

# A Data-Centric Optimization Workflow for the Python Language

## Alexandros Nikolaos Ziogas, Torsten Hoefer

**Motivation** Python is the language of choice for scientific computing due to its high productivity. However, its execution is often slow. How do we make Python the language of choice for writing highly performant code too?

### From Python ...

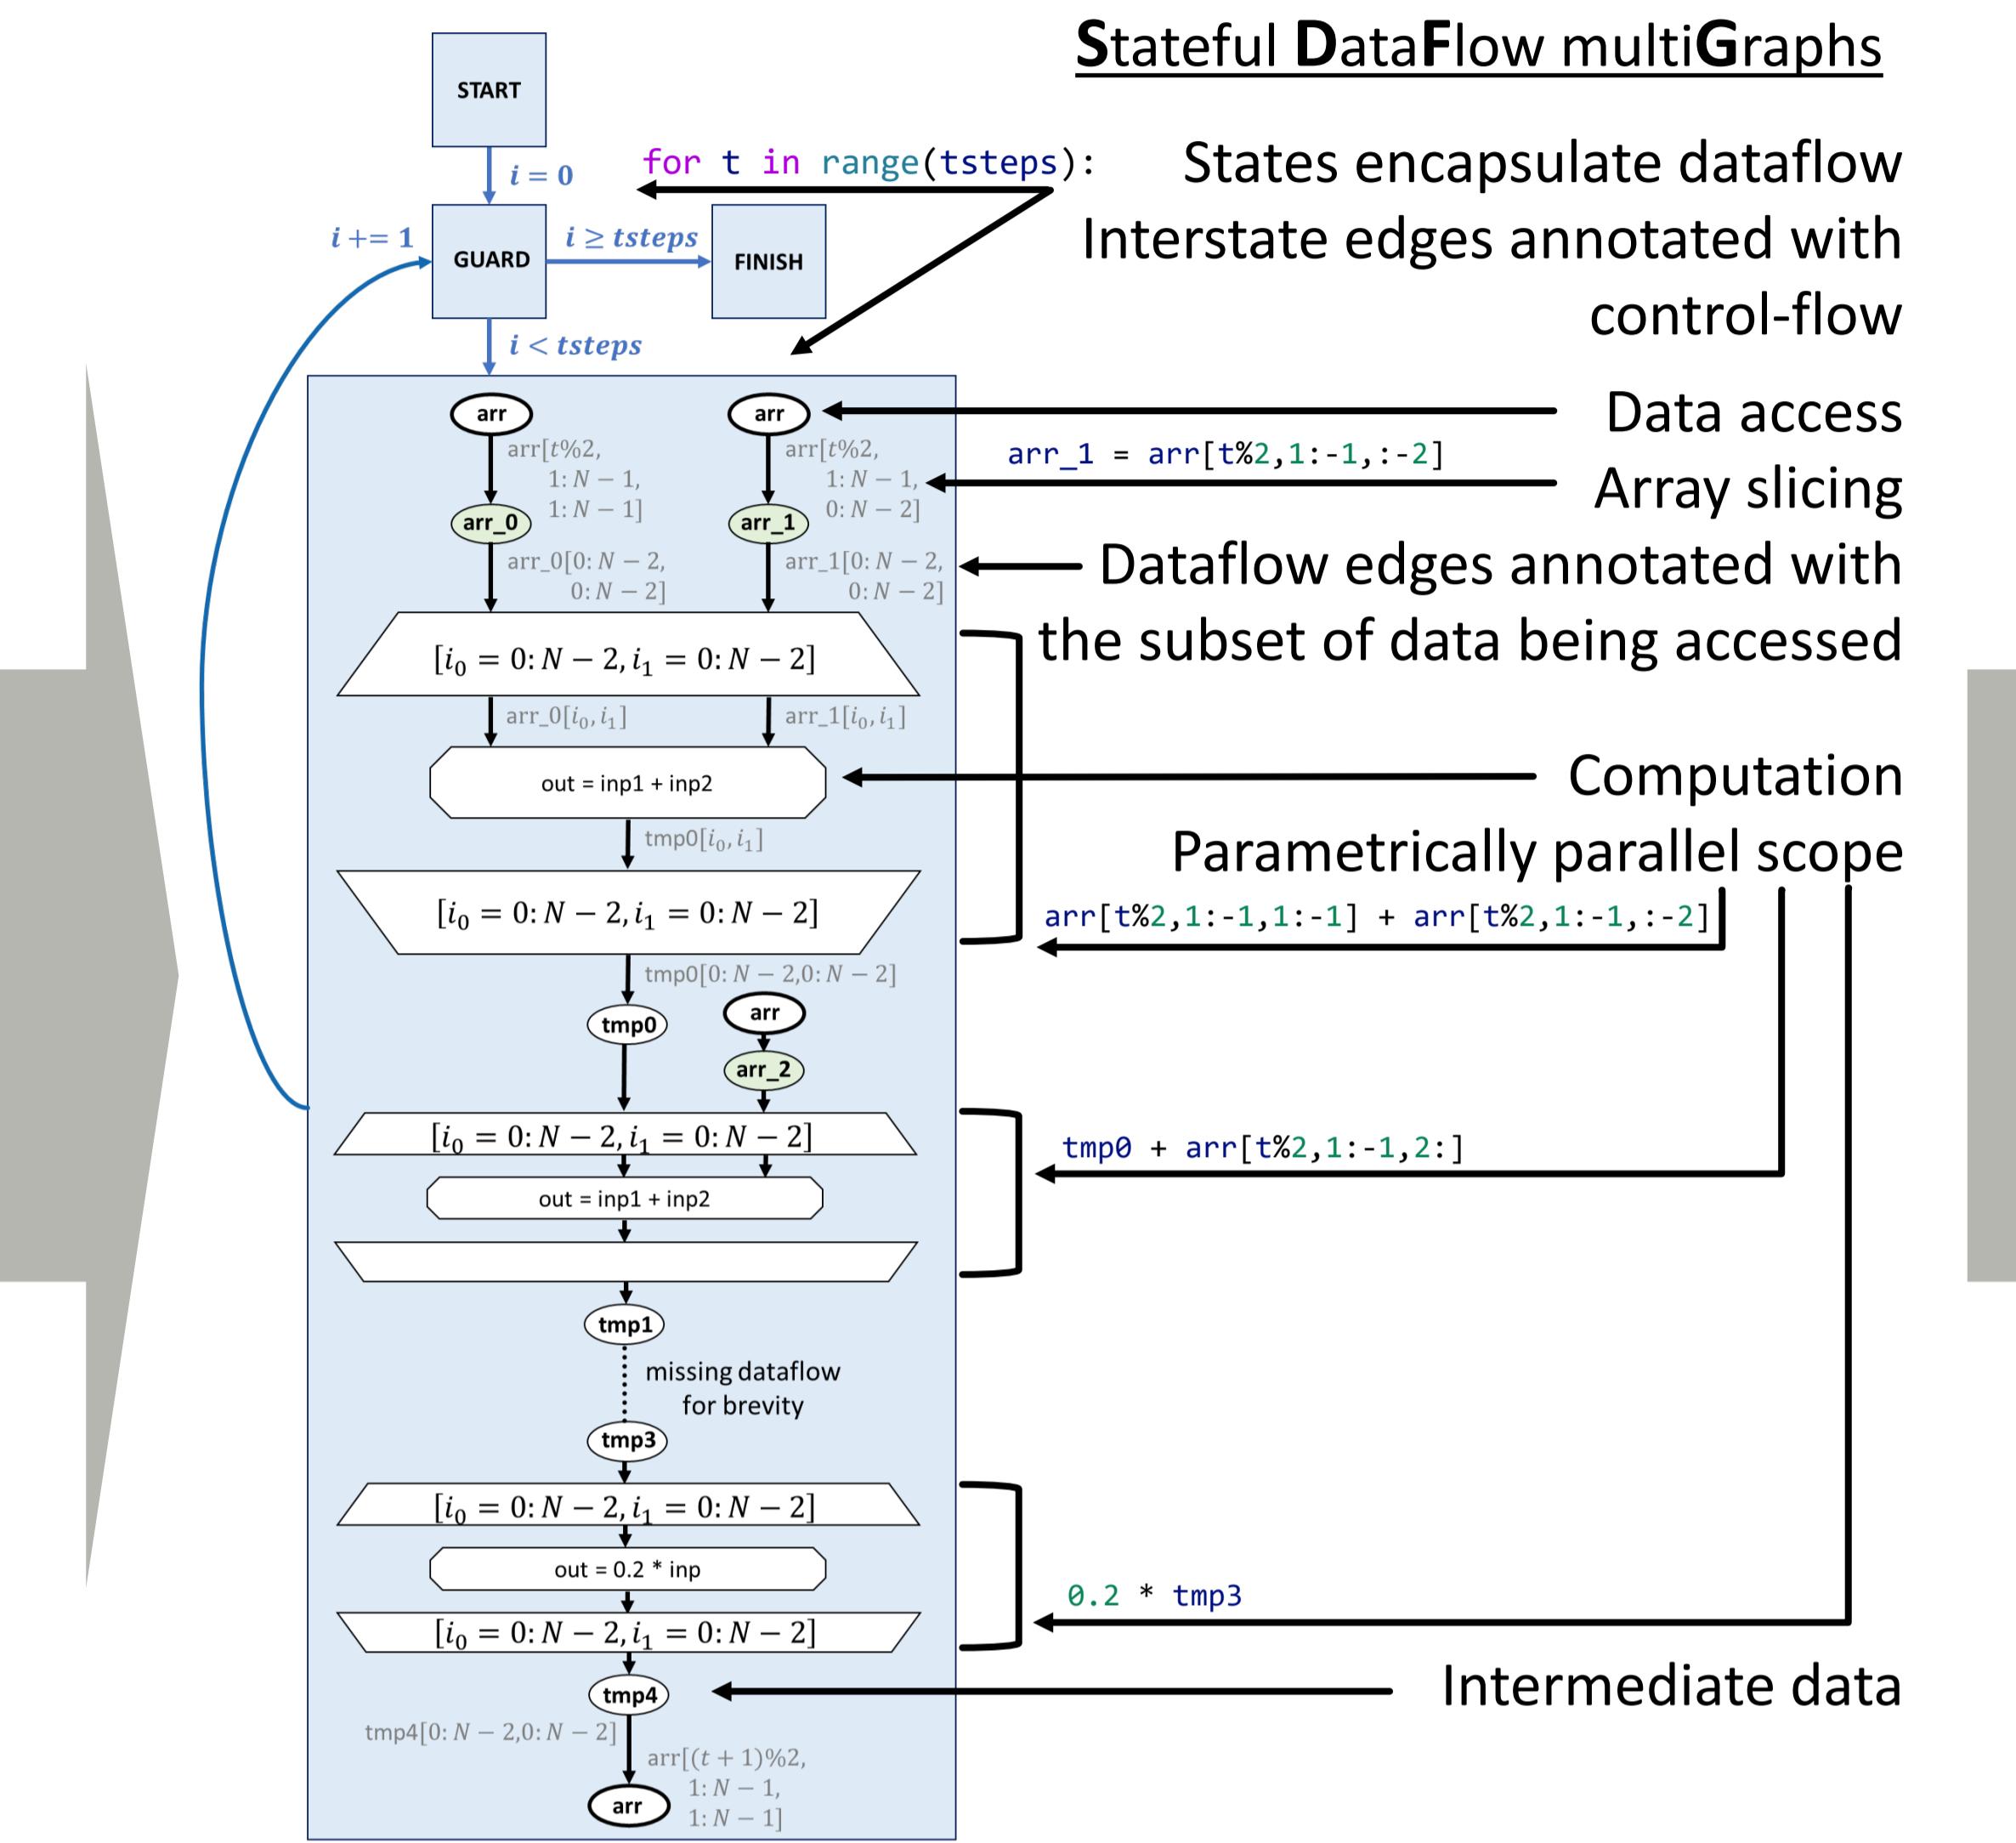
#### Python with type annotations for Ahead-Of-Time compilation

```
@dace.program
def jacobi2d(tsteps: dace.int32, arr: dace.float64[2, N, N]):
    for t in range(tsteps):
        arr[(t+1)%2, 1:-1, 1:-1] = 0.2 * (arr[t%2, 1:-1, 1:-1] +
                                             arr[t%2, 1:-1, -2] +
                                             arr[t%2, 1:-1, 2:] +
                                             arr[t%2, 2:, 1:-1] +
                                             arr[t%2, -2:, 1:-1])
```

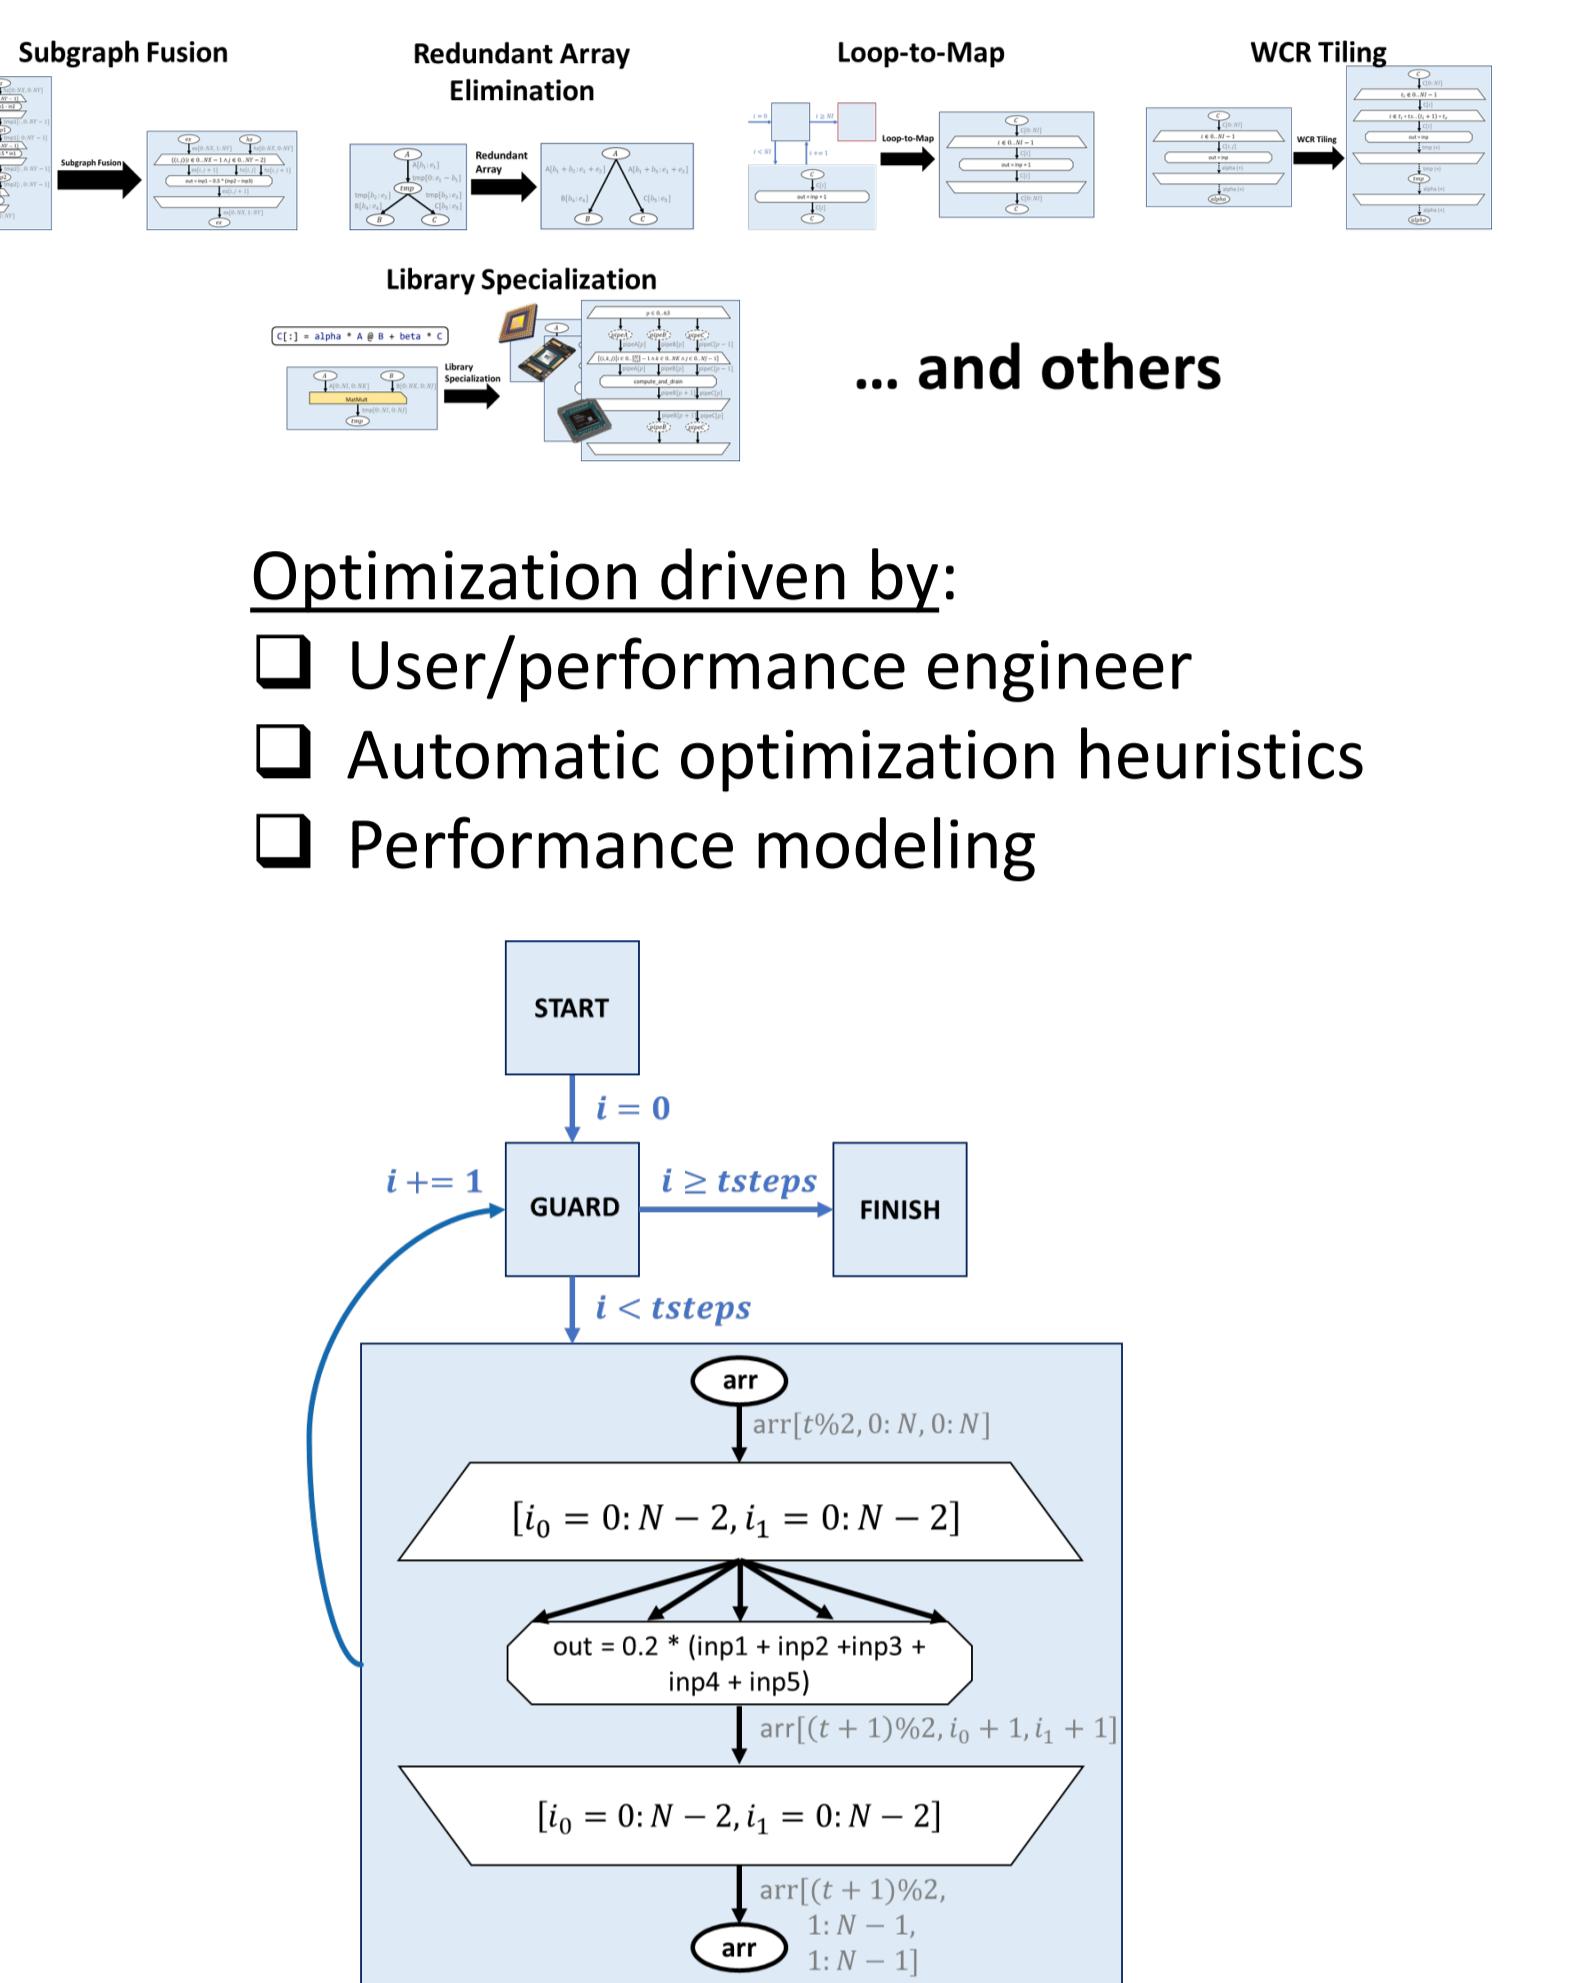
#### MPI-compatible syntax for distributed-memory programming

```
@dace.program
def jacobi2d_distr(tsteps: dace.int32, arr: dace.float64[2,Nx2,Ny+2]):
    req = np.empty((8,), dtype=dace.comm.Request)
    for t in range(tsteps):
        dace.comm.Isend(arr[t%2, 1, 1:-1], north, 0, req[0])
        ...
        dace.comm.Irecv(arr[t%2, 1:-1, -1], east, 2, req[7])
        dace.comm.Waitall(req)
        arr[(t+1)%2, 1+woff:-1-soff, 1+woff:-1-eoff] = 0.2 * (...)
```

### ... to a Data-Centric Representation ...



### ... optimized via Graph Transformations ...



### ... generating Architecture-Specific Code

**CPU**

```
for (t = 0; (t < tsteps); t = t + 1) {
    #pragma omp parallel for
    for (auto _io = 0; _io < (N - 2); _io += 1) {
        for (auto _in1 = arr[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1 + 1]; _in2_0 = arr[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1];
        double _in2_1 = arr[(((N * N) * (t % 2)) + (N * (_io + 2))) + _ii1 + 2];
        double _in2_0 = arr[(((N * N) * (t % 2)) + (N * (_io + 2))) + _ii1 + 1];
        double _in2_2 = arr[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1 + 1];
        double _out;
        _out = (0.2 * (((_in1 + _in2_0) + _in2_1) + _in2_0) + _in2);
        arr[(((N * N) * ((t + 1) % 2)) + (N * (_io + 1))) + _ii1 + 1] = _out;
    }
}
```

**GPU**

```
_global_ void outer_fused_0_0_2(double * __restrict__ gpu_arr, int N, long long t) {
    int _in1 = (blockIdx.x * 32 + threadIdx.x);
    int _io = (blockIdx.x * 1 + threadIdx.y);
    if (_in1 < (N - 2)) {
        double _in1 = gpu_arr[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1 + 1];
        double _in2_0 = gpu_arr[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1];
        double _in2_1 = gpu_arr[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1 + 2];
        double _in2_0 = gpu_arr[(((N * N) * (t % 2)) + (N * (_io + 2))) + _ii1 + 1];
        double _in2_2 = gpu_arr[(((N * N) * (t % 2)) + (N * (_io + 2))) + _ii1 + 1];
        double _out;
        _out = (0.2 * (((_in1 + _in2_0) + _in2_1) + _in2_0) + _in2));
        gpu_arr[(((N * N) * ((t + 1) % 2)) + (N * (_io + 1))) + _ii1 + 1] = _out;
    }
}
```

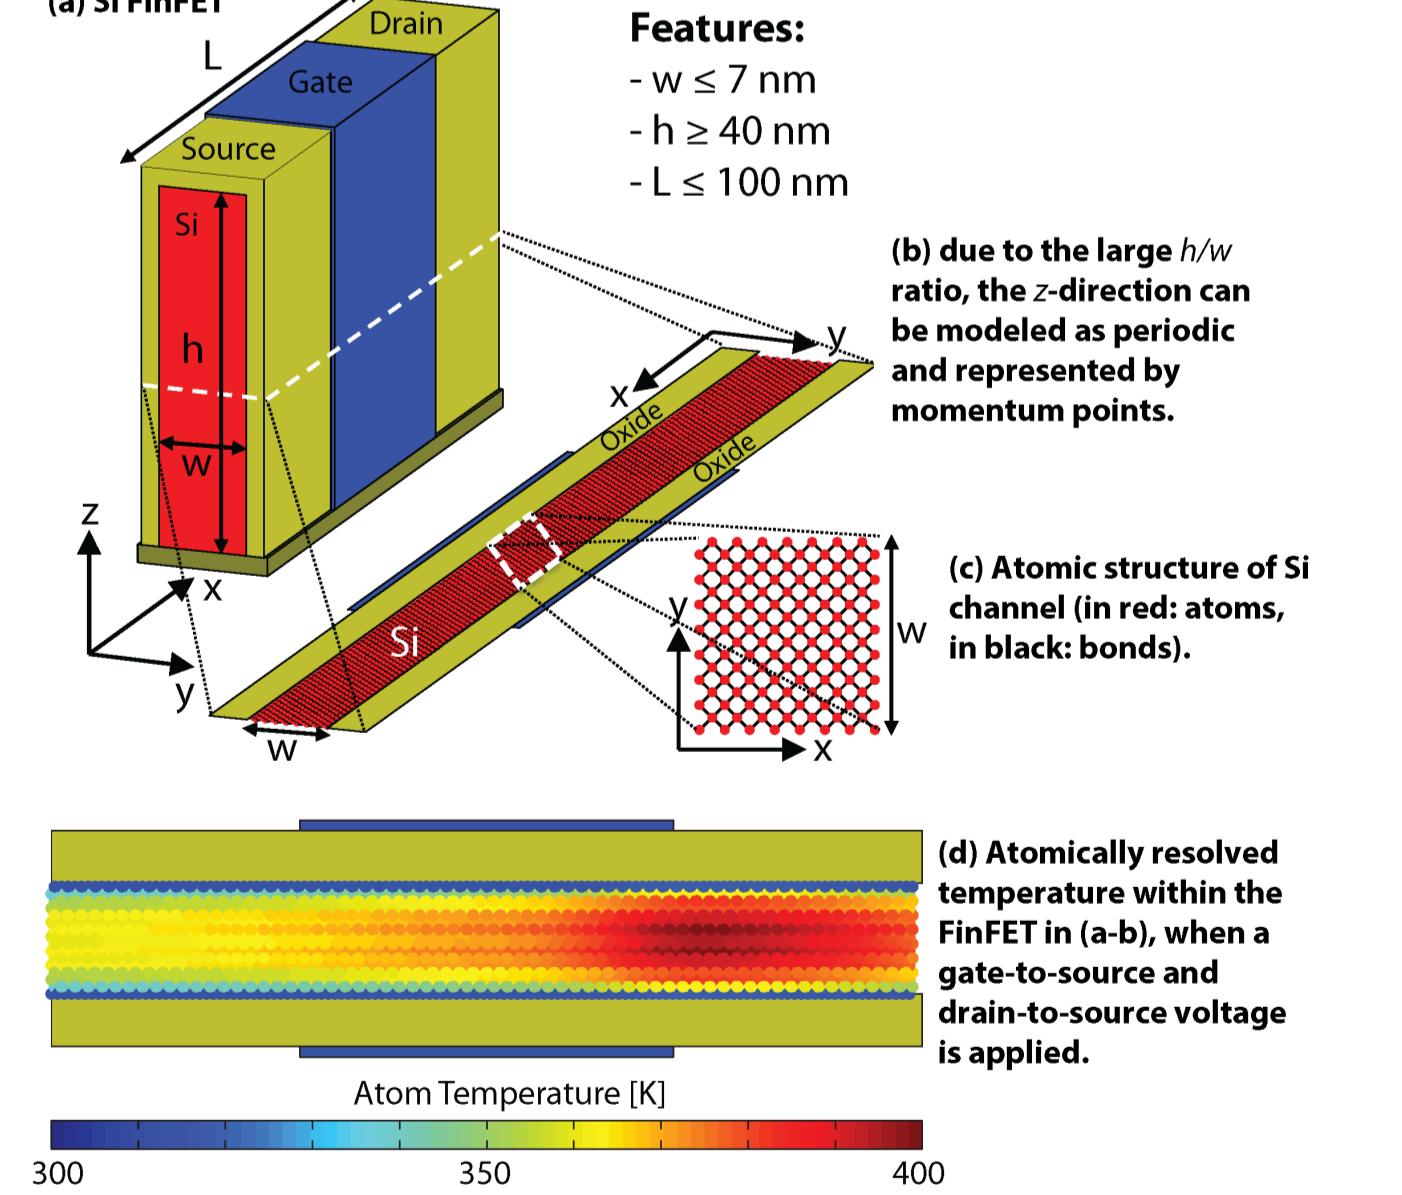
**FPGA**

```
for (t = 0; (t < tsteps_in); t = (t + 1)) {
    for (int _io = 0; _io < (N - 2); _io += 1) {
        #pragma HLS PIPELINE II=1
        #pragma HLS LOOP_FLATTEN
        {
            double _in1 = _arr_in[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1 + 1];
            double _in2_0 = _arr_in[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1];
            double _in2_1 = _arr_in[(((N * N) * (t % 2)) + (N * (_io + 2))) + _ii1 + 2];
            double _in2_0 = _arr_in[(((N * N) * (t % 2)) + (N * (_io + 2))) + _ii1 + 1];
            double _in2_2 = _arr_in[(((N * N) * (t % 2)) + (N * (_io + 1))) + _ii1 + 1];
            double _out;
            _out = (0.2 * (((_in1 + _in2_0) + _in2_1) + _in2_0) + _in2));
            *_arr_out + (((N * N) * ((t + 1) % 2)) + (N * (_io + 1))) + _ii1 + 1) = _out;
        }
    }
}
```

### Case Studies

#### Quantum Transport Simulation (ACM Gordon Bell Award 2019)

Implementation in Python and optimization of Quantum Transport Solver OMEN (original application written in C++)



#### NPBench

A benchmarking suite for Python, with over 50 code samples across a wide range of scientific domains.

