

Studying Bloated Dependencies in the Maven Ecosystem

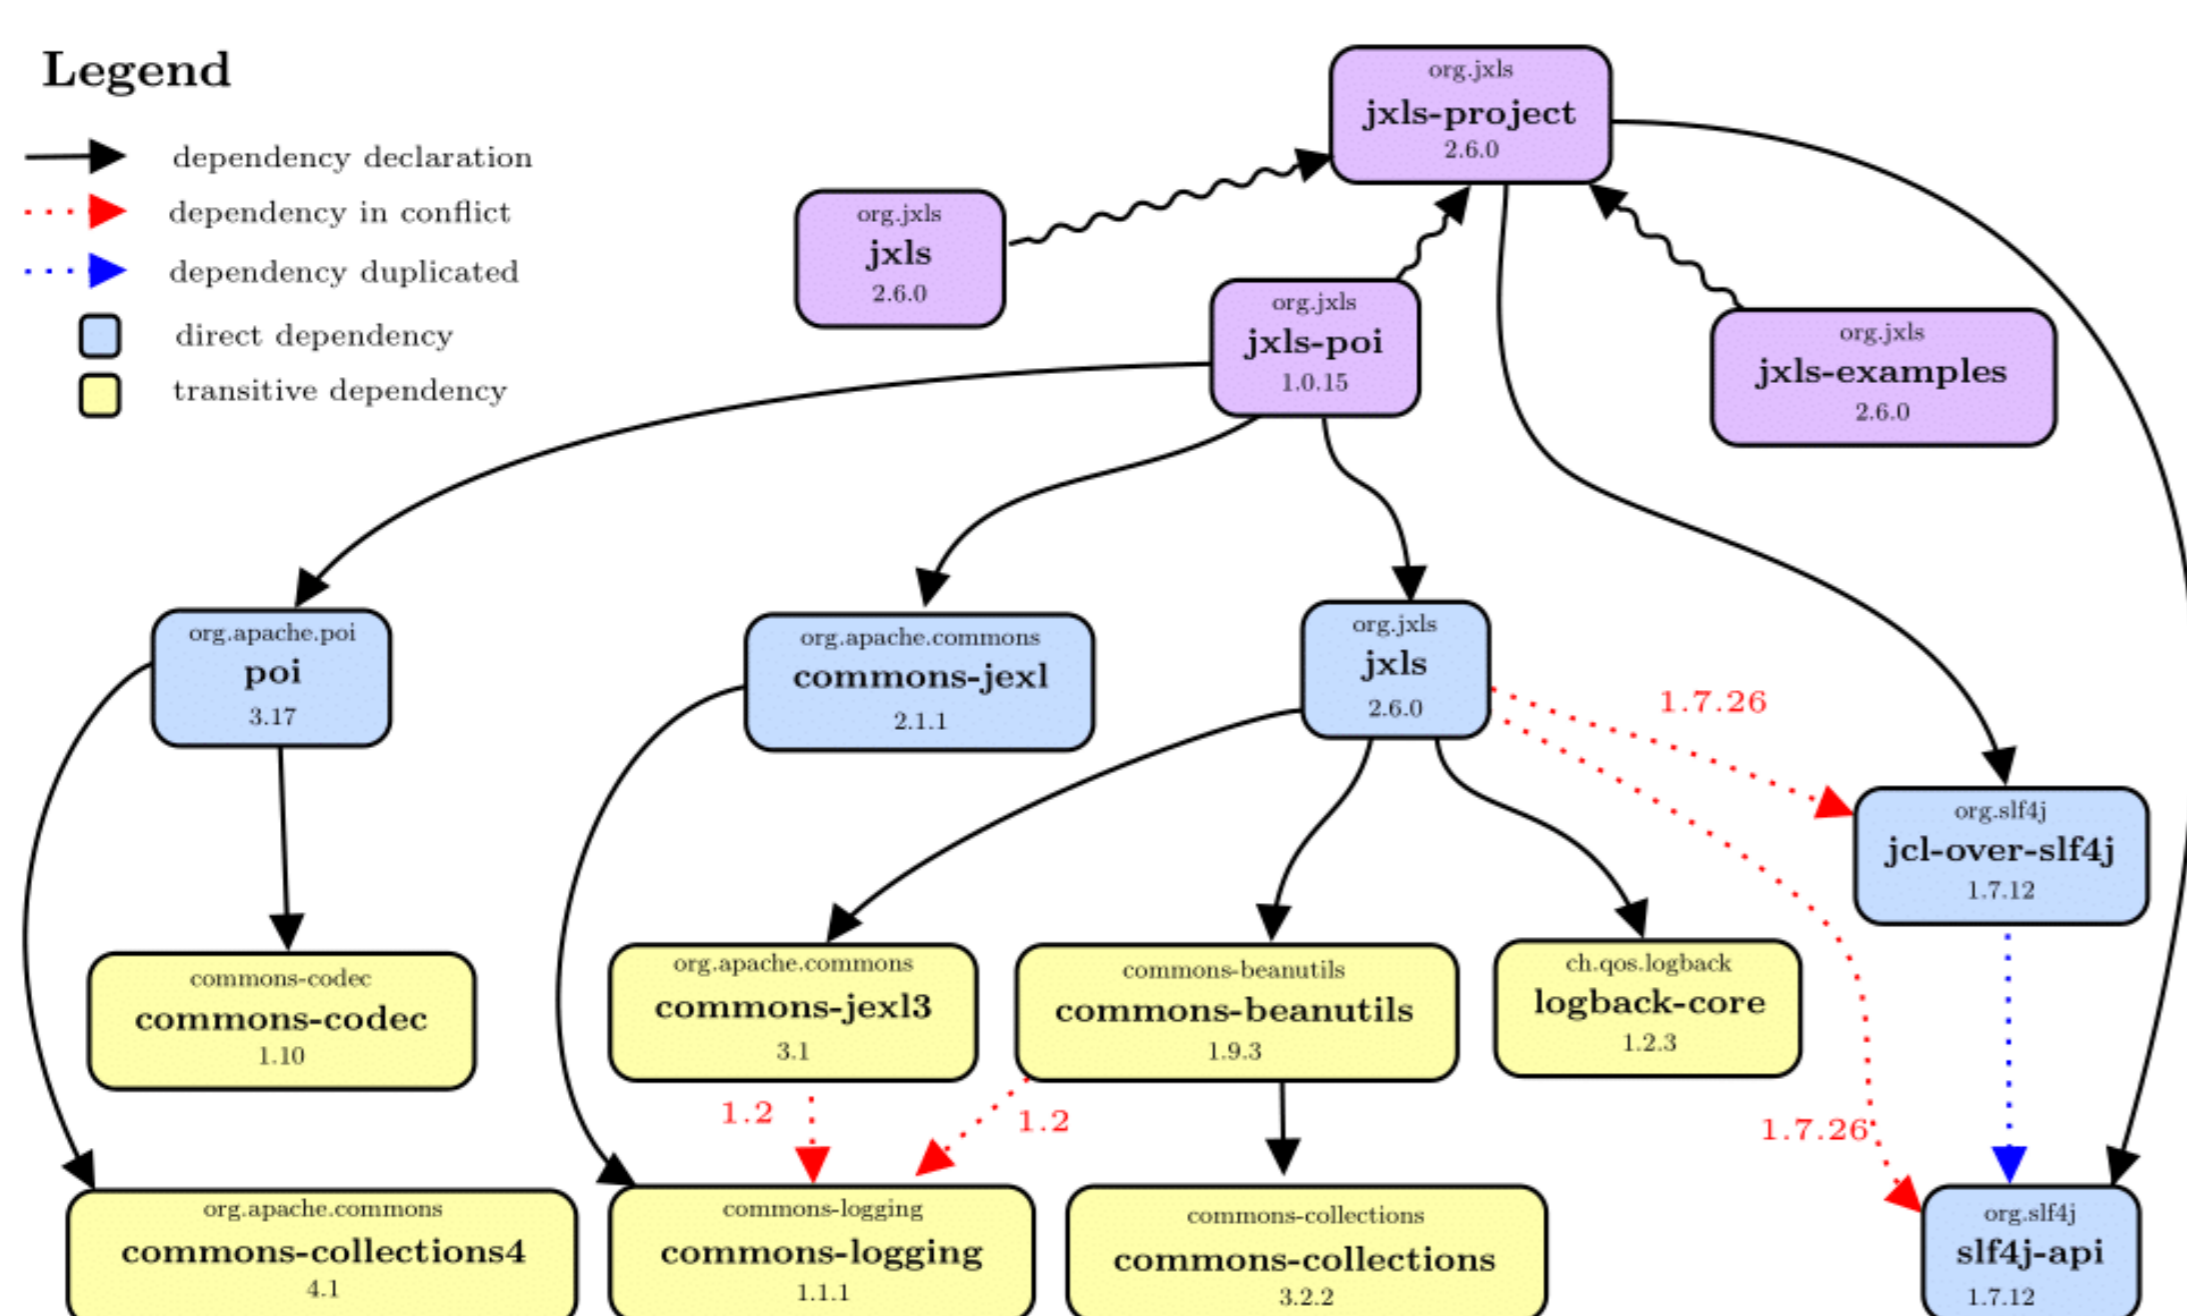
César Soto-Valero [cesarsv@kth.se], supervised by Martin Monperrus and Benoit Baudry
KTH Royal Institute of Technology



Bloated dependencies are added to the project's dependency tree, yet no single method of its API is being used. Consequently, removing such unnecessary dependencies does not affect the project build. We performed a novel, large-scale study of bloated dependencies in the Maven ecosystem. Our study reveals the existence of a pervasive number of bloated dependencies in the Maven Central repository.

Motivation

- Developers declare dependencies that they do not use in their projects



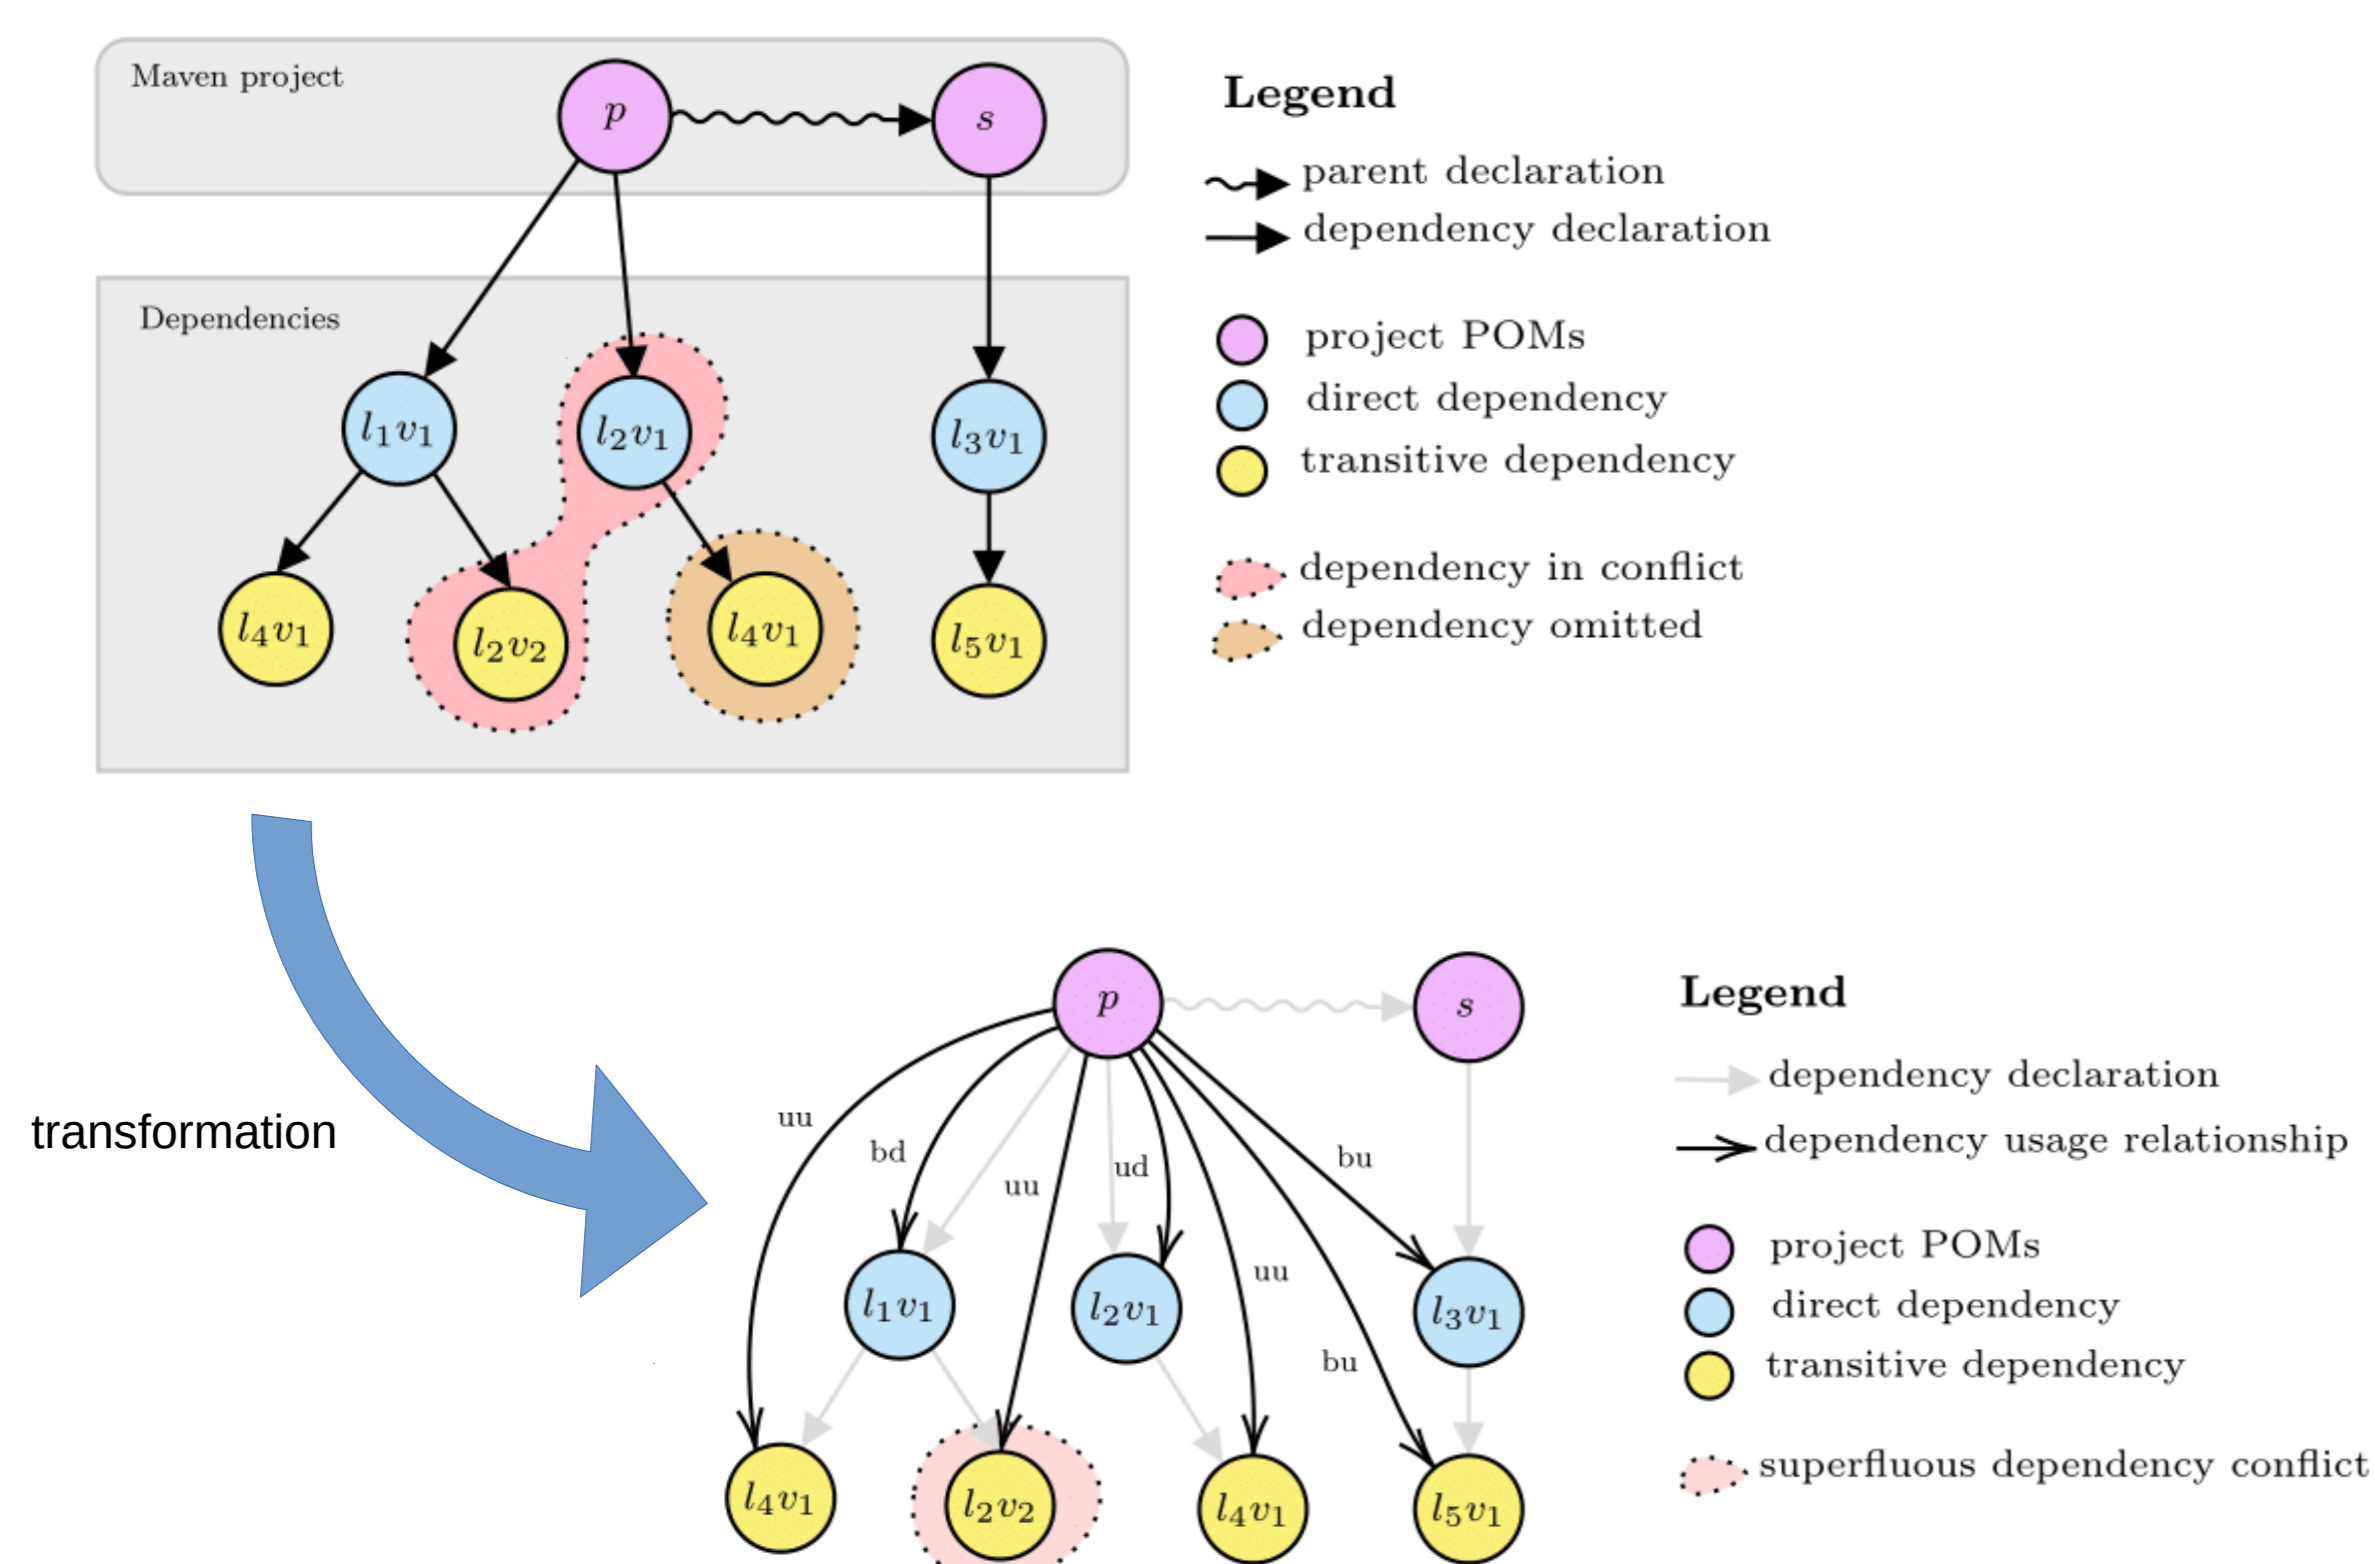
Research Questions

The presence of bloated dependencies in Maven projects is an intriguing phenomenon:

- How frequent is the occurrence of bloated dependencies in the Maven ecosystem?
- Which are the causes of the existence of bloated dependencies?
- To what extent dependency conflicts are superfluous due to bloated dependencies?
- How much can developers benefit from removing their bloated dependencies?

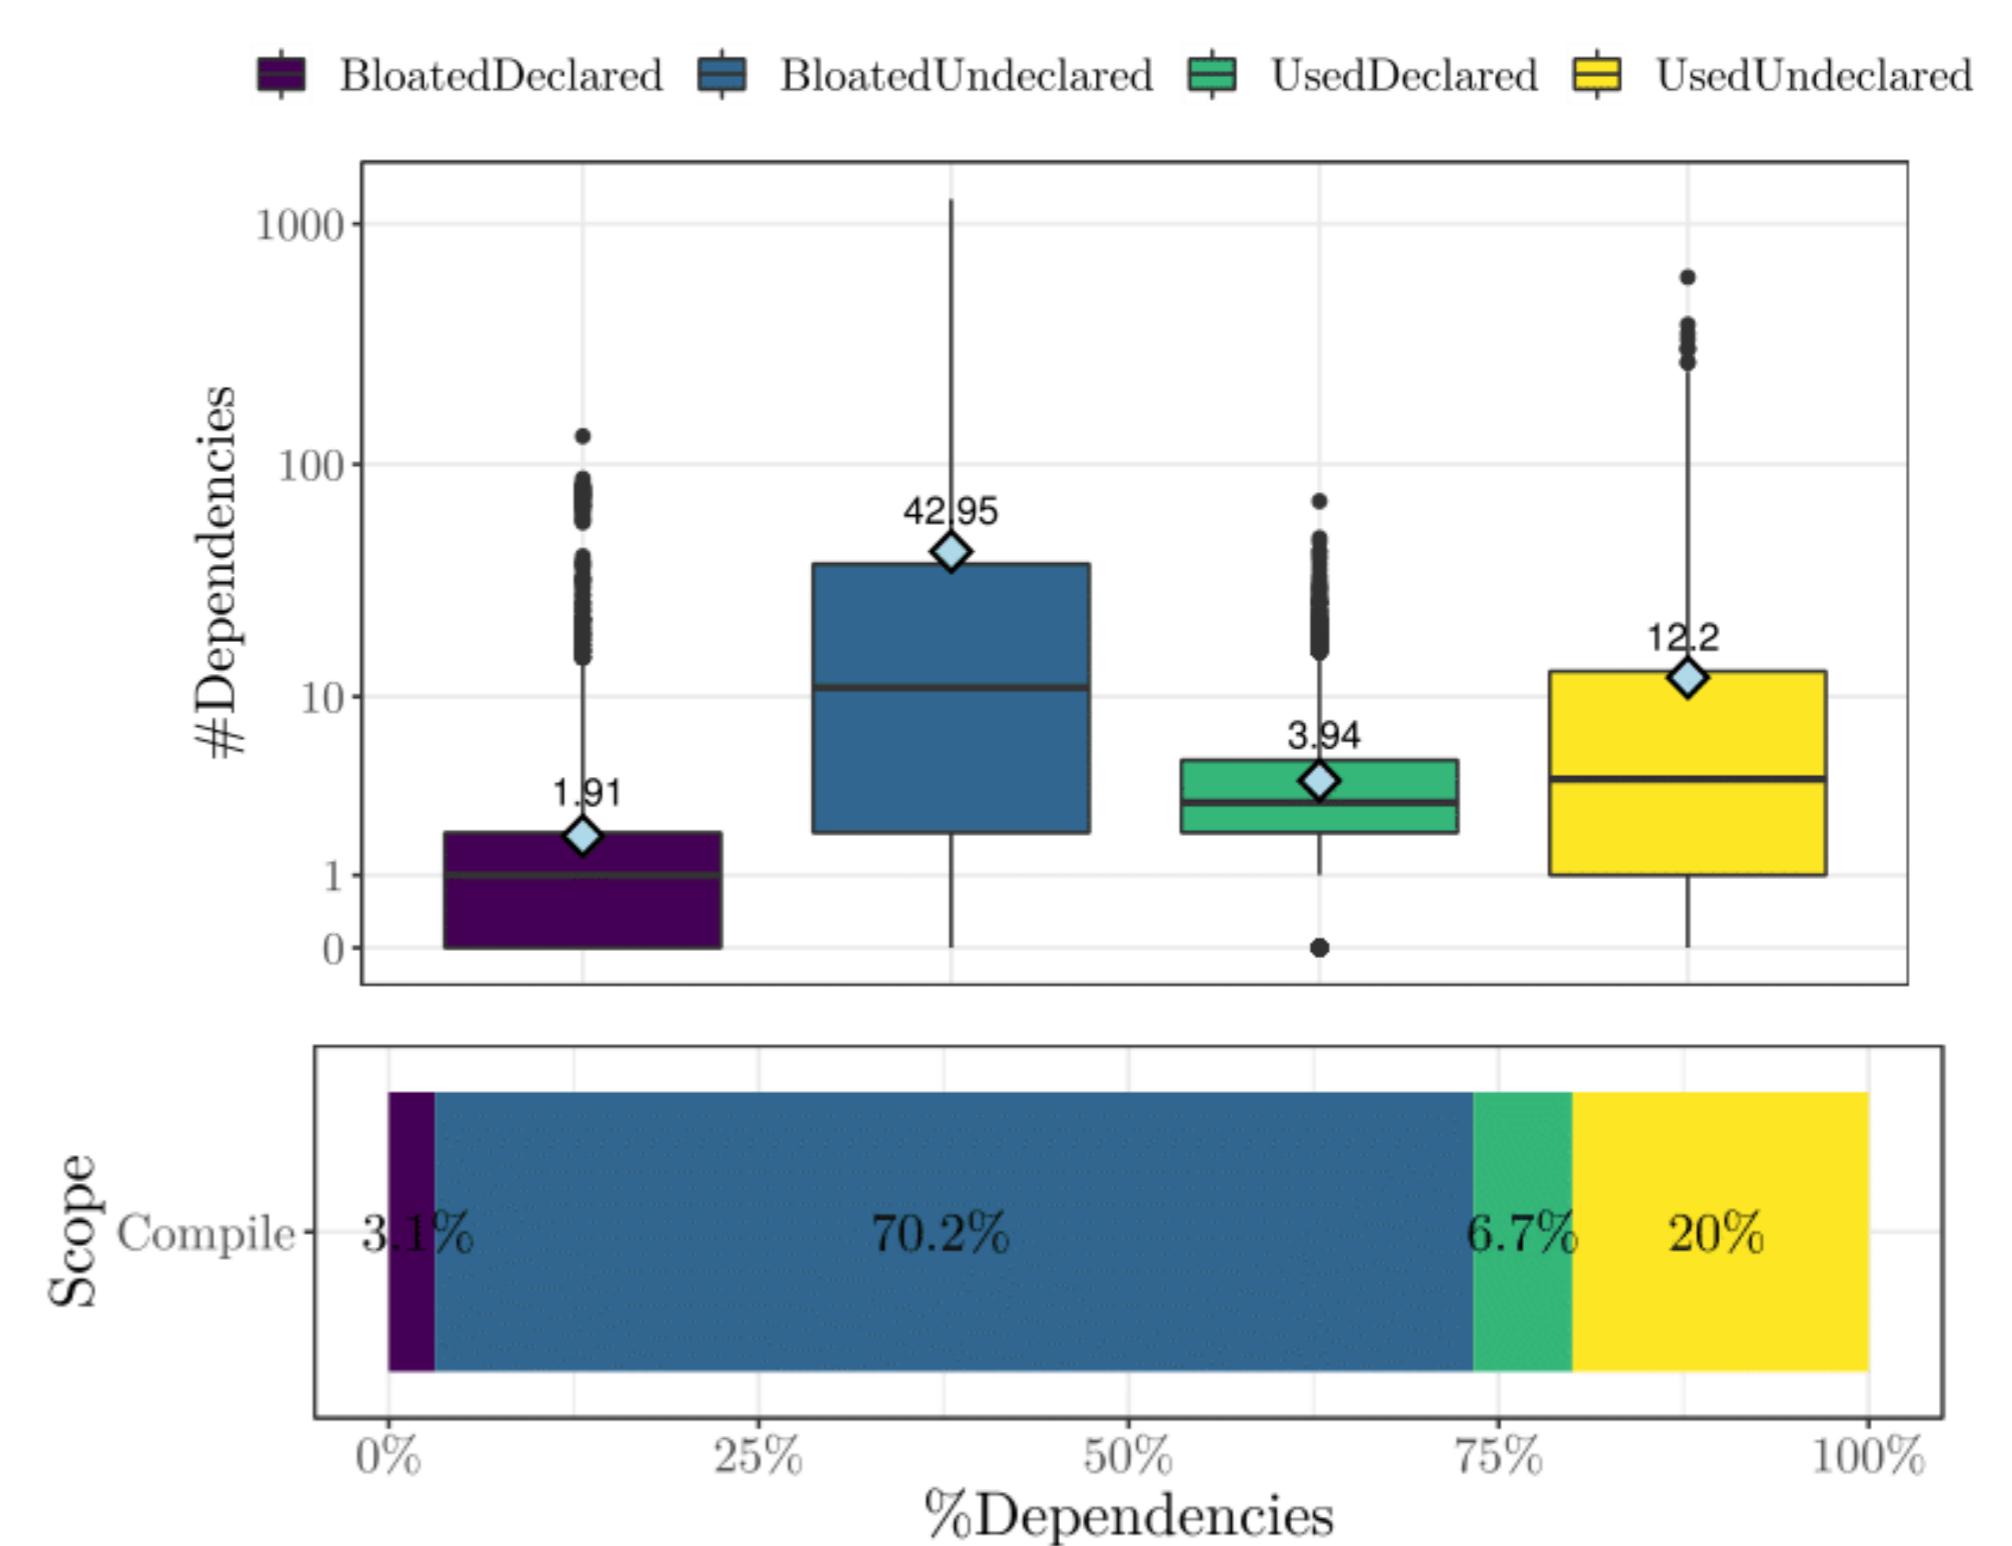
Methods

- We extend the Maven Dependency Graph by labeling edges between artifacts depending on their usage status



Results

- 73% of dependencies in our dataset are bloated
- the average number of bloated dependencies per project is 45



Road map

Experimental framework to collect and analyze bloated dependencies in the Maven ecosystem

