On the tree-transformation power of XSLT

Wim Janssen   Alexandr Korlyukov[†]   Jan Van den Bussche[∗]

Abstract

XSLT is a standard rule-based programming language for express-ing transformations of XML data. The language is currently in tran-sition from version 1.0 to 2.0. In order to understand the computa-tional consequences of this transition, we restrict XSLT to its pure tree-transformation capabilities. Under this focus, we observe that XSLT 1.0 was not yet a computationally complete tree-transformation language: every 1.0 program can be implemented in exponential time. A crucial new feature of version 2.0, however, which allows node sets over temporary trees, yields completeness. We provide a formal opera-tional semantics for XSLT programs, and establish confluence for this semantics.