



University  
of Victoria

Graduate Studies

Notice of the Final Oral Examination  
for the Degree of Master of Science

of

**EVAN WILDE**

BSc (University of Victoria, 2016)

**“Merge-Trees: Visualizing the Integration of Commits into Linux”**

Department of Computer Science

Friday, August 31, 2018

2:30 P.M.

Engineering and Computer Science Building  
Room 468

Supervisory Committee:

Dr. Daniel German, Department of Computer Science, University of Victoria (Supervisor)  
Dr. Margaret-Anne Storey, Department of Computer Science, UVic (Member)

External Examiner:

Dr. Bram Adams, Computer and Software Engineering, Polytechnique Montreal

Chair of Oral Examination:

Dr. Ralf St. Claire, Department of Curriculum and Instruction, UVic

Dr. David Capson, Dean, Faculty of Graduate Studies

## **Abstract**

Version control systems are an asset to software development, enabling developers to keep snapshots of the code as they work. Stored in the version control system is the entire history of the software project, rich in information about who is contributing to the project, when contributions are made, and to what part of the project they are being made. Presented in the right way, this information can be made invaluable in helping software developers continue the development of the project, and maintainers to understand how the changes to the current version can be applied to older versions of projects.

Maintainers are unable to effectively use the information stored within a software repository to assist with the maintenance of older versions of that software in highly-collaborative projects. The Linux kernel repository is an example of such a project. This thesis focuses on improving visualizations of the Linux kernel repository, developing new visualizations that help answer questions about how commits are integrated into the project. Older versions of the kernel are used in a variety of systems where it is impractical to update to the current version of the kernel. Some of these applications include the controllers for spacecrafts, the core of mobile phones, the operating system driving internet routers, and as Internet-Of-Things (IOT) device firmware. As vulnerabilities are discovered in the kernel, they are patched in the current version. To ensure that older versions are also protected against the vulnerabilities, the patches applied to the current version of the kernel must be applied back to the older version. To do this, maintainers must be able to understand how the patch that fixed the vulnerability was integrated into the kernel so that they may apply it to the old version as well.

This thesis makes four contributions: (1) a new tree-based model, the Merge-Tree, that abstracts the commits in the repository, (2) three visualizations that use this model, (3) a tool called *Linvis* that uses these visualizations, (4) a user study that evaluates whether the tool is effective in helping users answer questions related to how commits are integrated about the Linux repository.

The first contribution includes the new tree-based model, the algorithm that constructs the trees from the repository, and the evaluation of the results of the algorithm. The second contribution demonstrates some of the potential visualizations of the repository that are made possible by the model, and how these visualizations can be used depending on the structure of the tree. The third contribution is an application that applies the visualizations to the Linux kernel repository.

The tool was able to help the participants of the study with understanding how commits were integrated into the Linux kernel repository. Additionally, the participants were able to summarize information about merges, including who made the most contributions, which files were altered the most, more quickly and accurately than with Gitk and the command line tools.