sum_{j=0}^{N-1} A_{i,j} \cdot x_j$$


Matrix-vector products

$$y = A \cdot x$$

$$y_i = \sum_{j=0}^{N-1} A_{i,j} \cdot x_j$$

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} e \\ f \end{pmatrix} = \begin{pmatrix} ae + bf \\ ce + df \end{pmatrix}$$



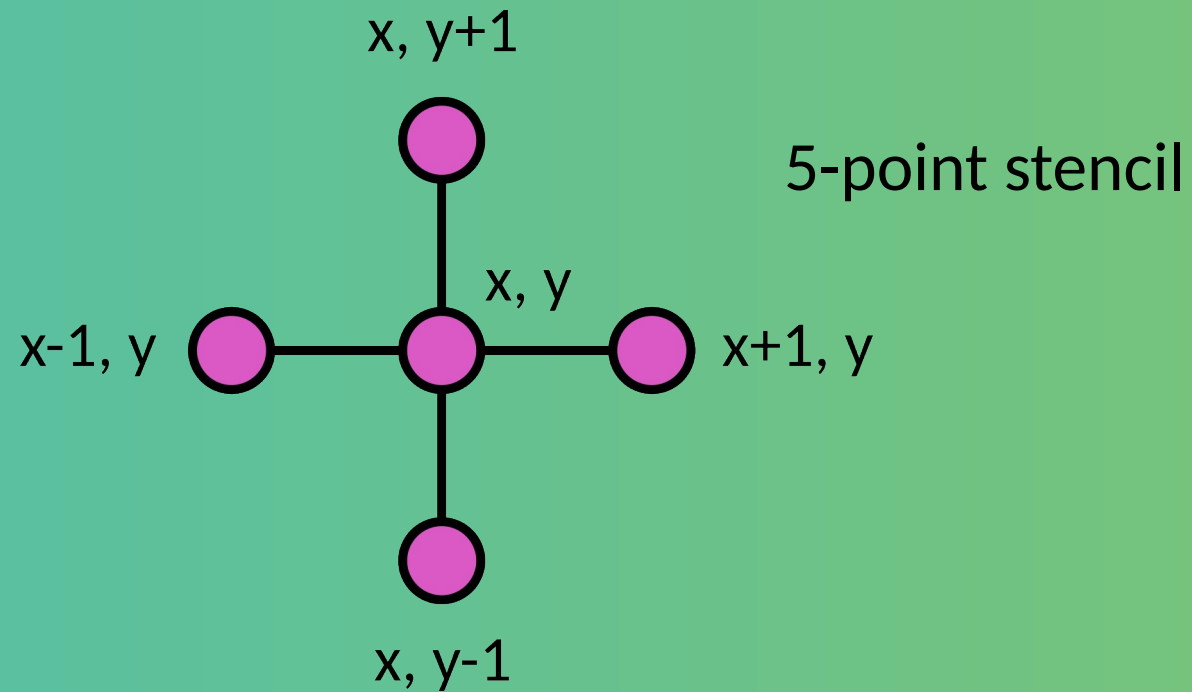


Lucky 

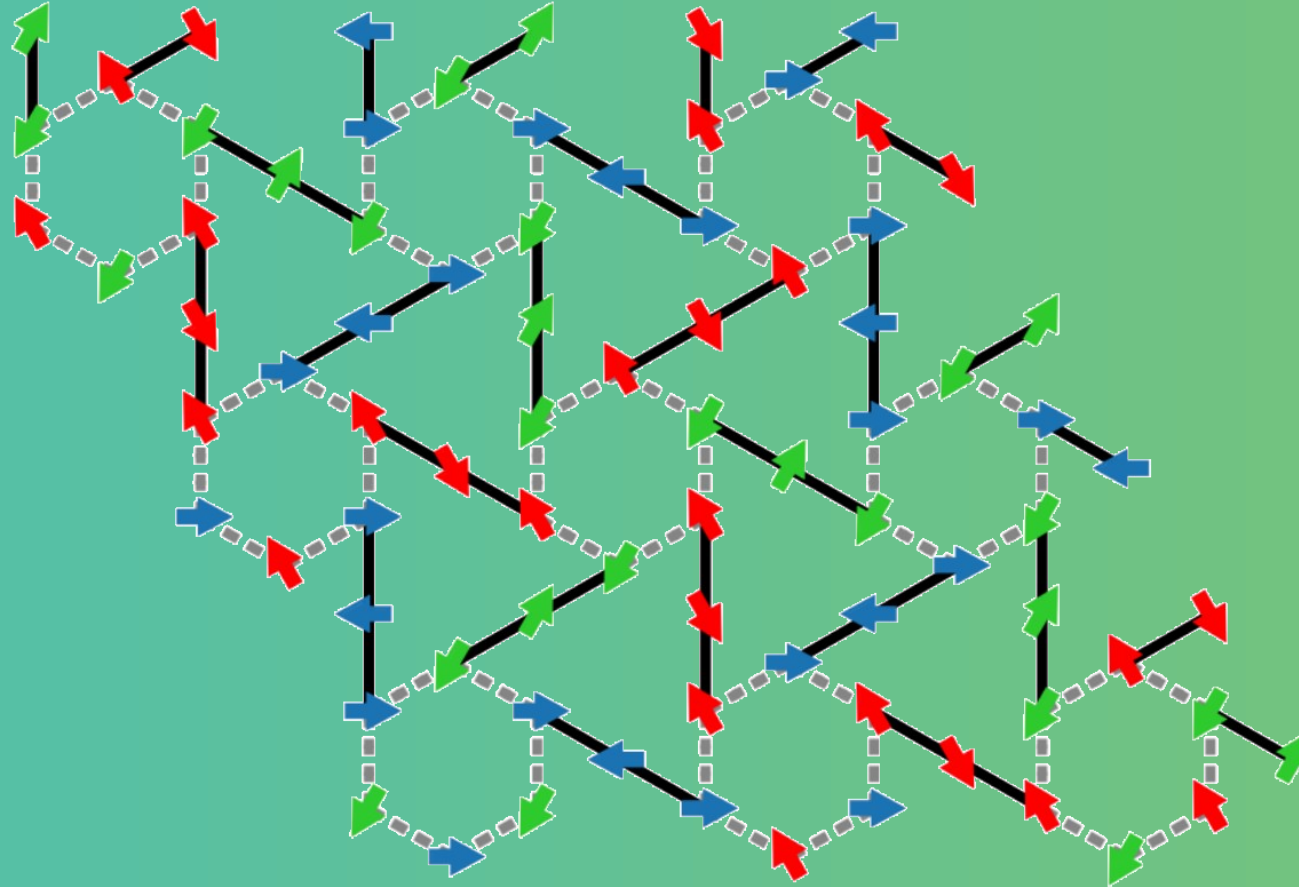


Lucky 


Finite differences methods



Quantum mechanics



Exact diagonalization




Physical system



$$\begin{bmatrix} 3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 2 & 0 & 2 & 0 & 0 & 0 \\ 0 & 2 & -1 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 2 & 2 & 0 \\ 0 & 2 & 2 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & -1 & 2 & 0 \\ 0 & 0 & 0 & 2 & 0 & 2 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3 \end{bmatrix}$$

Hamiltonian matrix

Exact diagonalization



Physical system




$$\begin{bmatrix} 3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 2 & 0 & 2 & 0 & 0 & 0 \\ 0 & 2 & -1 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 2 & 2 & 0 \\ 0 & 2 & 2 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & -1 & 2 & 0 \\ 0 & 0 & 0 & 2 & 0 & 2 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3 \end{bmatrix}$$

Hamiltonian matrix



Dimension: $\mathcal{O}(2^N)$

Exact diagonalization




Physical system

$$\begin{bmatrix} 3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 2 & 0 & 2 & 0 & 0 & 0 \\ 0 & 2 & -1 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 2 & 2 & 0 \\ 0 & 2 & 2 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & -1 & 2 & 0 \\ 0 & 0 & 0 & 2 & 0 & 2 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3 \end{bmatrix}$$

Hamiltonian matrix

$$\hat{H} = \sum_{i=0}^2 \sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z$$

Exact diagonalization



$$\begin{aligned} & \left(\sigma^x_0 \sigma^x_1 + \sigma^y_0 \sigma^y_1 + \sigma^z_0 \sigma^z_1 + \right. \\ & \left. \sigma^x_1 \sigma^x_2 + \sigma^y_1 \sigma^y_2 + \sigma^z_1 \sigma^z_2 + \right. \\ & \left. \sigma^x_2 \sigma^x_0 + \sigma^y_2 \sigma^y_0 + \sigma^z_2 \sigma^z_0 \right) \end{aligned}$$

Physical system

Expression



$$\hat{H} = \sum_{i=0}^2 \sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z$$

Procedure


Store the Hamiltonian
as an equation

$$\begin{aligned} & \left(\sigma^x_0 \sigma^x_1 + \sigma^y_0 \sigma^y_1 + \sigma^z_0 \sigma^z_1 + \right. \\ & \left. \sigma^x_1 \sigma^x_2 + \sigma^y_1 \sigma^y_2 + \sigma^z_1 \sigma^z_2 + \right. \\ & \left. \sigma^x_2 \sigma^x_0 + \sigma^y_2 \sigma^y_0 + \sigma^z_2 \sigma^z_0 \right) \end{aligned}$$

Solve



Describe the algorithm
with linear algebra primitives



Receive the
Nobel Prize







Memory for one vector: $\mathcal{O}(2^N)$ \rightarrow

distributed implementation
required



Memory for one vector: $\mathcal{O}(2^N)$ \rightarrow distributed implementation required

A Ph.D. or postdoc contract is short



lattice-symmetries (LS)



lattice-symmetries (LS)



lattice-symmetries (LS)

 **Haskell**
Halide



lattice-symmetries (LS)



 **Haskell**

Halide




lattice-symmetries (LS)



Halide


lattice-symmetries (LS)





Nix


— a package manager and a programming language



Nix

— a package manager and a programming language

```
toContainer = package: singularity-tools.buildImage {  
  name = package.pname;  
  contents = [ package ];  
};
```




Nix

— a package manager and a programming language

```
toContainer = package: singularity-tools.buildImage {  
  name = package.pname;  
  contents = [ package ];  
};
```

(just 200 MB with Chapel & Haskell runtimes,
Infiniband support etc.)



Nix

— a package manager and a programming language

```
toContainer = package: singularity-tools.buildImage {  
  name = package.pname;  
  contents = [ package ];  
};
```

(just 200 MB with Chapel & Haskell runtimes,
Infiniband support etc.)

```
mkShell {  
  buildInputs = [  
    lattice-symmetries.kernels  
    lattice-symmetries.haskell  
    hdf5 hdf5.dev  
  ];  
  nativeBuildInputs = [  
    chapel chapelFixupBinary  
    gcc pkg-config  
  ];  
};
```



$$y_i = \sum_j H_{i,j} x_j$$



$$y_i = \sum_j H_{i,j} x_j$$

```
1 forall i in y.domain {
2   var acc = 0.0;
3   for (Hij, j) in getRow(H, i) {
4     acc += Hij * x[j];
5   }
6   y[i] = acc;
7 }
```



parallelism over
nodes and cores

$$y_i = \sum_j H_{i,j} x_j$$

```
1 forall i in y.domain {  
2   var acc = 0.0;  
3   for (Hij, j) in getRow(H, i) {  
4     acc += Hij * x[j];  
5   }  
6   y[i] = acc;  
7 }
```




parallelism over
nodes and cores

$$y_i = \sum_j H_{i,j} x_j$$

```
1 forall i in y.domain {  
2   var acc = 0.0;  
3   for (Hij, j) in getRow(H, i) {  
4     acc += Hij * x[j];  
5   }  
6   y[i] = acc;  
7 }
```


potentially remote
memory access

Symmetries



- 0 | $\downarrow\downarrow\downarrow\downarrow$ \rangle
- 1 | $\downarrow\downarrow\downarrow\uparrow$ \rangle
- 2 | $\downarrow\downarrow\uparrow\downarrow$ \rangle
- 3 | $\downarrow\downarrow\uparrow\uparrow$ \rangle
- 4 | $\downarrow\uparrow\downarrow\downarrow$ \rangle
- 5 | $\downarrow\uparrow\downarrow\uparrow$ \rangle
- ⋮
- 15 | $\uparrow\uparrow\uparrow\uparrow$ \rangle

Symmetries



- 0 | ↓↓↓↓⟩
- 1 | ↓↓↓↑⟩
- 2 | ↓↓↑↓⟩
- 3 | ↓↓↑↑⟩
- 4 | ↓↑↓↓⟩
- 5 | ↓↑↓↑⟩
- ⋮
- 15 | ↑↑↑↑⟩

impose translational
invariance



- 0 { | ↓↓↓↓⟩ }
- 1 { | ↓↓↓↑⟩, | ↓↓↑↓⟩, | ↓↑↓↓⟩, | ↑↓↓↓⟩ }
- 2 { | ↓↓↑↑⟩, | ↓↑↑↓⟩, | ↑↓↑↑⟩, | ↑↑↓↑⟩ }
- 3 { | ↓↑↓↑⟩, | ↑↓↑↓⟩ }
- 4 { | ↓↑↑↑⟩, | ↑↓↑↑⟩, | ↑↑↓↑⟩, | ↑↑↑↓⟩ }
- 5 { | ↑↑↑↑⟩ }

$$y_i = \sum_j H_{i,j} x_j$$

```
1 forall i in basisStates.domain {
2   const beta = basisStates[i];
3   var acc = 0.0;
4   for (Hij, alpha) in getRow(H, beta) {
5     const ref srcLocale = Locales[localeIdxOf(alpha)];
6     on srcLocale { // ← remote task spawn
7       const j = stateToIndex(basisStates, alpha);
8       acc += Hij * x[j];
9     }
10  }
11  y[i] = acc;
12 }
```

$$y_i = \sum_j H_{i,j} x_j$$

```
1 y = 0;
2 forall j in basisStates.domain {
3   const alpha = basisStates[j];
4   for (Hij, beta) in getColumn(H, alpha) {
5     const ref destLocale = Locales[localeIdxOf(beta)];
6     const coeff = Hij * x[j];
7     on destLocale {
8       const i = stateToIndex(basisStates, beta);
9       y[i] += coeff; // atomically
10    }
11  }
12 }
```

$$y_i = \sum_j H_{i,j} x_j$$

```
1 y = 0;
2 forall j in basisStates.domain {
3   const alpha = basisStates[j];
4   for (Hij, beta) in getColumn(H, alpha) {
5     const ref destLocale = Locales[localeIdxOf(beta)];
6     const coeff = Hij * x[j];
7     on destLocale {
8       const i = stateToIndex(basisStates, beta);
9       y[i] += coeff; // atomically
10    }
11  }
12 }
```

} Producer

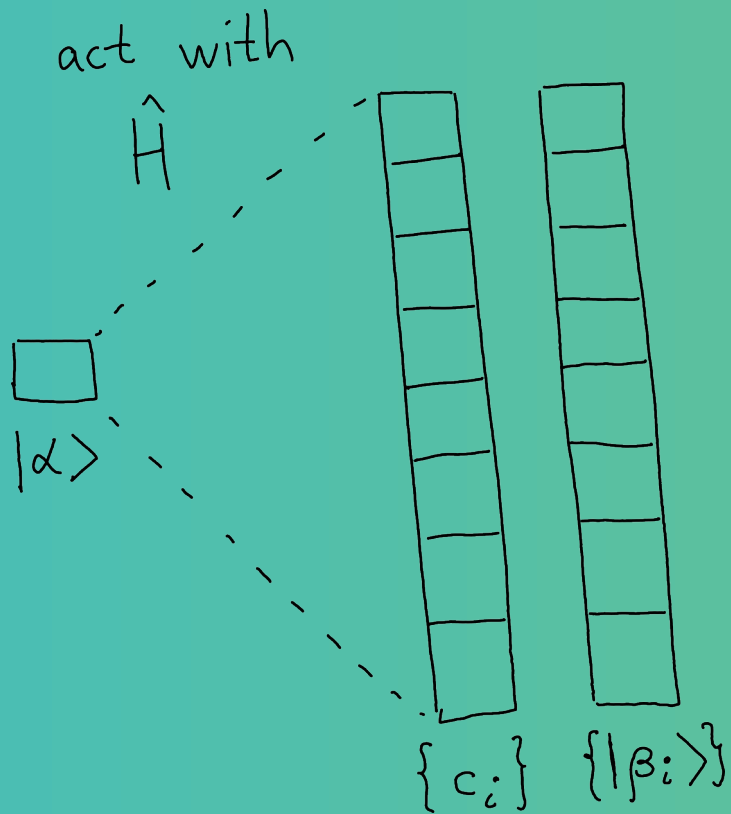
$$y_i = \sum_j H_{i,j} x_j$$

```
1 y = 0;
2 forall j in basisStates.domain {
3   const alpha = basisStates[j];
4   for (Hij, beta) in getColumn(H, alpha) {
5     const ref destLocale = Locales[localeIdxOf(beta)];
6     const coeff = Hij * x[j];
7     on destLocale {
8       const i = stateToIndex(basisStates, beta);
9       y[i] += coeff; // atomically
10    }
11  }
12 }
```

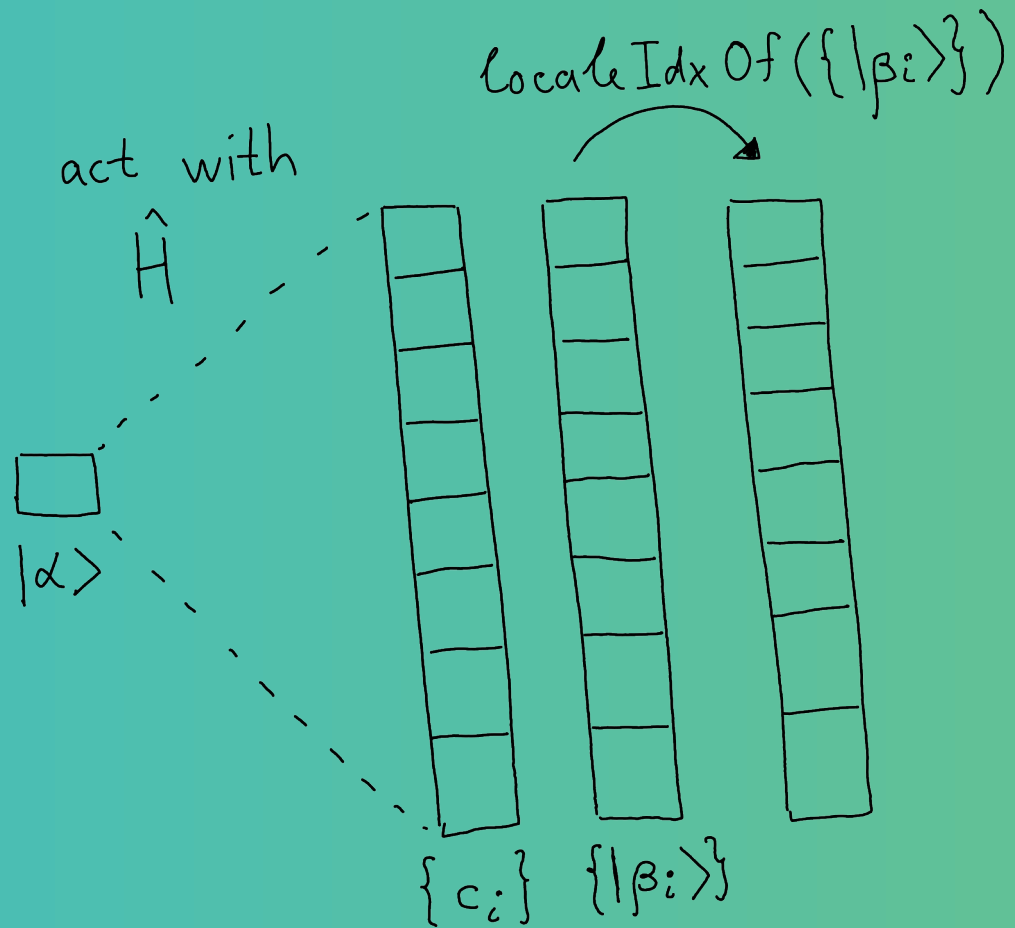
} Producer

} Consumer


Producer



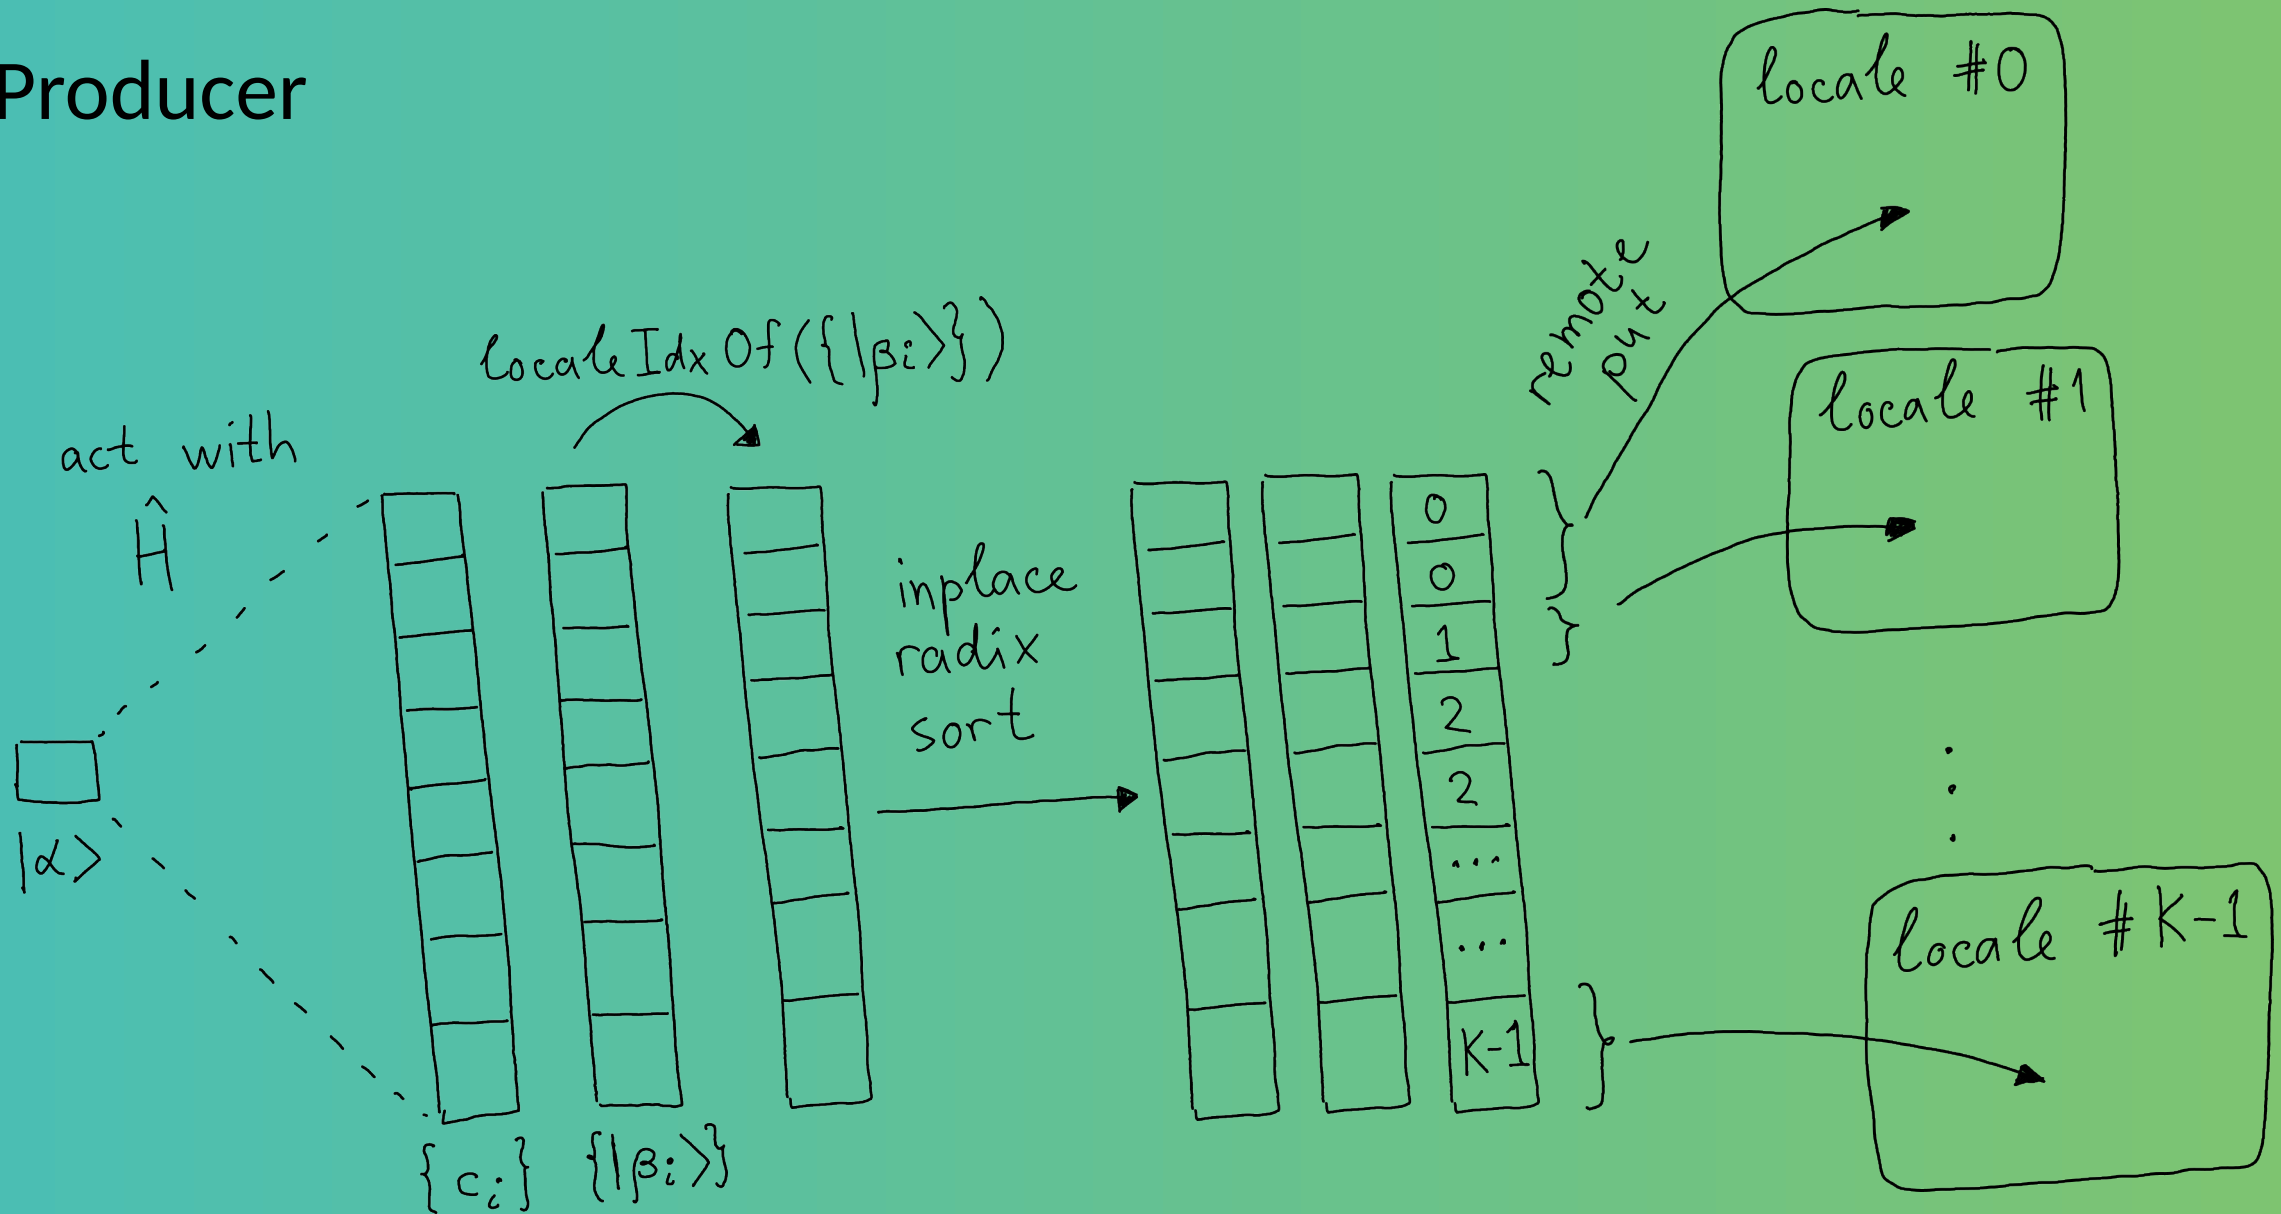
Producer



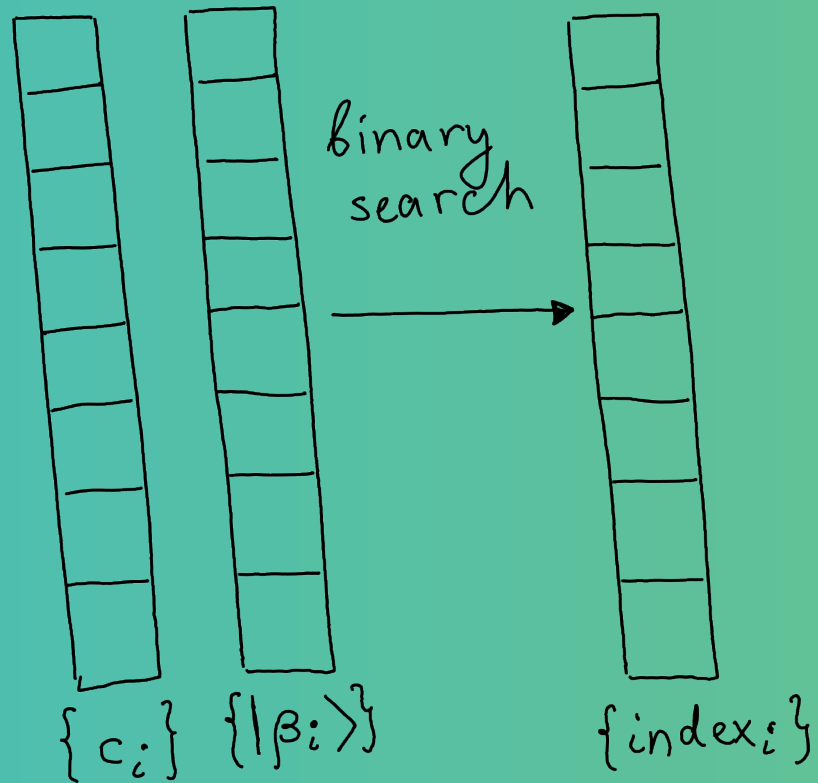
Producer



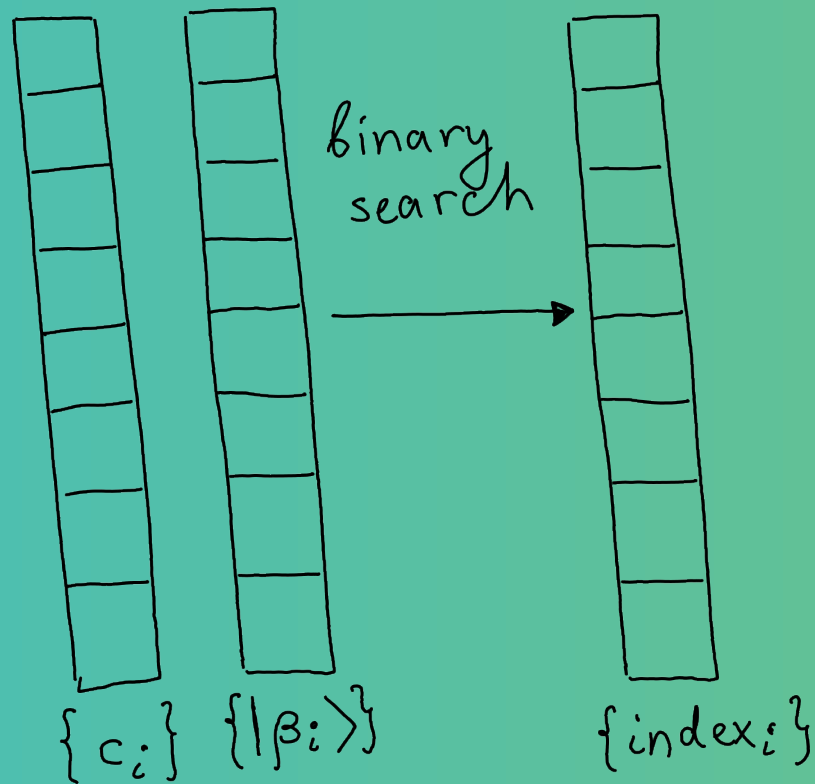
Producer



Consumer



Consumer



```
for i in 0 ..# count {  
  // atomic +=  
  y[index[i]].add(c[i], memoryOrder.relaxed);  
}
```

Scaling

Snellius

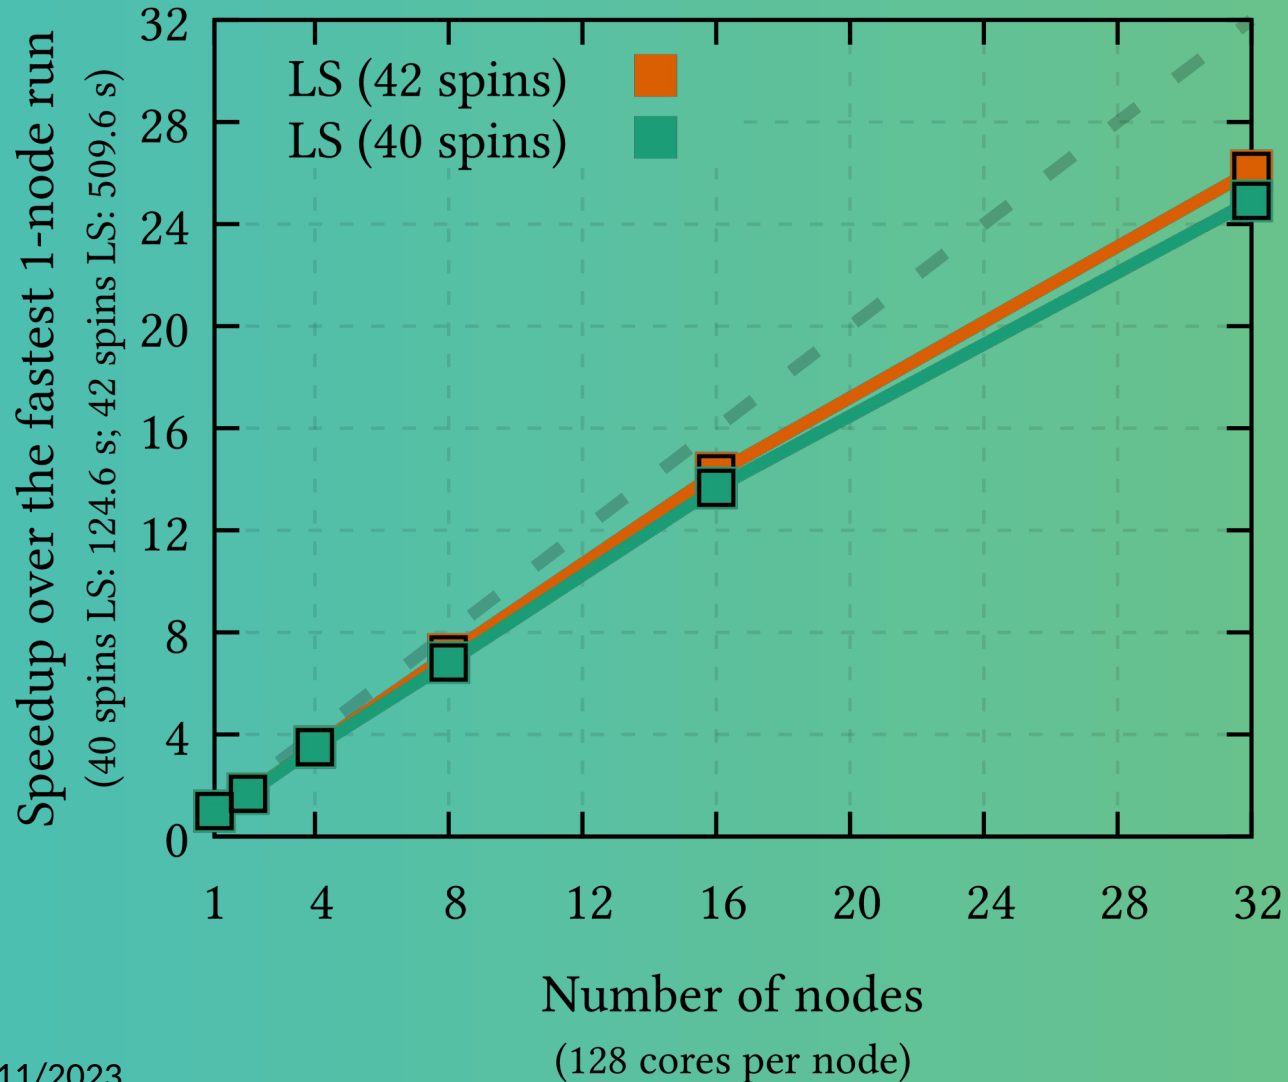
2× AMD Rome 7H12 @2.6GHz

64 cores/socket

100 Gb/s Infiniband

	No symmetries	With symmetries
40 spins	1.0×10^{12}	8.6×10^8
42 spins	4.4×10^{12}	3.2×10^9

Scaling



Snellius

2× AMD Rome 7H12 @2.6GHz
64 cores/socket
100 Gb/s Infiniband

	No symmetries	With symmetries
40 spins	1.0×10^{12}	8.6×10^8
42 spins	4.4×10^{12}	3.2×10^9

State-of-the-art MPI-based solution

SPINPACK

Written in C

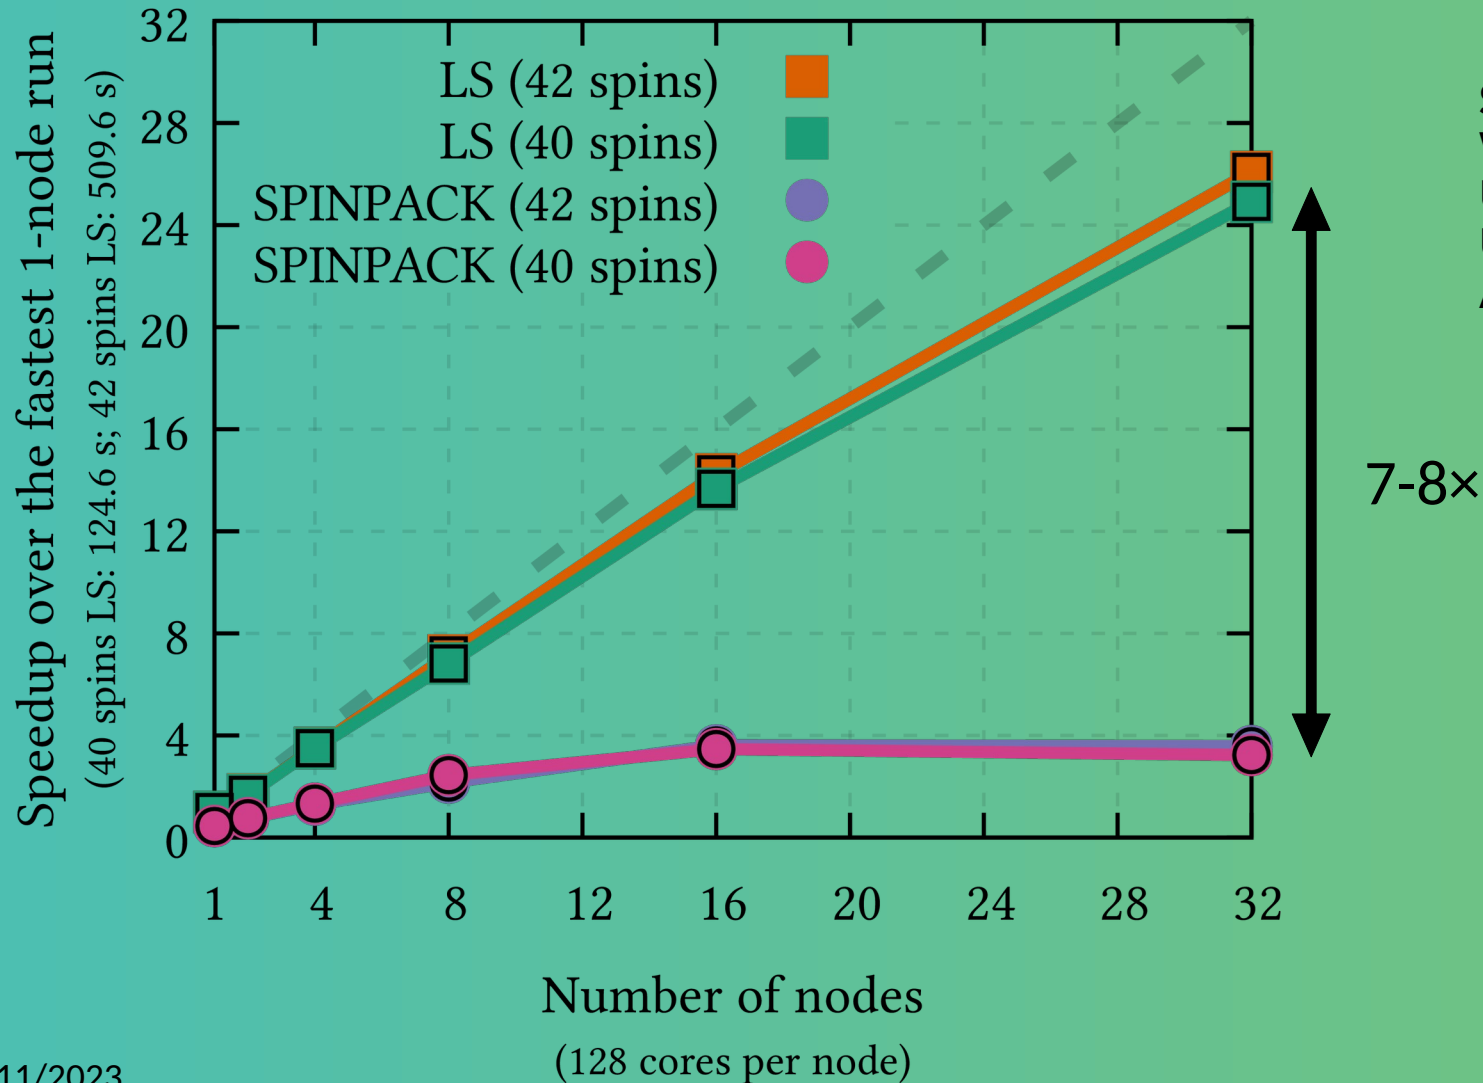
Uses MPI

Used in publications since 1995

About 26 000 software lines of code



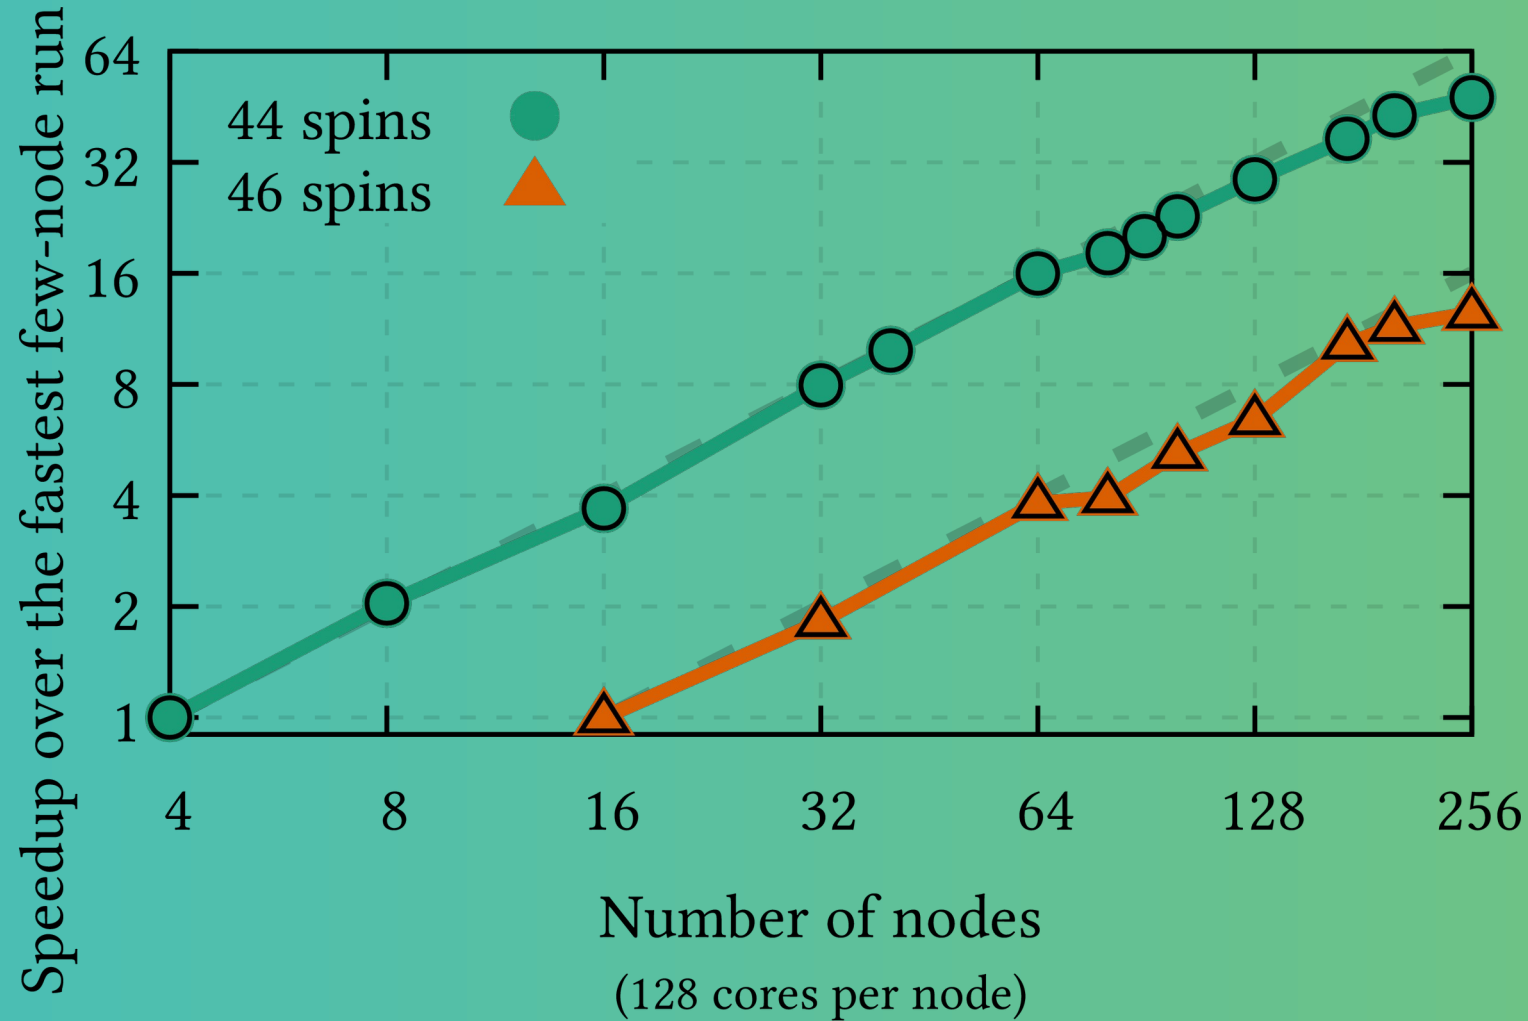
State-of-the-art MPI-based solution



SPINPACK
Written in C
Uses MPI
Used in publications since 1995
About 26 000 software lines of code

7-8x


Further scaling



	No symmetries	With symmetries
44 spins	1.8×10^{13}	1.1×10^{10}
46 spins	7.0×10^{13}	4.5×10^{10}
48 spins	2.8×10^{14}	1.7×10^{11}

32 768 cores

Further scaling



	No symmetries	With symmetries
44 spins	1.8×10^{13}	1.1×10^{10}
46 spins	7.0×10^{13}	4.5×10^{10}
48 spins	2.8×10^{14}	1.7×10^{11}

2.2 PB 🤯

32 768 cores

Conclusion

- Chapel-based implementation scales well to 256 nodes (or 32 768 cores)
- 7-8× improvement over the state of the art at 32 nodes
- 3× fewer SLOC than the state of the art
- Nix + Apptainer is an easy deployment solution
- Happy with Chapel — no desire to rewrite everything using MPI

Non-standard tools may still lead to competitive results

