

# Programming Techniques for Supercomputers Tutorial

Erlangen National High Performance Computing Center

Department of Computer Science


FAU Erlangen-Nürnberg

Sommersemester 2026

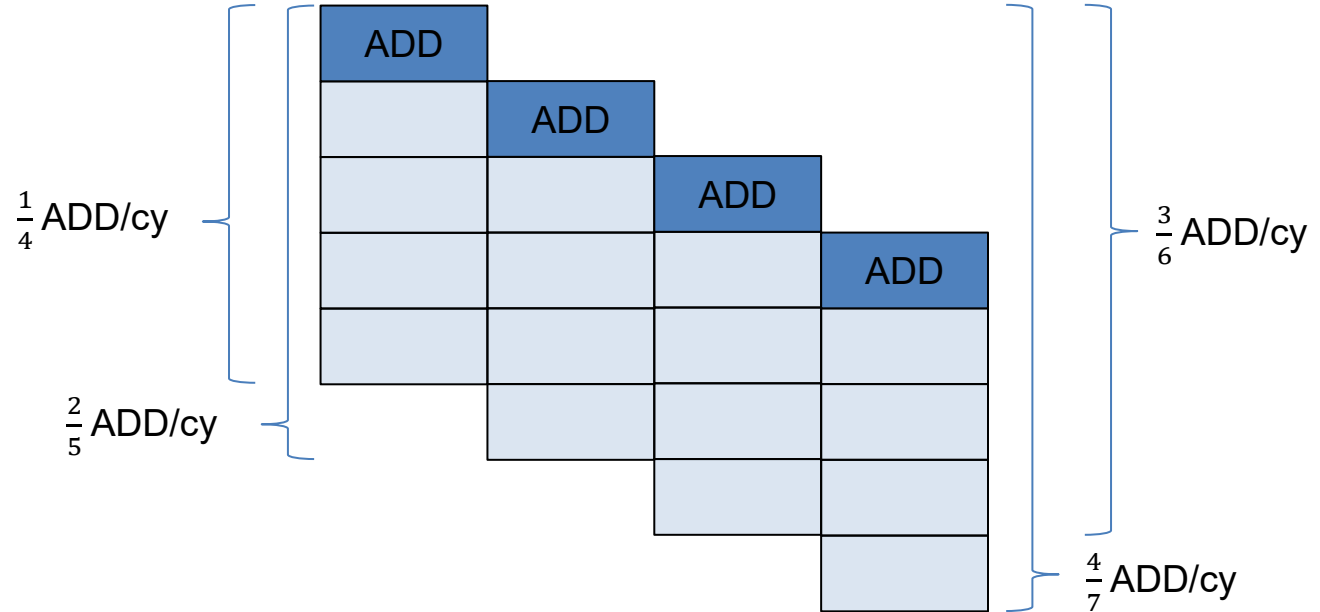


# Assignment 1 – Task 1

## Pipelines



- 4 cy latency
- 1 result per cy





$$TP_{pipe} = \frac{N}{N+m-1} \text{ with } TP_{pipe} = 0.9 \text{ inst/cy and } m = 4:$$

$$0.9(N + 3) = N \rightarrow 2.7 = 0.1N \rightarrow N = 27 \text{ with two pipelines: } \mathbf{N=54}$$

# Assignment 1 – Task 2

## More pipelines a)

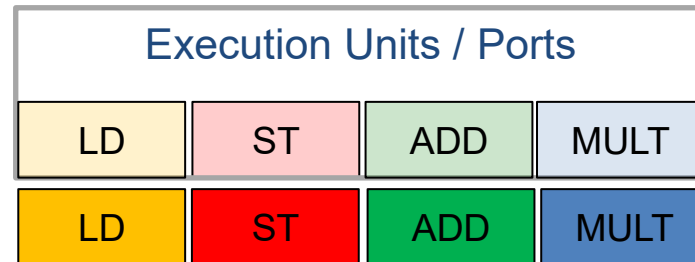
```
double a[...];
double s=0.1;
// a[], b[], c[] contain sensible data
for(int i=2; i<N; ++i) {
    a[i] += s * a[i-2];
}
for(int i=2; i<N; ++i) {
    b[i] += s * b[i-2];
}
for(int i=2; i<N; ++i) {
    c[i] += s * c[i-2];
}
```



# Assignment 1 – Task 2

## More pipelines a)

```
double a[...];  
double s=0.1;  
// a[], b[], c[] contain sensible data  
for(int i=2; i<N; ++i) {  
    a[i] += s * a[i-2];  
}  
//...
```



- 1 ADD, 6cy latency
- 1 MULT, 8cy latency
- 1 LOAD
- 1 STORE

Does this mean we can do 1 iter. per cycle (i.e., 2 flops in 1 cy)?

**NO**

Remember: Pipelining and dependencies from lecture (slide set 3a)



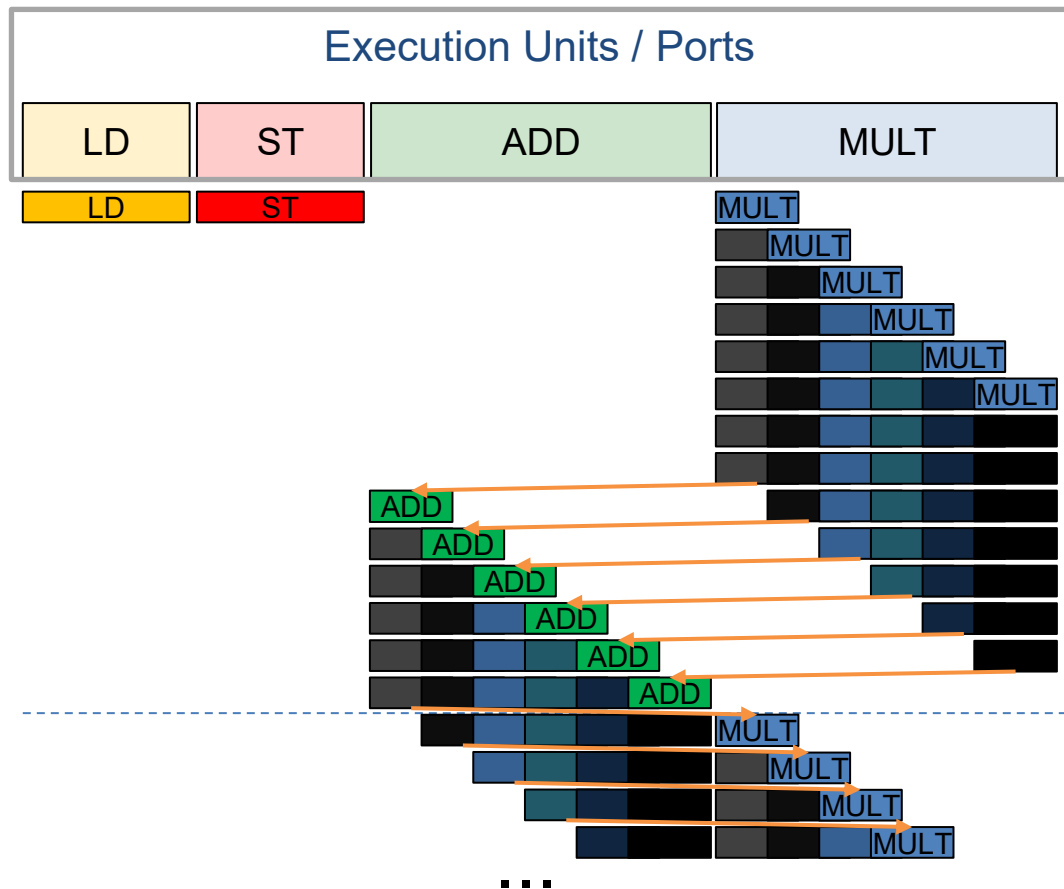
# Assignment 1 – Task 2

## More pipelines b)

```
double a[...];
double s=0.1;
// a[], b[], c[] contain sensible data
for(int i=2; i<N; ++i) {
    a[i] += s * a[i-2];
    b[i] += s * b[i-2];
    c[i] += s * c[i-2];
}
```

**6 instead of 2 independent loop-carried dependencies!**

$$\begin{aligned} \rightarrow P_{max} &= 6 \text{ iter}/14 \text{ cy} = 12 \text{ flops}/14 \text{ cy} \\ &= 0.857 \text{ flops/cy} \end{aligned}$$



# Assignment 1 – Task 3

## Square root in depth a)

```
sum = 0.;
for(int i=0; i<N; i++) {
    x = (i + 0.5) * delta_x;
    sum += 4.0 * sqrt(1.0 - x*x);
}
```



```
sum = 0.; sum1 = 0.; sum2 = 0.;
for(int i=0; i<N; i+=2) {
    x1 = (i + 0.5) * delta_x;
    x2 = (i+1+0.5) * delta_x;
    sum1 += 4.0 * sqrt(1.0 - x1*x1);
    sum2 += 4.0 * sqrt(1.0 - x2*x2);
}
sum = sum1 + sum2;
```

- Bottleneck: SQRT with 22cy latency and 3cy reciprocal throughput
- BUT: Modulo Variable Expansion (MVE) can be beneficial  
→ **different order of accumulation** → SQRTs can overlap → hide latency
- **partial sums must be reduced to one after the loop**

# Assignment 1 – Task 3

## Square root in depth b)

```
sum = 0.;  
for(int i=0; i<N; i++) {  
    x = (i + 0.5) * delta_x;  
    sum += 4.0 * sqrt(1.0 - x*x);  
}
```

- int2float
- 3 ADDs
- 3 MULTs
- 1 SQRT
- 3 instr. loop mechanics

IPC = 11 instr/6 cy  
= 1.83 instr/cy

Measurement (N=10<sup>9</sup>): 6.02cy

```
.LBB1_7:  
vaddsd    %xmm5, %xmm2, %xmm3  
vmulsd    %xmm0, %xmm3, %xmm3  
vfnmadd213sd %xmm4, %xmm3, %xmm3  
vsqrtsd   %xmm3, %xmm3, %xmm3  
vfmadd231sd %xmm6, %xmm3, %xmm1  
vaddsd    %xmm4, %xmm2, %xmm2  
decl      %eax  
jne       .LBB1_7
```

IPC = 8 instr/6 cy  
= 1.33 instr/cy