High-Performance Matrix Computations

Prof. **Paolo Bientinesi**

pauldj@aices.rwth-aachen.de

Lecture 1: Computer architecture

- CPU, ALU = arithmetic logic unit
- frequency, cycles, ticks
- ALU - latency, throughput (addition vs. multiplication)
  pipelining
- multiplication by 1, multiplication by 0
  see code in the next slide
- registers, memory hierarchy, caches
  size vs access time
- cache line, (hardware/software) prefetching

Assembly

```
int main ( )
{
   double a = 343.5436523;
   double b = a * 0.0;
   return b;
}
```

```
int main ( int argc, char *argv[] )
{
   double a = atof(argv[1]);
   double b = a * 0.0;
   return b;
}
```

Compile with `gcc -S`; compare `-O0` with `-O2`