STEVENS
INSTITUTE OF TECHNOLOGY
1870

Ph.D. DISSERTATION DEFENSE

Candidate:	Ramana Nagasamudram
Degree:	Doctor of Philosophy
School/Department:	Charles V. Schaefer, Jr. School of Engineering and Science / Computer Science
Date:	Tuesday, July 29th, 2025
Time/Location:	12:00 p.m. / Gateway North 303
Title:	Auto-active relational verification and alignment completeness
Chairperson:	Dr. David Naumann, Department of Computer Science, Institute of Technology
Committee Members:	Dr. Anindya Banerjee, Department of Computer Science, Dartmouth College Dr. Lennart Beringer, Department of Computer Science, Princeton University Dr. Eduardo Bonelli, Department of Computer Science, Stevens Institute of Technology Dr. Eric Koskinen, Department of Computer Science, Stevens Institute of Technology

ABSTRACT

Establishing relations between programs arises as a task in various verification contexts such as relating new versions of programs with older versions or proving the correctness of program transformations. Existing tools for relational verification provide a high degree of automation at the cost of restricting the class of programs handled. Auto-active tools such as Dafny and Why3, on the other hand, require more user interaction and support verification of a broad class of programs, including those that act on pointers. However, they don't provide native facilities for relational verification. In the first part of this thesis, we introduce WhyRel, an auto-active tool for relational verification that bridges this gap. We evaluate WhyRel on challenging case studies including establishing representation independence of data types and proving the correctness of program optimizations.

In the second part of this thesis, we study the relational Hoare logics (RHLs) that underly tools like WhyRel. The key to compositional reasoning is the alignment of computation steps. RHLs provide a considerable number of rules that embody various kinds of alignments, some seemingly more expressive than others. However, a single degenerate alignment rule that reduces relational reasoning to unary reasoning suffices to make a RHL complete. Thus, the usual notion of completeness doesn't offer a way to distinguish between RHLs or shed light on the rules a relational logic should include. In prior work, we introduced alignment completeness as a more satisfactory measure of RHLs and proved alignment completeness of a few RHL rules with respect to ad hoc forms of alignment. We extend these results and prove alignment completeness for a RHL with respect to a very general class of alignments. Finally, we introduce a new relational program logic for forward simulation and prove it is both alignment complete and complete in the usual sense.