BCRI MINI-SYMPOSIUM

LOGIC, SEMANTICS AND COMPLEXITY IN COMPUTER SCIENCE

Tuesday 17th October 2023

University College Cork
Room 2.26 Western Gateway Building

Organiser

Professor Michel Schellekens
Department of Computer Science

Programme

09:30 – 10:00	Opening Address
10:15 – 11:15	**Stephen Brookes** *(Professor Computer Science, Carnegie Mellon University)* A denotational framework for relaxed memory concurrent programs
11:15 – 11:45	Tea & Coffee
11:45 – 12:45	**Michel Schellekens** (*Professor of Computer Science, University College Cork)* On sharp multitarget minimean bounds
12:45 – 14:00	Lunch
14:00 – 15:00	**Peter O'Hearn**, (*Engineering Manager FaceBook & Professor of Computer Science University College London UCL)* From Concurrent Separation Logic to Incorrectness Logic

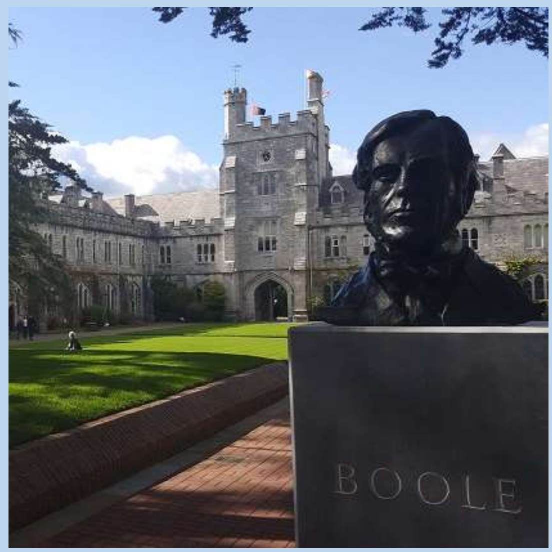

Organiser Contact

Michel Schellekens,
Department of Computer Science UCC
(m.schellekens@ucc.ie)

All interested are most welcome to attend this event.
Please register with m.dwane@ucc.ie for catering purposes

Sponsored by the Boole Centre for Research in Informatics

UCC University College Cork, Ireland Coláiste na hOllscoile Corcaigh | School of Computer Science and Information Technology

BOOLE CENTRE FOR RESEARCH IN INFORMATICS

Abstracts

A denotational framework for relaxed memory concurrent programs

Stephen Brookes Carnegie Mellon University

We introduce a denotational framework for coherent relaxed memory models, in which executions guarantee that all threads "see" the writes to each individual variable in the same relative order. Each memory model also allows per-thread relaxation of the relative order of certain kinds of action. The semantics is "truly concurrent", using partially ordered multisets of actions (pomsets), a natural generalization of traces. The semantic constructions, such as sequential and parallel composiition, are parameterized by the memory model, embodied by a combination of a relaxation relation and a linearity guarantee. We focus on three historically important memory models: SC, TSO and PSO, which guarantee (progressively weaker forms of) coherence. For SC we recover the usual trace semantics familiar from traditional accounts. We illustrate how the semantic models work by examining litmus tests, small programs whose executional behavior differs across memory models. We discuss related work, and point the way to future developments.

From Concurrent Separation Logic to Incorrectness Logic

Peter O'Hearn, Engineering Manager FaceBook & Professor of Computer Science University College London

In the early 2000s Steve Brookes proved soundness of Concurrent Separation Logic, for which I will always be grateful: the problem was extremely difficult, beyond my capabilities, and his proof saved the logic. Later, in in 2016, I was working in industry, and I had the itch to apply CSL, and I hatched a plan to apply it to the internal codebases at Facebook. My aim was to automatically prove race freedom of hundreds of thousands or even millions of lines of code. Product demands led us to pivot from this idealistic goal, but what we achieved by compromising ended up being much better than the idealistic goal, and eventually ended in Incorrectness Logic. In this talk I'll tell the story of the twists and turns along the way that happened when program logic theory bumped up against industrial software practice.

Sharp multitarget minimean bounds

M. Schellekens University College Cork School of Computer Science and IT

Abstract Partial order production yields generalised complexity lower-bounds for comparison-based algorithms that need not be a sort or a search. Input topological sorts over a source order α produce output topological sorts over a single refined target β. Andrew Yao's minimean complexity lower-bound is the variation H(α)−H(β) between source and target entropy, for H(γ) = log2(|γ|) and |γ| the topological sort count over γ. This ratio-bound R = log2(|α|/|β|) reduces to the log2(n!) entropy bound for the sorting case, but is not sharp in general. This raises the question when sharp bounds arise, while Yao queried the multi-target production complexity. A multi-target bound R is derived for two classes of calibrated refiners–algorithms producing targets in a multiset partition. The bound coincides with Yao's in case of a single target. Perfect calibrators producing uniformly calibrated multisets, such as Quicksort's Partitioning, have a minimean sharply R-bounded by the variation H(α) − H(β1, . . . , βm) between source and *average* target entropy. Quicksort's redundancy-free fragmentation sheds new light on its superior average-time performance. Fractional calibrators, including the perfect calibrators, satisfy a generalised ratio bound. Fragmentation theoretic analysis sharpens Yao's Find-Max bound and yields Heap-Fusion's sharp entropy-bound.

∗ The work was carried out in part during a Fulbright Scholarship at Stanford University, Apr 5 - Aug 31, 2019, research host: D. Knuth.