



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Application Specific Computing (ASC)


Robert Strzodka

Application Specific Computing
ZITI, Heidelberg University

<http://asc.ziti.uni-heidelberg.de>


Application Specific Computing (ASC)

Accelerator



Data

Many-core




Algorithm

FPGAs



Embedded



Software

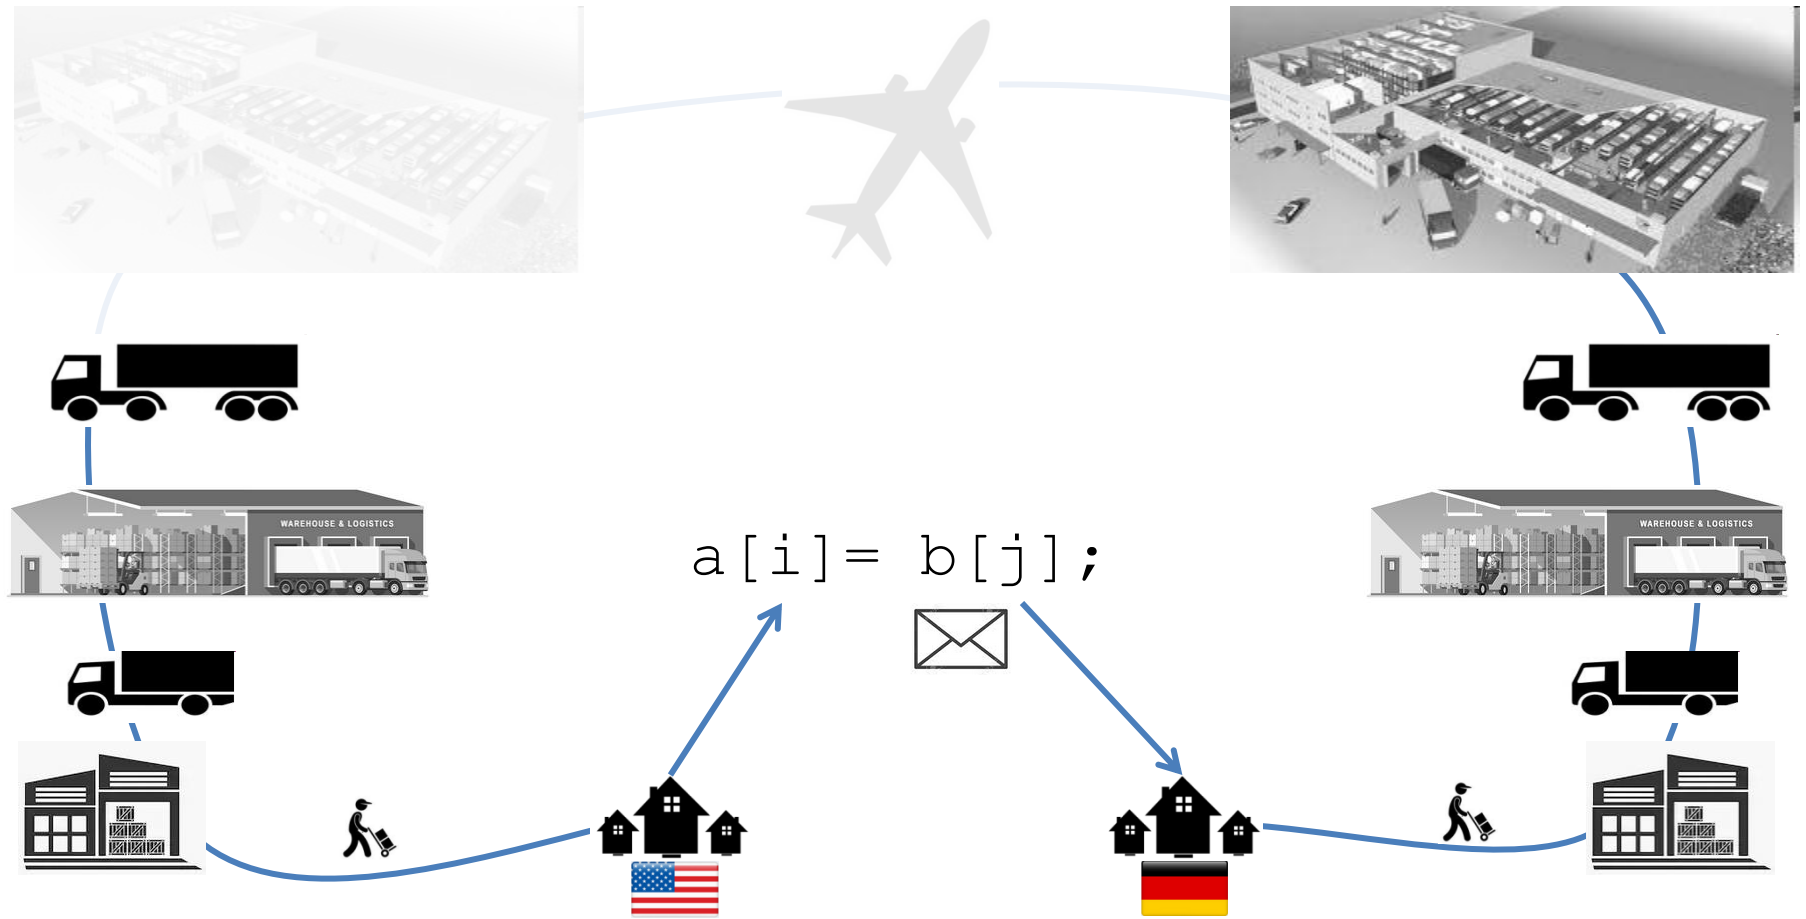
#	T	Site	Manufacturer	Computer	Country	HPCG [Pflop/s]	Rmax [Pflop/s]	HPCG/ Peak	HPCG/ HPL
1	1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.07 GHz, Volta GV100, EDR	USA	2.9258	122.3	1.6%	2.4%
2	3	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1 GHz, Volta GV100, EDR	USA	1.7957	71.6	1.5%	2.5%
3	16	RIKEN Advanced Institute for Computational Science	Fujitsu	K Computer SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	0.6027	10.5	5.3%	5.7%
4	9	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries	USA	0.5461	14.1	1.2%	3.9%
5	6	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	0.4864	19.6	1.9%	2.5%
6	2	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	0.4808	93.0	0.4%	0.5%
7	12	JCAHPC Joint Center for Advanced HPC	Fujitsu	Oakforest-PACS PRIMERGY CX1640 M1, Intel Xeons Phi 7250 68C 1.4 GHz, OmniPath	Japan	0.3855	13.6	1.5%	2.8%
8	10	Lawrence Berkeley National Laboratory	Cray	Cori Cray XC40, Intel Xeons Phi 7250 68C 1.4 GHz, Aries	USA	0.3554	14.0	1.3%	2.5%
9	14	Commissariat a l'Energie Atomique (CEA)	Bull	Tera-1000-2 Bull Sequana X1000, Intel Xeon Phi 7250 68C 1.4 GHz, Bull BXI 1.2	France	0.3338	12.0	1.4%	2.8%
10	8	Lawrence Livermore National Laboratory	IBM	Sequoia BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	0.3304	17.2	1.6%	1.9%

Source: www.top500.org

Why is Everything so Slow?

- **Data:** `a[i] = b[j]`
 - Number representation
 - Data ordering and layout
- **Loops:** `for (...)`
 - Tilings
 - Hierarchies
 - Fusion
- **Branches:** `if (...)`
 - Extraction
 - Pre-evaluation
 - Predication

Single Data Access as Postal Delivery



Loops, single access

Foreach and Tiling

```
for (i=0; i<N; ++i) {  
    S(i);  
}
```

- Thread parallelism in outer loop
- SIMD parallelism in inner loop

- Many different tiling choices


```
foreach ( i $\in$ [0,N) ) {  
    S(i);  
}
```

```
foreach ( i $\in$ [0,N/T) ) {  
    foreach ( k $\in$ [0,T) ) {  
        S(k+i*T);  
    }  
}
```

Tiling with Dependencies

```
for (i=0; i<N; ++i) {
    S(i);
}
```

```
foreach ( i ∈ [0, N/T) ) {
    for (k=0; k<T; ++k) {
        S(k+i*T);
    }
}
```



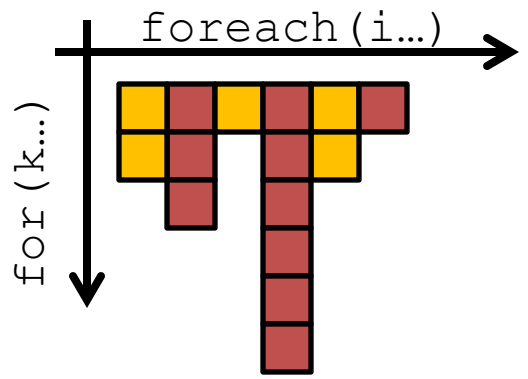
```
for (k=0; k<T; ++k) {
    foreach ( i ∈ [0, N/T) ) {
        S(k+i*T);
    }
}
```


Segments

```

foreach ( i ∈ [0, M) ) {
  for (k=0; k < L(i); ++k) {
    S(i, k);
  }
}

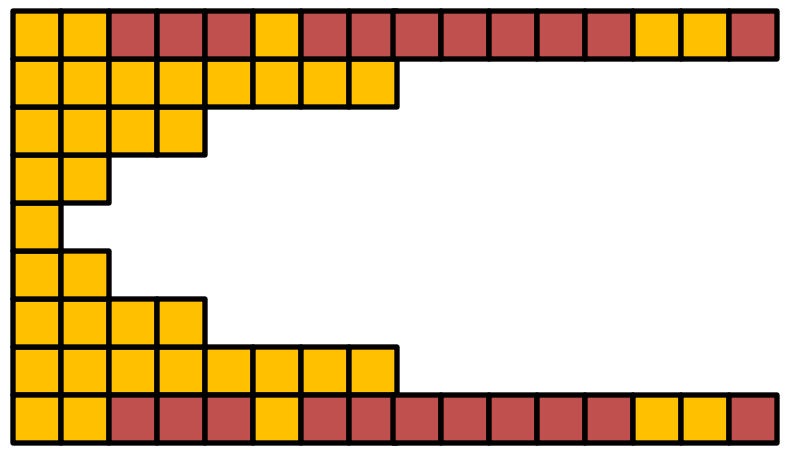
```



```

for (lev...) {
  foreach (k...) {
    P(lev, k);
  }
}

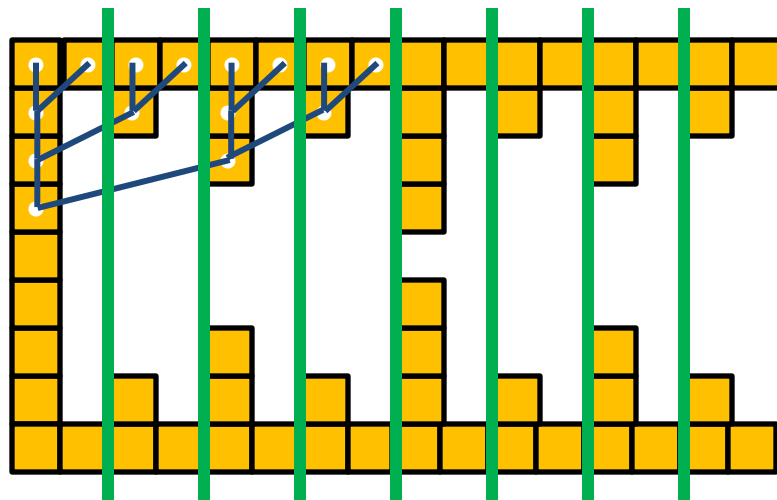

```



Hierarchies and Communication

```
for (lev...) {  
  foreach (k...) {  
    P(lev, k);  
  }  
}
```

- Reduction, scan, barrier
- Mipmaps, wavelets, multigrid
- Data distribution important




Loops, multiple accesses

Fusion of Two

```

foreach (  $i \in I_0$  ) {
     $S_0(i)$ ;
}
foreach (  $i \in I_1$  ) {
     $S_1(i)$ ;
}

```



```

foreach (  $i \in I_0 \setminus I_1$  ) {
     $S_0(i)$ ;
}
foreach (  $i \in I_0 \cap I_1$  ) {
     $S_0(i)$ ;  $S_1(i)$ ;
}
foreach (  $i \in I_1 \setminus I_0$  ) {
     $S_1(i)$ ;
}


```

Fusion of Many

```

for (n=0; n<N; ++n) {
  foreach ( i ∈ In ) {
    Sn(i);
  }
}

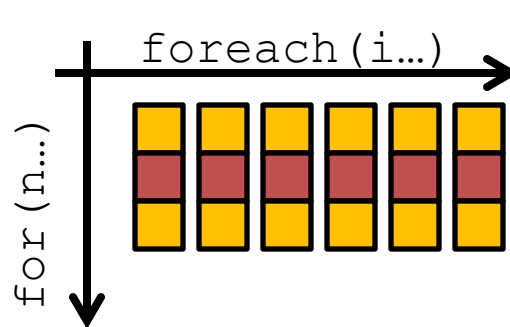
```



```

...
foreach ( i ∈ I0 ∩ ... ∩ IN-1 ) {
  for (n=0; n<N; ++n) {
    Sn(i);
  }
}
...

```




Fusion with Neighbor Dependencies

```

for (n=0; n<N; ++n) {
  foreach ( i∈I ) {
    Sn(i);
  }
}

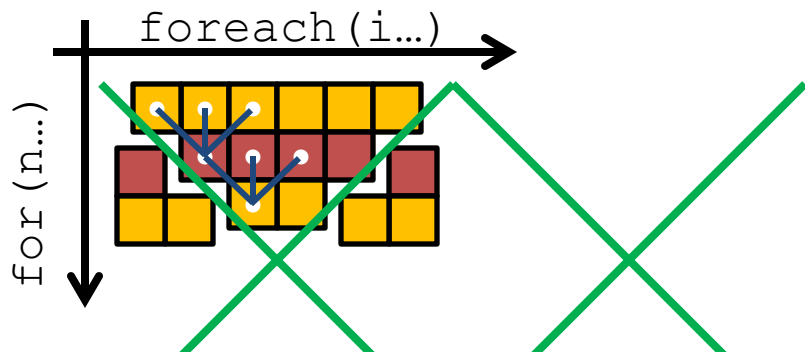
```



```

foreach (tile...) {
  for (n...) {
    foreach ( i∈tile ) {
      Sn(i);
    }
  }
}

```





Fusion with Reductions

```

foreach (  $I \rightarrow I$  ) { ... }
for (  $I \rightarrow \alpha$  ) { ... }
foreach (  $I \rightarrow I$  ) { ...  $\alpha$  ... }
  
```

```

foreach (  $I \rightarrow I$  ) { ... }
foreach (  $I \rightarrow I$  ) { ... }
for (  $I \rightarrow \alpha$  ) { ... }
...  $\alpha$  ...
  
```



Branches

Extraction

```
foreach (  $i \in I$  ) {  
  if (  $\text{cond}_0(i)$  ) {  
     $S_0(i)$  ;  
  } else if (  $\text{cond}_1(i)$  ) {  
     $S_1(i)$  ;  
  } else {  
     $S_2(i)$  ;  
  }  
}
```

```
foreach (  $i \in I_0(\text{cond}_0)$  ) {  
   $S_0(i)$  ;  
}  
foreach (  $i \in I_1(\text{cond}_1)$  ) {  
   $S_1(i)$  ;  
}  
foreach (  $i \in I_2(\text{cond}_2)$  ) {  
   $S_2(i)$  ;  
}
```

Pre-Evaluation

```
foreach(  $i \in I$  ) {  
    if(  $cond_0(i)$  ) {  
         $S_0(i)$ ;  
    } else if(  $cond_1(i)$  ) {  
         $S_1(i)$ ;  
    } else {  
         $S_2(i)$ ;  
    }  
}  
  
compute(  $cond_0[]$ ,  $cond_1[]$  );  
  
foreach(  $i \in I$  ) {  
    if(  $cond_0[i]$  ) {  
         $S_0(i)$ ;  
    } else if(  $cond_1[i]$  ) {  
         $S_1(i)$ ;  
    } else {  
         $S_2(i)$ ;  
    }  
}
```

Predication

```
foreach (  $i \in I$  ) {  
    if (  $\text{cond}_0(i)$  ) {  
         $v = S_0(i)$ ;  
    } else if (  $\text{cond}_1(i)$  ) {  
         $v = S_1(i)$ ;  
    } else {  
         $v = S_2(i)$ ;  
    }  
}
```

```
foreach (  $i \in I$  ) {  
     $v_0 = S_0(i)$ ;  
     $v_1 = S_1(i)$ ;  
     $v_2 = S_2(i)$ ;  
     $v = \text{selector}(i, v_0, v_1, v_2)$ ;  
}
```

Simple Rules

Golden Rules for High Performance Code

- Do not use branches
- Do not use loops
- Do not copy data

~~if(...) {}~~

~~for(...) {}~~

~~A= B;~~

Rules of Thumb for High Performance Code

- **Avoid** branches in loops `if(...) {}`
- **Avoid** general loops `for(...) {}`
- **Avoid** data copies `A= B;`
- **Employ** clever algorithms

Conclusions

- Hardware(HW) – Software(SW)
 - Hardware designs highly specialized
 - Language constructs too general
 - Mismatch of SW to HW
- Algorithms
 - Important to avoid obvious dongs
 - Some very efficient patterns exist
 - General cases execute poorly
- **Many things to be discovered !**



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Application Specific Computing (ASC)

Robert Strzodka

Application Specific Computing
ZITI, Heidelberg University

<http://asc.ziti.uni-heidelberg.de>