

# OpenCL lab 02

## 1 Task 1 (identical to Task5 from previous sheet)

Implement matrix multiply using openCL. You may use `matmul.c` as a starting point. Play with the parameters of `matmul` and observe their effect on the runtime.

## 2 Task 2

Try to improve your matrix multiply.

- Look at the effect of using a different thread mapping.
- On conventional machines, the *star* algorithm has superior performance. It computes the matrix product of the two matrices **a** and **b** by:

```
for( i=0, i < count; i++) {
    for( j=0, j < count; j++) {
        a_i_j = a[i][j];
        for( k=0, k < count; k++) {
            c[i][k] += a_i_j * b[j][k] ;
        }
    }
}
```

Implement this algorithm in openCL and compare its performance to the previous versions. Think about how you want to map your threads. How does the best version that you create compare to the earlier versions? Can you think of reasons for this behaviour?