

02_cache/02_mm-vs-mv

Hands-on: Matrix-matrix product: Well-tuned DGEMM vs. others

- Choose either C/C++ (`c.boost_eco`) or Fortran (`fortran.boost_eco`) samples. Both of them are fine, as well.
 - `boost_eco`: run a job using the boost eco mode (`freq=2200, eco_state=2`)
 - `normal`: run a job using the normal mode (`normal=2000, eco_state=0`)

C/C++

How to compile and how to execute

1. Run a script of creating working space

- edit `task.sh` (modify `--gname` option).
- See `objst` in `create_project.sh`.
 - It indicates a list of the settings of compile options.
- Running `create_project.sh`, the executable file (`run.x`) is generated in `Obj_***` directory.
 - This directory is also automatically generated.
 - When you desire to use the own compiler, set `MAKE_DIR` variable in `create_project.sh` from `config` to `config.own`.
- Example:

```
$ cd c.boost_eco
$ bash create_project.sh
$ ls
Obj_fj-ssl2  Obj_fj-ssl2so
```

2. Run program

- A job script to execute the program is located in `Obj_***/results`.
- You can run the program either:

```
## To run as a batch job
$ cd Obj_fj-ssl2/results
$ pjsub task.sh
## Or, to run in an interactive job
$ cd Obj_fj-ssl2/results
$ bash task.sh
```

- Each of the samples in the Exercises will be completed within 6 minutes.
 - For safety, we set the elapsed time of the job script as 10 minutes.

- One job script contains the execution with different kinds of matrix dimension. We treat only square matrices.
 - On the other hand, in `c.fapp` we fix a single specific matrix dimension.

Exercises A

- E1: Examine difference between a direct use of DGEMM and other implementations of matrix-matrix product in `main.c`.
- E2: Compare the performance (GFlop/s) of DGEMM to that in the others when changing matrix dimension, in `Obj_fj-ssl2/`.
 - Also, compare those to the peak value of GFlop/s of the single core in the boost eco mode.

Exercises B (advanced)

- E3: Try to use the dynamic library of Fujitsu SSL2, instead of the static one. Check `config/Makefile.fj-ssl2so`.
- E4: In `c.normal`, `task.sh` file is set to use the normal mode (`--rsc-list "freq=2000, eco_state=0"`). Run jobs in the directory `c.normal` and compare the performance with those in directory `c.boost_eco`.
- E5: In `c.fapp`, perform the basic CPUPA analysis using `fapp`. Using the CPUPA reports, compare `dgemm` to `dgemv` from the viewpoints of use of cache.
 - On the basic CPUPA, `fapp`-measurements are performed five times (i.e., `-Hevent=pa1`, `-Hevent=pa2`, `-Hevent=pa3`, `-Hevent=pa4`, and `-Hevent=pa5`).
 - Run the job with `task-Bpa.sh`. After the job completion, perform analyses using `perf-Bpa.sh`. Using the resultant CSV files and `cpu_pa_report.xlsm`, you can obtain the CPUPA report.
 - Indeed, you can find that there are various kinds of difference between `dgemm` and `dgemv`, other than cache.
- E6: Check the value of floating-point operation peak ratio in the CPUPA reports in `c.fapp`. Compare the values of the ratios between the normal mode and the boost eco mode.

Fortran

How to compile and how to execute

1. Run a script of creating working space

- edit `task.sh` (modify `--gname` option).
- See `objst` in `create_project.sh`.
 - It indicates a list of the settings of compile options.
- Running `create_project.sh`, the executable file (`run.x`) is generated in `Obj_***` directory.
 - This directory is also automatically generated.
 - When you desire to use the own compiler, set `MAKE_DIR` variable in `create_project.sh` from `config` to `config.own`.
- Example:

```
$ cd fortran.boost_eco
$ bash create_project.sh
```

```
$ ls
Obj_fj-ssl2  Obj_fj-ssl2.fast  Obj_fj-ssl2so
```

2. Run program

- A job script to execute the program is located in `Obj_***/results`.
- You can run the program either:

```
## To run as a batch job
$ cd Obj_fj-ssl2/results
$ pjsub task.sh
## Or, to run in an interactive job
$ cd Obj_fj-ssl2/results
$ bash task.sh
```

- Each of the samples in the Exercises will be completed within 5 minutes.
 - For safety, we set the elapsed time of the job script as 10 minutes.
 - One job script contains the execution with different kinds of matrix dimension. We treat only square matrices.
 - On the other hand, in `fortran.fapp` we fix a single specific matrix dimension.

Exercises A

- E1: Examine difference between a direct use of DGEMM and other implementations of matrix-matrix product in `main.F90`.
- E2: Compare the performance (GFlop/s) of DGEMM to that in the others when changing matrix dimension, in `Obj_fj-ssl2/`. Also, compare those to the peak value of GFlop/s of the single core in the boost eco mode.

Exercises B (advanced)

- E3: Try the case of `fortran/Obj_fj-ssl2.fast`, in which `-Kfast` option is used. Are there any differences from the case of `Obj_fj-ssl2`?
- E4: Try to use the dynamic library of Fujitsu SSL2, instead of the static one. Check `confing/Makefile.fj-ssl2so`.
- E5: In `fortran.normal`, `task.sh` file is set to use the normal mode (`--rsc-list "freq=2000, eco_state=0"`). Run jobs in the directory `fortran.normal` and compare the performance with those in directory `fortran.boost_eco`.
- E6: In `fortran.fapp`, perform the basic CPUPA analysis using `fapp`. Using the CPUPA reports, compare `dgemm` to `dgemv` from the viewpoints of use of cache.
 - On the basic CPUPA, `fapp`-measurements are performed five times (i.e., `-Hevent=pa1`, `-Hevent=pa2`, `-Hevent=pa3`, `-Hevent=pa4`, and `-Hevent=pa5`).
 - Run the job with `task-Bpa.sh`. After the job completion, perform analyses using `perf-Bpa.sh`. Using the resultant CSV files and `cpu_pa_report.xlsm`, you can obtain the CPUPA report.
 - Indeed, you can find that there are various kinds of difference between `dgemm` and `dgemv`, other than cache.

- E7: Check the value of floating-point operation peak ratio in the CPUPA reports in `fortran.fapp`. Compare the values of the ratios between the normal mode and the boost eco mode.
- E8: Consider the case of using Fortran `matmul` function as a matrix-matrix-product routine.
 - See `fortran/config/Makefile.fj-ssl2.mamul`. We have two cases:
 - `matmul` with `-Nalloc_assign -O3` option. The job execution time might become longer. (~10 minutes)
 - `matmul` with `-Kfast`, but not specifying `-Nalloc_assign`.