

# Towards a modern CMake workflow

Heinz-Peter Liechtenecker<sup>1</sup> and Raphael Riebl<sup>2</sup>

<sup>1</sup> TU Graz, Institut für Elektronik, Inffeldgasse 12, 8010 Graz, Austria

<sup>2</sup> TH Ingolstadt, Esplanade 10, 85049 Ingolstadt, Germany

## Abstract

Modern CMake offers the features to manage versatile and complex projects with ease. With respect to OMNeT++ projects, a workflow relying on CMake enables projects to combine discrete event simulation and production code in a common development environment. Such a combination means less maintenance effort and thus potentially more sustainable and long-living software. This paper highlights the significant improvements since the first attempt of using CMake in OMNeT++ projects. In particular, a state-of-the-art integration of OMNeT++ in Visual Studio Code including support for debugging and multi-platform compilation is presented. Last but not least, an exemplary use case demonstrates the powerful mix of production and simulation code in a common software architecture supported by the OMNeT++ CMake package.

## 1 Introduction

Modular software architectures are of great help in designing platform-independent and long-living software, which allows developers to reuse components. Furthermore, complex functionality can be broken down into smaller chunks to ease handling and testing. In the context of event-driven simulation, it is desirable to reuse features such as protocol decoders or application logic: Besides reduced development effort, the resulting simulation model closely resembles the application as found in production environments. For example, the INET Framework relies on some FFmpeg libraries for its **Voice over IP (VoIP)** feature instead of reinventing the wheel.

A well-designed software architecture and build environment reduces the management overhead and thus allows even small teams to maintain and improve complex projects. CMake is a battle-proven tool supporting developers to create such a powerful build environment. With its recent advances modern CMake [1] can, among others, be employed for the management of various build targets, automation of pre- and post-build tasks, and cross-platform toolchain orchestration. Furthermore, CMake is capable of generating various build contexts, ranging from classic Makefiles to the younger Ninja [2] build system.

Usage of CMake as an option to handle software dependencies in OMNeT++ simulations has first been presented at the OMNeT++ Community Summit in 2015 [3]. Maintenance of the presented tooling has not ceased since then; it took place within the Artery project, however. This intermingling with Artery made it unnecessarily hard for other projects to adopt CMake as well. Meanwhile, Thor K. Høgås extracted the related sources into a separate repository in 2020. Since 2021, this repository has moved under the umbrella of the OMNeT++ organisation on GitHub<sup>1</sup>. Thor and the authors of this paper are now jointly maintaining this standalone OMNeT++ CMake package for (hopefully) wide usage by the OMNeT++ community.

---

<sup>1</sup><https://github.com/omnetpp/cmake>

Beyond these organisational changes, this OMNeT++ CMake package has also been drastically improved in terms of features. This paper highlights some of these improvements: Section 2 revisits how to deal with OMNeT++ projects in CMake and how the related mechanisms have changed. Furthermore, we outline in Section 3 how OMNeT++ projects can be developed with alternative **Integrated Development Environments (IDEs)** such as **Visual Studio Code (VSC)**. In particular, we bring convenient features such as OMNeT++ debug targets to **VSC**. Section 4 shows an exemplary project structure, which leverages CMake to have build targets for production and simulation software in a common setup. Such a setup makes it easy to run continuous integration tests in a simulation environment to strengthen confidence in the software executed in a production environment.

## 2 Depending on OMNeT++ Projects

One of CMake’s key selling points is its ability to handle third-party software: It is a piece of cake for developers to add such dependencies to their projects – if someone has extended CMake’s *find* capabilities accordingly. With the help of our OMNeT++ CMake package, `find_package(OmnetPP)` is readily available which provides *import targets* for all OMNeT++ libraries and *variables* to OMNeT++ tools such as the message compiler.

If a project depends not only on OMNeT++ itself but also intends to reuse models by OMNeT++ frameworks such as INET or SimuLTE, these frameworks need to be covered by CMake as well. In the OMNeT++ CMake package’s first incarnation, a project had to use the `import_opp_target` macro to add another OMNeT++ project as its dependency. As detailed in [3], the *Makefile* produced by `opp_makemake` is parsed by our `opp_cmake` script generating an adequate CMake *import target* out of it. However, the original build process of the OMNeT++ project was still in charge of creating the project’s binaries via its *Makefile*.

Nowadays, `add_opp_target` provides an alternative mechanism: Instead of merely importing binaries built by an external toolchain, `add_opp_target` lets CMake itself build the OMNeT++ project’s binaries. It can not only handle the C++ source files (`*.cc`) but also OMNeT++ messages (`*.msg`), *i.e.* it invokes the OMNeT++ message compiler as an intermediate build step. Similar to its `import_opp_target` sibling, the resulting target has its `NED_FOLDERS` property configured. Hence, the project’s NED paths are automatically considered when a depending simulation model is run.

The main advantage of `add_opp_target` over `import_opp_target` is the consistent build chain. A CMake project employing `import_opp_target` has to trigger the execution of the OMNeT++ *Makefile*, even if the CMake project itself employs another build generator such as Ninja. With `add_opp_target`, all build steps are fully controlled by CMake itself. Hence, one can easily exploit the speed-up enabled by parallelised Ninja builds. Furthermore, the compiler options used for OMNeT++ build artefacts are entirely configurable via the common CMake mechanisms.

## 3 Alternative IDEs

The official OMNeT++ **IDE** is the natural choice for many OMNeT++ users as it neatly integrated OMNeT++ specifics. Especially novice users thus do not have to worry about

peculiarities such as setting up NED paths. As long as the simulation models remain rather self-contained, i.e. no dependencies to third-party components exist, the workflow in the OMNeT++ IDE is rather straightforward.

As soon as users head for more sophisticated setups, our OMNeT++ CMake package can help out to tackle the management of dependencies as pointed out in Section 2. Unfortunately, the built-in OMNeT++ IDE plugins cannot handle such a CMake environment as conveniently as the classic *opp\_makemake* projects. However, a CMake-based OMNeT++ project can be handled by any IDE with decent CMake support. In the following, we explore the state-of-the-art in developing OMNeT++ projects with VSC [4].

A modern and robust IDE has to fulfil several requirements, including code highlighting and syntax checks, auto-completion and suggestions, integration with version control systems, execution of builds, debugging and execution and analysis of tests. In recent years, VSC has become a popular choice for a plethora of use cases. Customisability and cross-platform support are key features of VSC, which allow us to neatly adopt it for OMNeT++ development as well. Especially with its extensions to build and debug C/C++ code managed by CMake, VSC is an ideal candidate to work with the OMNeT++/CMake ecosystem. The authors suggest the following extensions making VSC a fully-featured IDE:

- *ms-vscode.cpptools* for C/C++ language support [5]
- *twxs.cmake* for CMake language support [6]
- *ms-vscode.cmake-tools* for CMake project integration [7]
- *schrej.omnetpp-ned* for OMNeT++ NED language support [8]
- *vadimcn.vscode-lldb* (CodeLLDB) for debugging C++ applications [9]

### 3.1 Debugging OMNeT++ Models

Besides code editing, debugging is a major requirement within any IDE. For launching an actual debug session, VSC depends on additional instructions given in the VSC-specific *launch.json* file [10]. In particular, debugging an OMNeT++ model requires instructions about the desired debugger, the location of the OMNeT++ runner executable (usually *opp\_run\_debug*), the OMNeT++ model and associated command-line parameters. These command-line parameters specify the look-up paths for NED files, the to be used simulation configuration file (an *omnetpp.ini* in most cases), and the OMNeT++ model's libraries. Managing these parameters manually is a tedious and error-prone task, especially for large and complex projects. For example, Artery configurations often employ multiple other OMNeT++ models, e.g. SimuLTE and INET, plus its own modules.

As a remedy for this maintenance burden, recent versions of the OMNeT++ CMake package can configure the VSC debug targets automatically. The build system then updates the debug configuration in the *launch.json* file according to the CMake project configuration. If this feature is enabled, VSC debug configurations are generated for every *add\_opp\_run* invocation. As before, *add\_opp\_run(<name>)* still adds *run\_<name>* (always), *debug\_<name>* (for debug builds), and *memcheck\_<name>* (if Valgrind is installed) targets. Now, it also invokes a Python script shipped with the OMNeT++ CMake package that updates the *launch.json* in the project's *.vscode* directory. Manual configurations by the

user remain untouched, though. Hence, it is perfectly fine to have generated and custom debug configuration side-by-side as long as their names are distinct.

At the moment, two flavours of debug configurations are supported: GDB [11] and LLDB [12], the debuggers by the GNU and LLVM projects, respectively. Listing 1 shows an exemplary configuration for the VSC CodeLLDB extension employing the LLDB debugger. We have found this extension to be advantageous for debugging across platforms, *i.e.* on Windows and Linux systems. The upcoming OMNeT++ 6.0 release will also ship with an OMNeT++ formatter script for the LLDB debugger. For these combinations, *i.e.* CodeLLDB and an OMNeT++ 6.0 pre-release, the generated debug configurations load this formatter script as part of the debugger's initialisation commands. Thus, the debugger pretty-prints the OMNeT++ data structures out-of-the-box then.

Listing 1: VSC debug configuration using CodeLLDB (*launch.json*)

---

```

1 {
2   "version": "0.2.0",
3   "configurations": [
4     {
5       "name": "Launch MyProject - CodeLLDB (OMNeT++)",
6       "type": "lldb",
7       "request": "launch",
8       "program": "path/to/bin/opp_run_dbg",
9       "args": [
10        "-n", "path/to/ned/folders",
11        "-l", "path/to/library",
12        "path/to/omnetpp.ini"
13      ],
14      "stopOnEntry": false,
15      "cwd": "path/to/working/directory",
16      "initCommands": [
17        "command script import path/to/lldb/formatters/omnetpp.py"
18      ]
19    }
20  ]
21 }
```

---

Unfortunately, the Python's native **JavaScript Object Notation (JSON)** package does not support C-style comments within **JSON** files and is not as fault-tolerant as the VSC's **JSON** decoder. Hence, when manually editing the *launch.json* file, the user must ensure to comply strictly with the **JSON** standards for now. We may be able to lift this limitation in the future.

### 3.2 Improving Developer Experience on Windows with VSC

VSC and the corresponding CMake extension allows to set up a modern and flexible cross-platform development environment. Especially developing OMNeT++ packages and simulations on Windows using alternatives to the shipped Eclipse-based IDE can often be error-prone as OMNeT++ comes with its own C++ toolchain environment. Integrating the virtual environment in alternative IDEs like VSC in conjunction with the OMNeT++ CMake

can be achieved quickly - for **VSC**, the CMake extension needs to be aware of the shipped OMNeT++ compilers. In addition, the virtual environment binary paths have to be part of the systems PATH environment variable. Paths can be either added globally or by executing an environment setup script. The **VSC** CMake extension offers the necessary feature: Utilising user local kits, one can add additional toolchains within a *cmake-kits.json* file [13]. Listing 2 gives an example of a *cmake-kits.json* file setting up a kit for OMNeT++ 6.0pre10 using the Clang compiler and executing a command script to set up the environment. The environment setup script, shown in Listing 3, sets the PATH variable and optionally activates a virtual python environment.

Listing 2: Example of a **VSC** CMake Kits configuration (*cmake-kits.json*)

---

```

1 [
2   {
3     "name": "Clang OMNeT++ 6.0pre10 with Python VENV",
4     "environmentSetupScript": "${workspaceFolder}/.vscode/omnetpp-6.0
      pre10env.cmd",
5     "compilers": {
6       "C": "path/to/omnetpp-6.0pre10/tools/win64/mingw64/bin/clang.exe",
7       "CXX": "path/to/omnetpp-6.0pre10/tools/win64/mingw64/bin/clang++.exe"
8     }
9   }
10 ]
```

---

Listing 3: Environment setup script (*omnetpp-6.0pre10env.cmd*)

---

```

1 set PATH=%PATH%;path\to\omnetpp-6.0pre10\tools\win64\mingw64\bin
2 set PATH=%PATH%;path\to\omnetpp-6.0pre10\bin
3 set PATH=%PATH%;path\to\omnetpp-6.0pre10\tools\win64\opt\mingw64\bin
4 set PATH=%PATH%;path\to\omnetpp-6.0pre10\tools\win64\usr\bin
5 set PATH=%PATH%;path\to\omnetpp-6.0pre10\lib
6
7 rem Optional: Activate a python virtual environment
8 set current_dir="%~dp0"
9 call %current_dir%..\venv\Scripts\activate.bat
```

---

Taking full advantage of the presented features, developing and testing simulations and packages against different OMNeT++ versions becomes possible by simply switching the CMake kit within **VSC**. Figure 1 gives an overview of the **IDE** experience: Besides building binaries, CMake also configures the **IDE** itself based on the given build-targets.

## 4 Use Case: Build Targets and Continuous Integration

Various factors lead to the necessity to reuse portions of source code within simulation and production environments. Especially for small teams, reuse is a vital aspect so projects can be maintained for years. With suitable software and build architecture in place, improvements found by simulation can be directly integrated into production code. Reducing

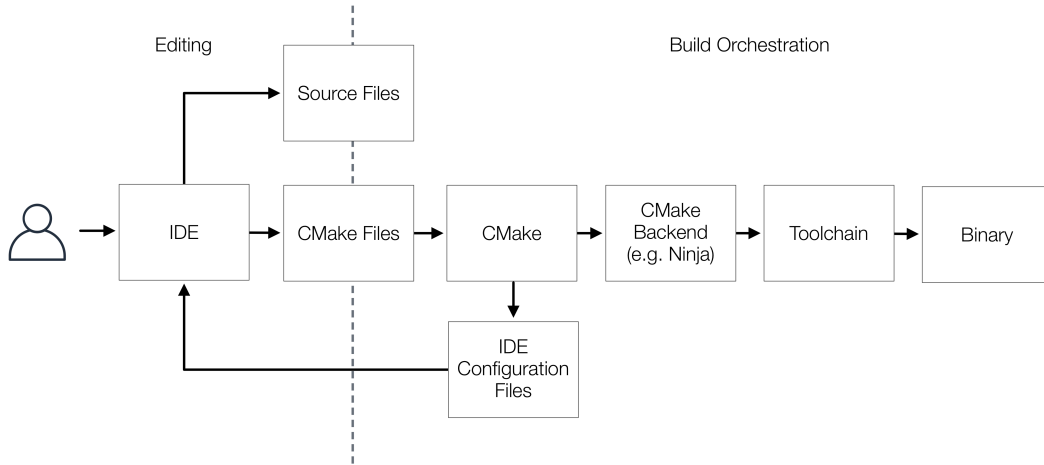


Figure 1: CMake Toolchain and IDE integration

redundancy between simulation and production code is a critical factor for development efficiency. In conjunction with a suitable build management tool such as CMake, build targets tailored to production can easily integrate existing source code. Typical build targets may comprise:

- Dynamically loaded libraries to run a simulation
- Test executables putting various code sections through their paces
- A standalone executable reusing parts of the code used for simulation within a production environment

Figure 2 shows an exemplary project structure featuring a *simulation library* with OMNeT++ modules, a *standalone executable* for productive operation, and *tests*. In the given example, a network application shall receive TCP traffic, process it and respond to it if necessary. The *application library* includes this core functionality, which is employed by all three types of build targets as mentioned above. The *platform abstraction layer* provides the networking functionality required by the application library via a platform-agnostic interface: From the application's perspective, it does not matter if the network is simulated or an actual network is accessed. The respective implementations of this interface can rely on OMNeT++/INET for a simulated runtime context or the ASIO library, which employs the operating system's network stack. ASIO [14] and INET [15], in particular its socket implementation, behave both like event-driven I/O systems. With this common behaviour of INET sockets and the ASIO library, a generalised interface can easily be offered to the application library.

Two variants are apparent to realise the platform abstraction layer within CMake: Either distinct (static or dynamic) library targets exist for each implementation, or all flavours are built into a single binary, e.g. using the pointer-to-implementation idiom at runtime [16]. With the former approach, each library depends only on some platform-specific libraries. For example, the "real" implementation depends on the ASIO library, while the "simulation"




Figure 2: CMake targets and their dependencies in a mixed simulation & production project

implementation depends only on OMNeT++ and INET. In any case, the platform variants are realised as CMake targets, and thus CMake takes care of all the (transitive) dependencies automatically.

Without changes at the application logic itself, the executable for productive operation can employ the ASIO-based platform while other components may use the OMNeT++-based platform. Results and improvements found by simulation and verification are directly contributed to a common code base represented by the application library. This pattern predestines for a continuous integration pipeline. Ideally, the build system is part of this pipeline to build code upon specific events, e.g. a Git commit or pull request. By dividing the application into several targets, it is possible only to build those of interest: the application library with the corresponding platform abstraction implementation, the standalone and the testing executables. To sum up, the proposed CMake workflow avoids the management overhead occurring when separate build systems need to be maintained for each platform.

## 5 Conclusion

Due to the presented advances in the OMNeT++ CMake package, developers familiar with CMake workflows can fully integrate existing OMNeT++ packages like INET into their projects. Extending the functionality of the CMake package to automate the setup of **VSC** as a fully-featured integrated development environment improves developer experience on Linux- and Windows-based systems.

Further work on the package includes optimisation of the environment setup when debugging simulations within **VSC** on Windows. The approach presented in Section 3.2 currently only sets the environment variables for the build steps; hence debugging requires adding OMNeT++ paths to the global Windows system *PATH* variable. Furthermore, adding examples to show the OMNeT++ CMake package usage, a project skeleton, and improved documentation is necessary to lower the entry bar for developers heading towards a modern CMake workflow with OMNeT++.

## References

- [1] Kitware Inc. *et al.*, “CMake.” [online]. <https://cmake.org/documentation/>, accessed on 7th July 2021.
- [2] E. Martin *et al.*, “Ninja build tool.” [online]. <https://ninja-build.org/manual.html>, accessed on 7th July 2021.
- [3] R. Riebl and C. Facchi, “Regain control of growing dependencies in OMNeT++ simulations,” *CoRR*, vol. abs/1509.03561, 2015.
- [4] Microsoft Corp. *et al.*, “Visual Studio Code.” [online]. <https://code.visualstudio.com/>, accessed on 7th July 2021.
- [5] Microsoft Corp. *et al.*, “C/C++ for Visual Studio Code.” [online]. <https://github.com/microsoft/vscode-cpptools/>, accessed on 7th July 2021.
- [6] “CMake language support for Visual Studio Code.” [online]. <https://github.com/twxs/vs.language.cmake>, accessed on 7th July 2021.
- [7] Microsoft Corp. *et al.*, “CMake tools for Visual Studio Code.” [online]. <https://github.com/microsoft/vscode-cmake-tools>, accessed on 7th July 2021.
- [8] J. Schrettenbrunner *et al.*, “OMNeT++ NED language support for Visual Studio Code.” [online]. <https://github.com/schrej/vscode-omnetpp-ned>, accessed on 7th July 2021.
- [9] V. Chugunov *et al.*, “CodeLLDB - LLDB debugger support for Visual Studio Code.” [online]. <https://github.com/vadimcn/vscode-lldb>, accessed on 7th July 2021.
- [10] Microsoft Corp. *et al.*, “Debugging in Visual Studio Code.” [online]. <https://code.visualstudio.com/docs/editor/debugging>, accessed on 7th July 2021.
- [11] Free Software Foundation Inc. and the GDB developers, “GDB - the gnu project debugger.” [online]. <https://www.gnu.org/software/gdb/>, accessed on 19th July 2021.
- [12] The LLDB Team, “GDB - the gnu project debugger.” [online]. <https://lldb.llvm.org/>, accessed on 19th July 2021.
- [13] Microsoft Corp. *et al.*, “CMake kits-management in Visual Studio Code.” [online]. <https://github.com/microsoft/vscode-cmake-tools/blob/develop/docs/kits.md>, accessed on 7th July 2021.
- [14] C. Kohlhoff *et al.*, “ASIO - C++ library for network and low-level I/O programming.” [online]. <https://think-async.com>, accessed on 20th July 2021.
- [15] A. Varga *et al.*, “INET framework.” [online]. <https://inet.omnetpp.org/>, accessed on 20th July 2021.
- [16] S. Meyers, *Effective Modern C++ - 42 Specific Ways to Improve Your Use of C++11 and C++14*. O’Reilly Media, 2014.