A collaborative peer review process for grading coding assignments in coursework

Pratik Nayak, Fritz Göbel, Hartwig Anzt

Karlsruhe Institute of Technology

June 18th
Workshop on Teaching Computational Science, ICCS 2021

Outline

- Why ?
- What ?
- Where ?
- How ?

Pratik Nayak, nayak@kit.edu

Outline

- Why ?
- What ?
- Where ?
- How ?
- Feedback
- Conclusions and Perspectives

Pratik Nayak, nayak@kit.edu

Why ?

- Almost all research employs some form of software.
- Software lifecycle often exceeds hardware lifecycles.
- Good sustainable software is **THE** key component of computational science.
- Ingraining good software practices in students is important to their careers in industry and in academia.

Pratik Nayak, nayak@kit.edu

What ?

- Version control.
- Continuous Integration.
- Automated testing
- Collaborative peer review.

Pratik Nayak, nayak@kit.edu

Version Control

- Offload the code “version” management.
- Eases collaboration and parallel work on a single code base.
- Examples: SVN, Mercurial, git, Perforce etc.

Pratik Nayak, nayak@kit.edu

Version Control

- “branch” your work from the main codebase.
- “commit” your distinct changes.
- Request your changes to be merged back into the main codebase.

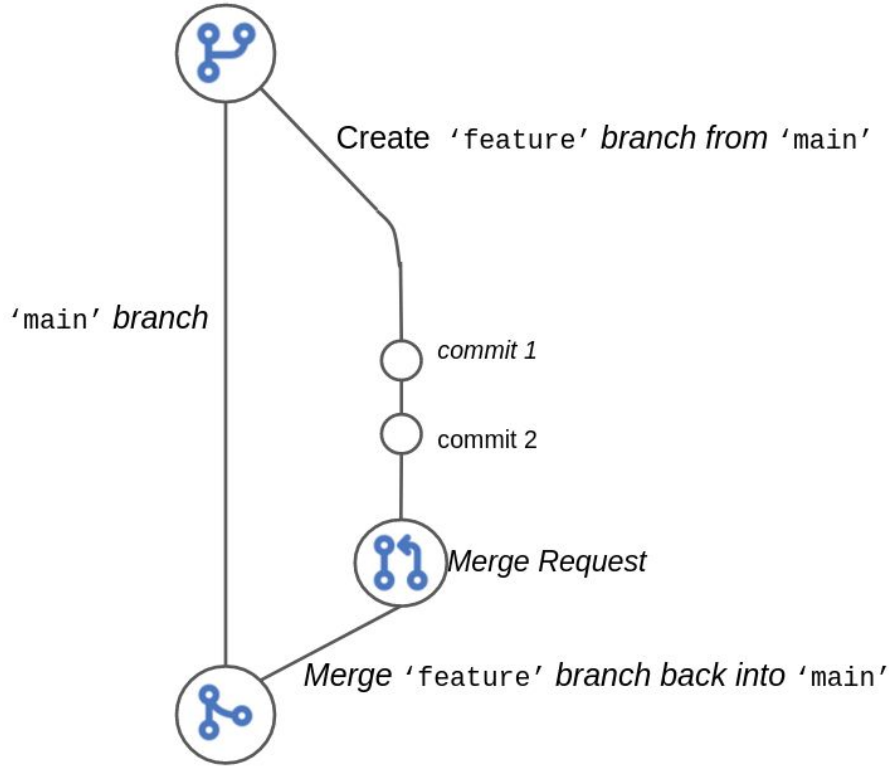

A common git workflow

Pratik Nayak, nayak@kit.edu

Automated Testing

- Verify code correctness and robustness.
- Allow easy identification of bugs in large code-bases.
- Streamline development and integrate verification into development.

```
208: [  PASSED  ] 32 tests.
208/209 Test #208: core/test/utils/unsort_matrix_test ..............   Passed    0.08 sec
test 209
        Start 209: core/test/utils/value_generator_test
209: Test command: /builds/ginkgo-project/ginkgo-public-ci/build/amd/clang/hip/release/static/c
209: Test timeout computed to be: 1500
209: Running main() from _deps/googletest-src/googletest/src/gtest_main.cc
209: [==========] Running 4 tests from 4 test suites.
209: [----------] Global test environment set-up.
209: [----------] 1 test from ValueGenerator/0, where TypeParam = float
209: [ RUN      ] ValueGenerator/0.OutputHasCorrectAverageAndDeviation
209: [       OK ] ValueGenerator/0.OutputHasCorrectAverageAndDeviation (0 ms)
209: [----------] 1 test from ValueGenerator/0 (0 ms total)
209:
209: [----------] 1 test from ValueGenerator/1, where TypeParam = double
209: [ RUN      ] ValueGenerator/1.OutputHasCorrectAverageAndDeviation
209: [       OK ] ValueGenerator/1.OutputHasCorrectAverageAndDeviation (0 ms)
209: [----------] 1 test from ValueGenerator/1 (0 ms total)
209:
209: [----------] 1 test from ValueGenerator/2, where TypeParam = std::complex<float>
209: [ RUN      ] ValueGenerator/2.OutputHasCorrectAverageAndDeviation
209: [       OK ] ValueGenerator/2.OutputHasCorrectAverageAndDeviation (1 ms)
209: [----------] 1 test from ValueGenerator/2 (1 ms total)
209:
209: [----------] 1 test from ValueGenerator/3, where TypeParam = std::complex<double>
209: [ RUN      ] ValueGenerator/3.OutputHasCorrectAverageAndDeviation
209: [       OK ] ValueGenerator/3.OutputHasCorrectAverageAndDeviation (0 ms)
209: [----------] 1 test from ValueGenerator/3 (0 ms total)
209:
209: [----------] Global test environment tear-down
209: [==========] 4 tests from 4 test suites ran. (1 ms total)
209: [  PASSED  ] 4 tests.
209/209 Test #209: core/test/utils/value_generator_test .............   Passed    0.05 sec
100% tests passed, 0 tests failed out of 209
Total Test time (real) = 297.79 sec
```

Continuous integration

- Compile and test code on a variety of platforms and machines.
- Automatically identify breaking changes.
- Perform static analyses and ensure high code quality.
- Verify code uniformity and formatting
- Deploy code and documentation.

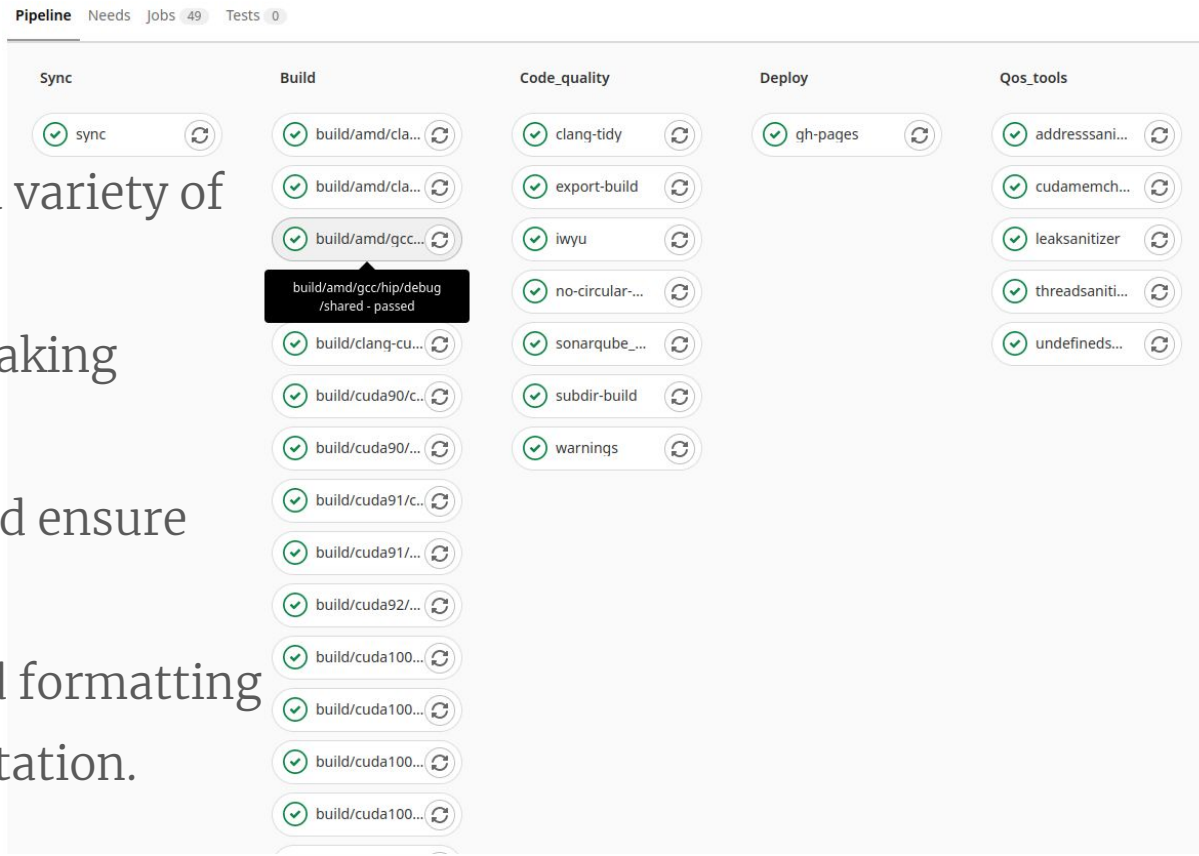

Peer review - Platforms

- Many platforms available for collaborative coding.
- Gitlab, Github, Bitbucket are examples of such platforms.
- They provide web-based interface
  - to interact with code,
  - give feedback on specific lines of code,
  - show pass/failure of the CI pipelines

Pratik Nayak, nayak@kit.edu

Peer review - code review

- Peer review is essential in collaborative software development.
- View code critically in context of the entire project.
- Making sure the code follows the code guidelines.
- Looking for performance and correctness gotchas.

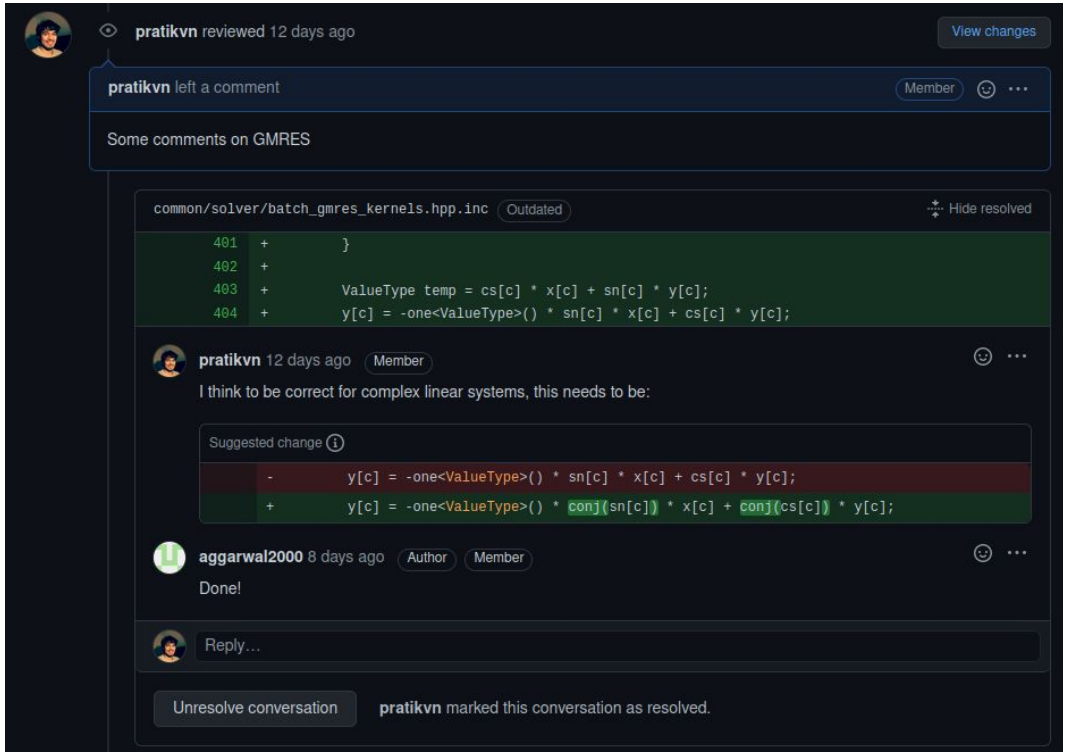

Pratik Nayak, nayak@kit.edu

Where ?

- Numerical Linear Algebra for High Performance Computing course at Karlsruhe Institute of Technology as a pilot programme.
- Course content:
  - Parallel programming basics.
  - OpenMP, CUDA and MPI programming models.
  - BLAS kernels and their implementations.
  - Iterative solvers.
  - Preconditioners.

Target audience: Master students in the Math and Computer Science departments.

Pratik Nayak, nayak@kit.edu

How ?

- Incorporate these practices into the course.
- Encourage students to experiment with algorithms and implementations and not be bogged down in build and platform issues.
- Homework schedule:
  - HW1: Basic Linear Algebra operations.
  - HW2: Dense Matrix-Vector multiplication with OpenMP and CUDA.
  - HW3: LU factorization with OpenMP tasking framework.
  - HW4: Sparse Matrix-Vector multiplication with OpenMP and CUDA.
  - HW5: An Iterative linear system solver with CUDA.

Pratik Nayak, nayak@kit.edu

How ?

Main repository

- Create a common Exercise framework for all to work on.
- Provide the building blocks: Compilation, testing and benchmarking frameworks and setup a Continuous Integration setup to automatically test the code on push.

https://github.com/pratikvn/nla4hpc-exercises-framework

Pratik Nayak, nayak@kit.edu

How ?

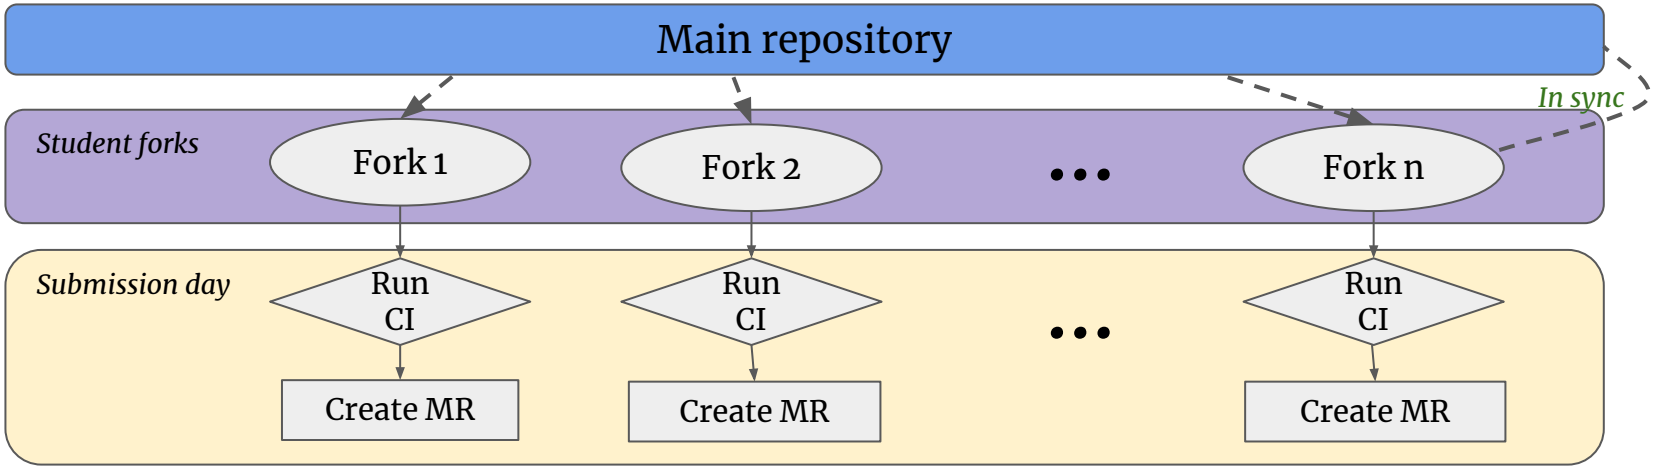

- Students create forks from the main repository, create separate branches for each HW and on submission date submit a Merge Request to merge the changes back to the main repository.

Pratik Nayak, nayak@kit.edu

How ?

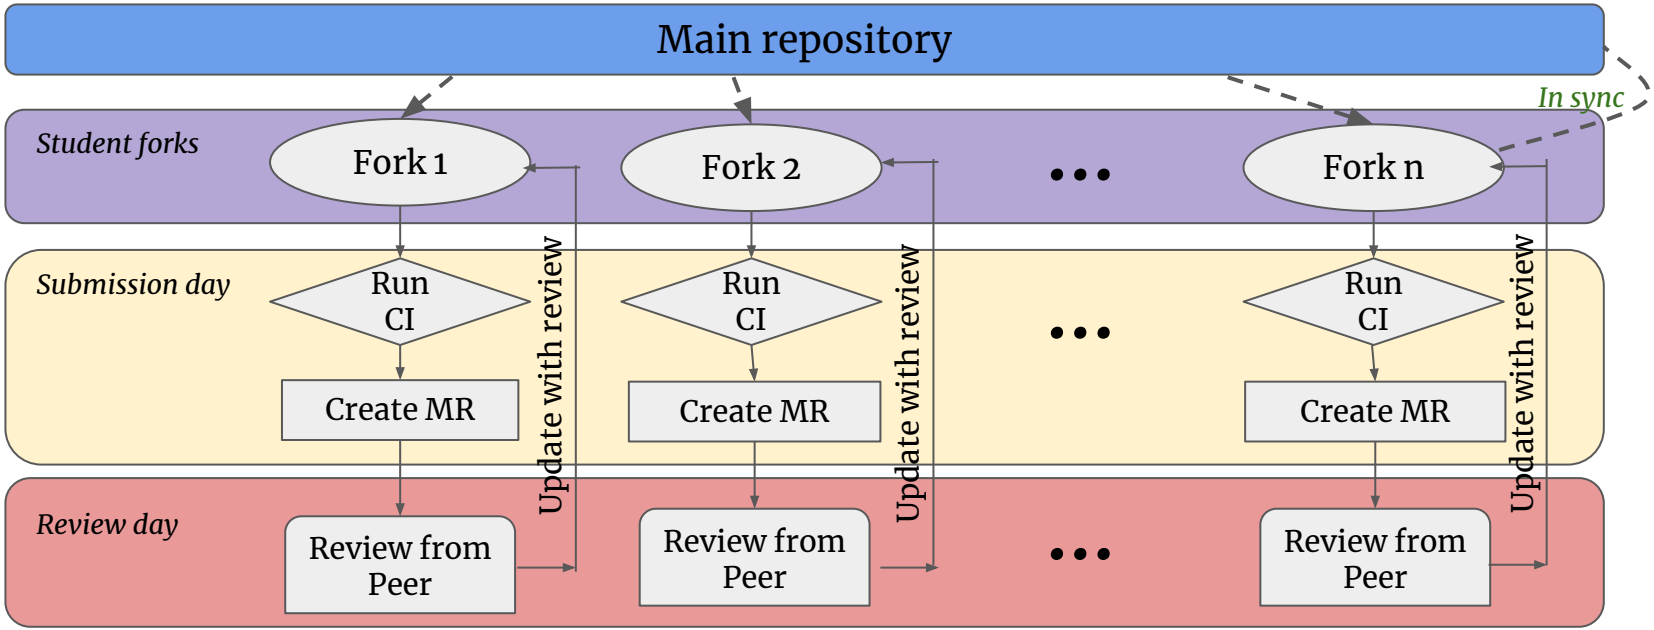

- They are assigned a Merge Request to review and have certain guidelines to follow.
- At the deadline, the Merge Requests are merged back into the main branch and graded.

Pratik Nayak, nayak@kit.edu

How ?

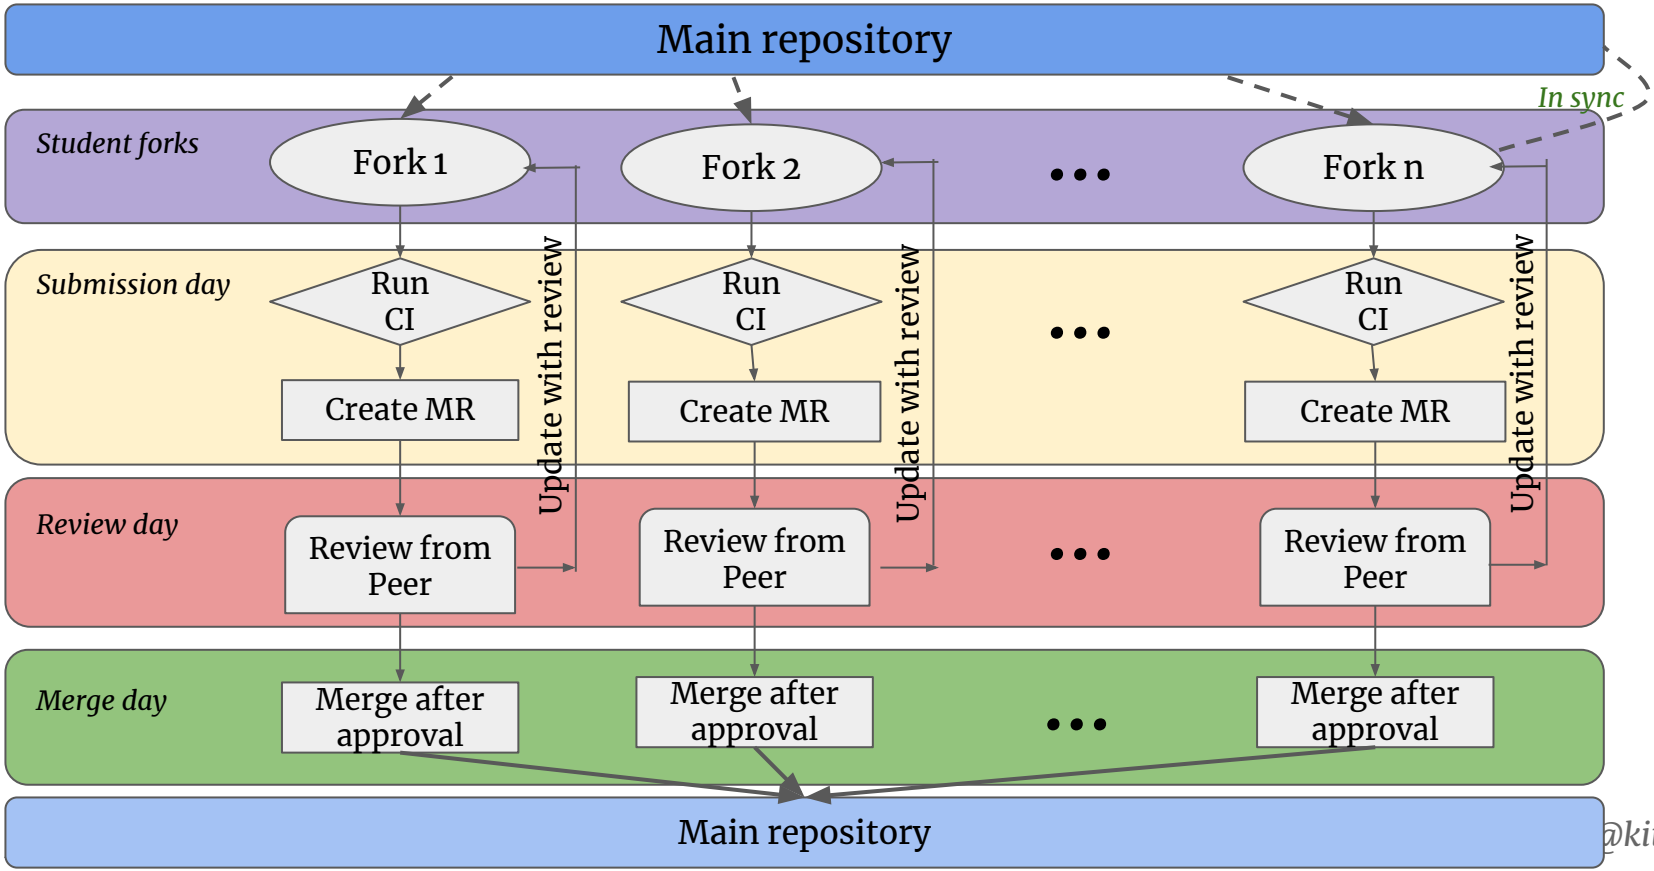
Main repository
In sync
Student forks
Fork 1
Fork 2
Fork n
Submission day
Run CI
Run CI
Run CI
Create MR
Create MR
Create MR
Update with review
Update with review
Review day
Review from Peer
Review from Peer
Review from Peer
Merge day
Merge after approval
Merge after approval
Merge after approval
Main repository

Peer review in action

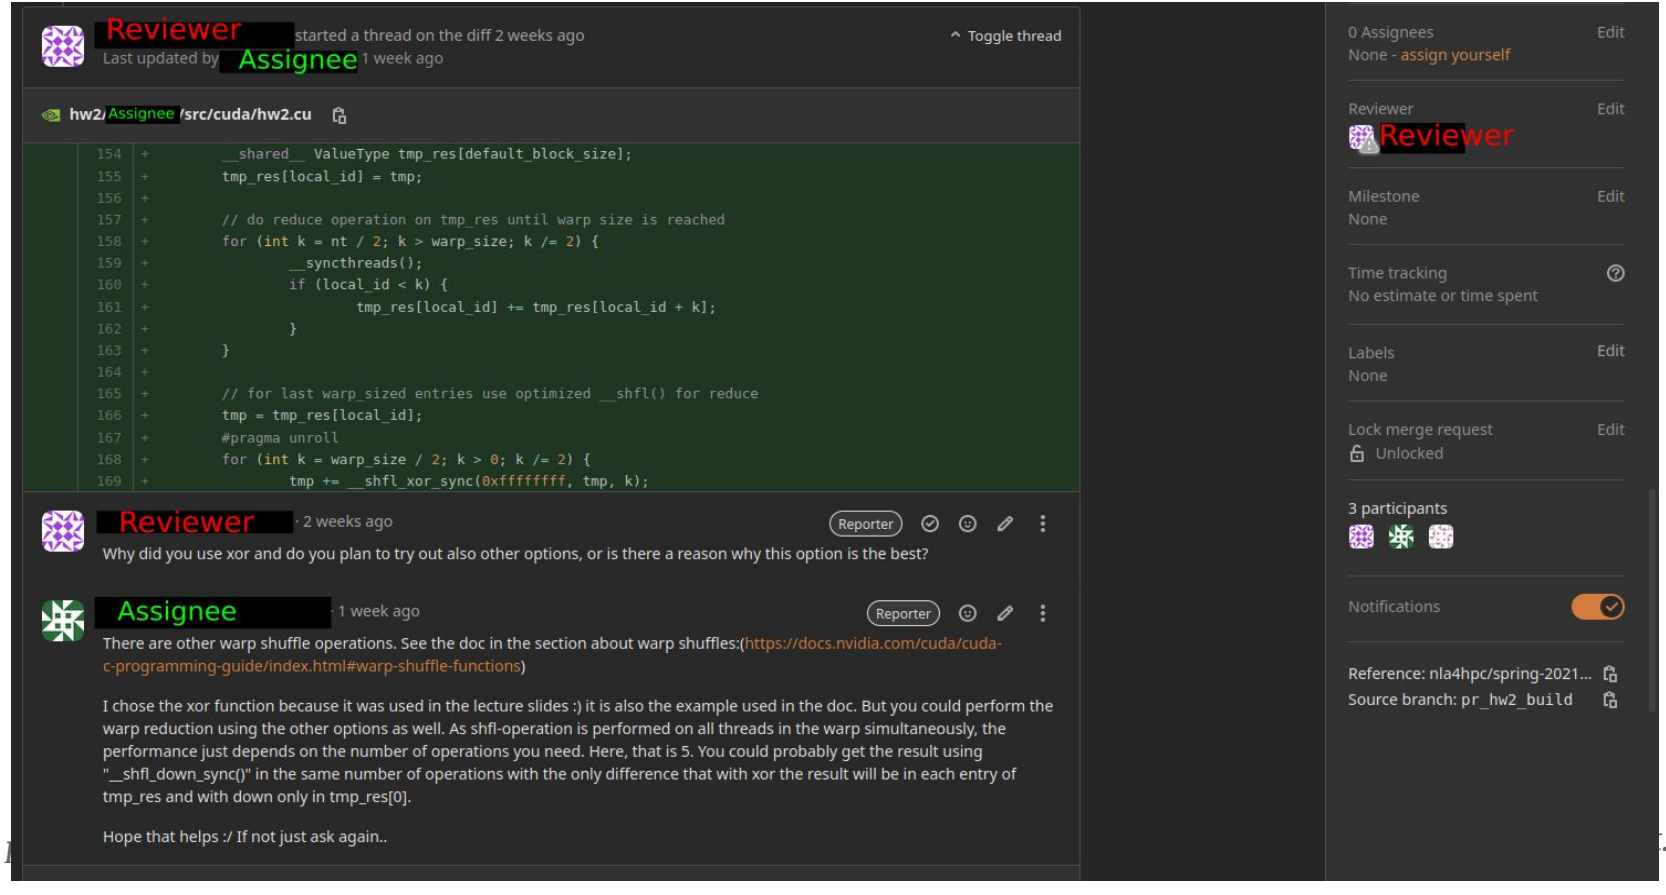

Peer review in action

- CI is run on the forks.
- Discussion on the gitlab web interface.
- Approval from reviewer after addressing the comments.
- Merge after approval.
- Green is assignee
- Red is reviewer.

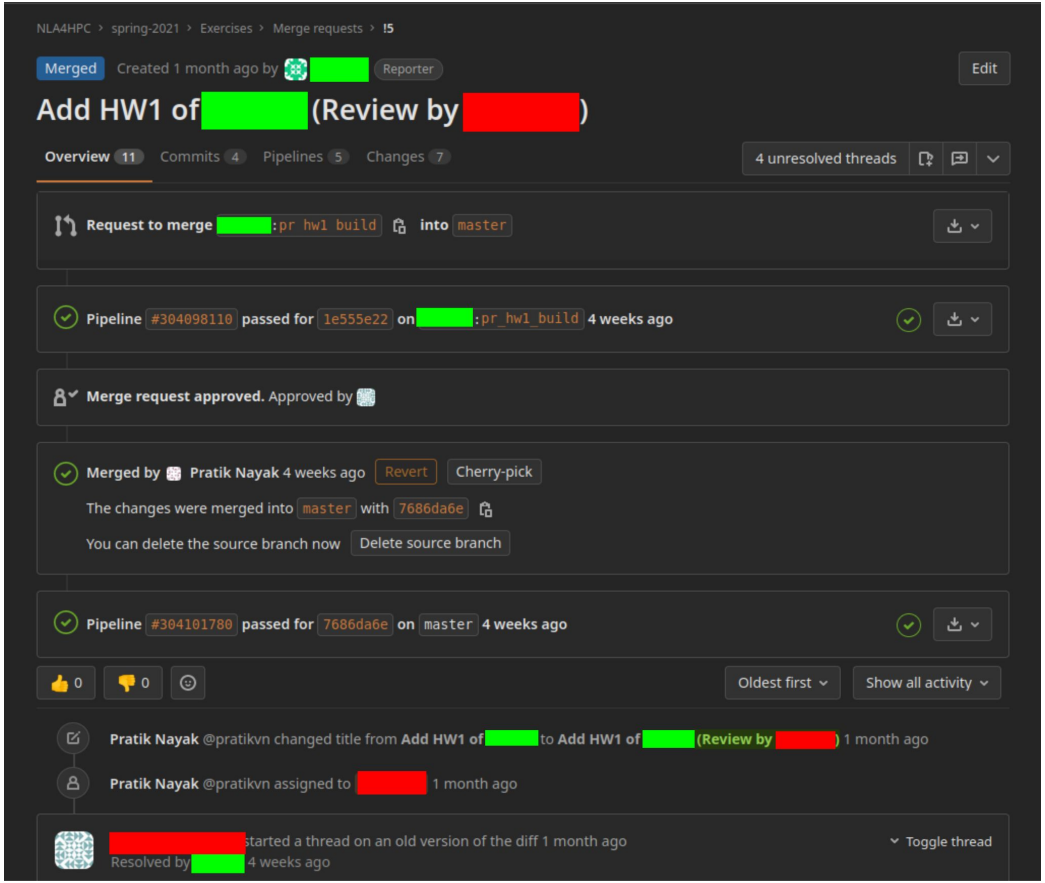

Feedback.

Question	Avg rating (1–5)
How easy was it to use the framework?	2
How useful did you find the exercises instructions?	2
How easy was it to compile and run the code as provided?	2.3
How useful was the code review from your peer?	1.6
How easy was the reviewing process?	3.6
Would you like to see this type of frameworks in other courses?	1

Some feedback from students. (N=4)

Very useful/ Very easy

1 ⟵ ⟶ 5

Not useful/ Not easy

Pratik Nayak, nayak@kit.edu

Conclusions and Perspectives.

- We saw a marked improvement in code quality as the course progressed, which was not the case in our previous course offerings.
- This approach seems more scalable even with a lot more students.
  - It can be almost completely automated.
- The students were able to focus on algorithms and optimizations rather than on build system and other orthogonal issues.
- It encourages students to showcase their code and makes them comfortable with contributing to open-source projects.

Pratik Nayak, nayak@kit.edu

Pratik Nayak, nayak@kit.edu

Thank you!

nayak@kit.edu

Backup

Pratik Nayak, nayak@kit.edu

How ?

- Create a common Exercise framework for all to work on.
- Provide the building blocks: Compilation, testing and benchmarking frameworks and setup a Continuous Integration setup to automatically test the code on push.
- Students create forks from the main repository, create separate branches for each HW and on submission date submit a Merge Request to merge the changes back to the main repository.
- They are assigned a Merge Request to review and have certain guidelines to follow.
- At the deadline, the Merge Requests are merged back into the main branch and graded.
- Students submit an additional report analyzing their code and performance.

Pratik Nayak, nayak@kit.edu

How ?

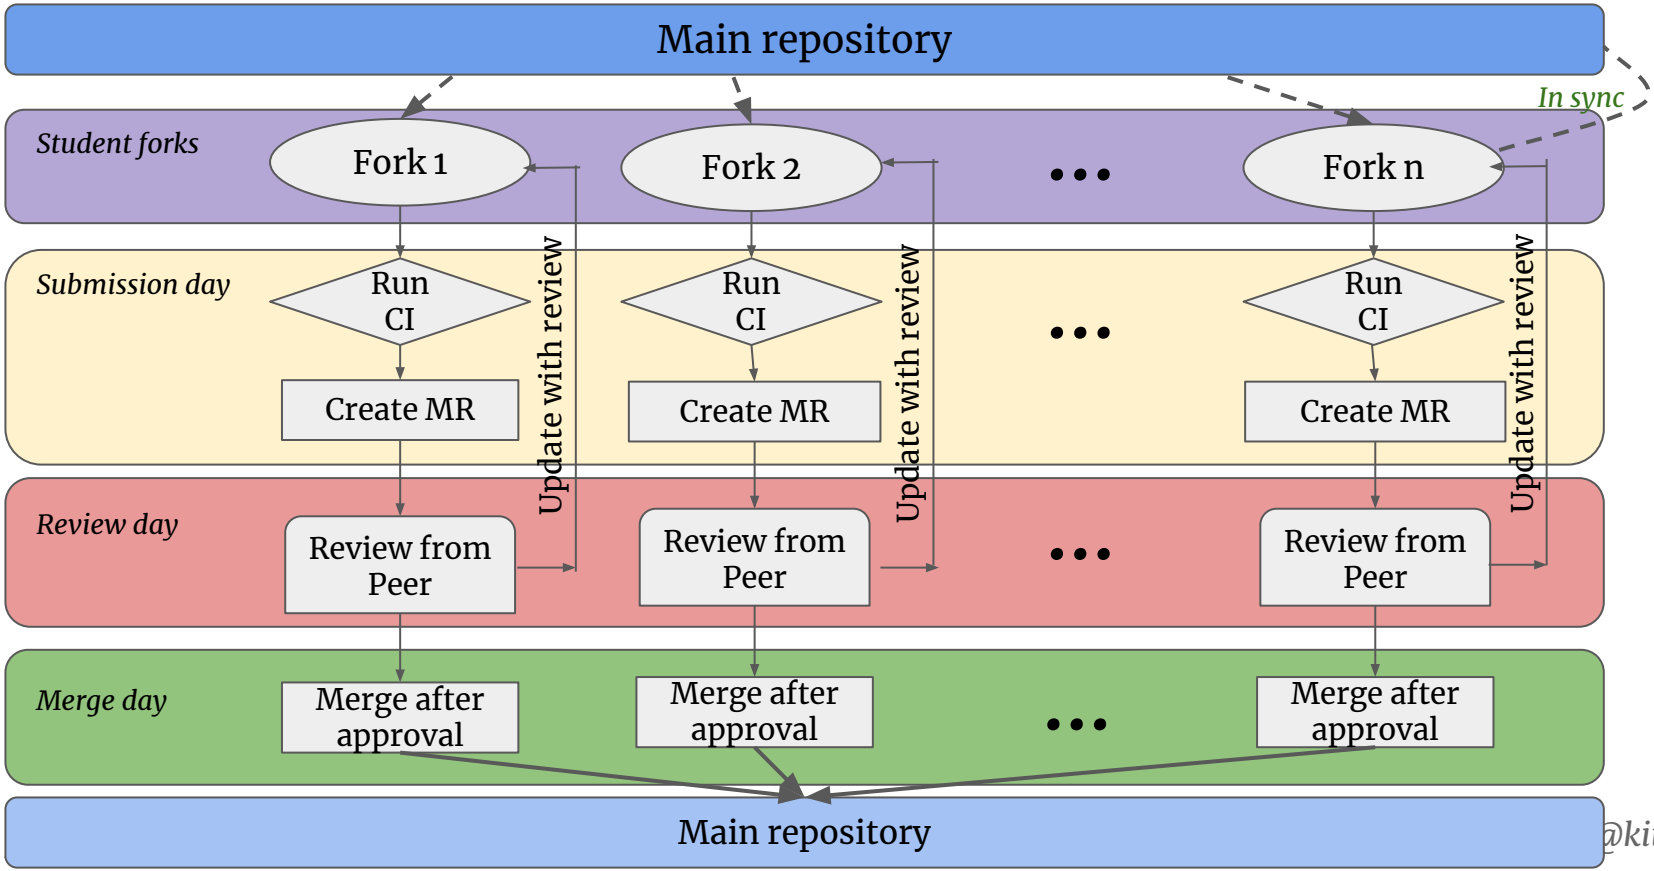
Main repository
In sync
Student forks
Fork 1
Fork 2
Fork n
Submission day
Run CI
Run CI
Run CI
Create MR
Create MR
Create MR
Update with review
Update with review
Review day
Review from Peer
Review from Peer
Review from Peer
Merge day
Merge after approval
Merge after approval
Merge after approval
Main repository