

This is an extraordinary paper, and nothing like this has been done previously on this scale – formally proving critical security properties of the specification for existing prototype Arm Morello hardware based on the U.Cambridge CHERI-Arm-Morello formal specification. The work was funded by projects that I happen to have been leading, with DARPA and some UK funding as well. Peter G. Neumann, PhD and *dr rerum naturalium*, Chief Scientist SRI International Computer Science Lab, Menlo Park, California 1-650-804-1802 [Neumann@CSL.SRI.COM](mailto:Neumann@CSL.SRI.COM).

**ABSTRACT:** Memory safety bugs continue to be a major source of security vulnerabilities in our critical infrastructure. The CHERI project has proposed extending conventional architectures with hardware-supported capabilities to enable fine-grained memory protection and scalable compartmentalisation, allowing historically memory-unsafe C and C++ to be adapted to deterministically mitigate large classes of vulnerabilities, while requiring only minor changes to existing system software sources. Arm is currently designing and building Morello, a CHERI-enabled prototype architecture, processor, SoC, and board, extending the high-performance Neoverse N1, to enable industrial evaluation of CHERI and pave the way for potential mass-market adoption. However, for such a major new security-oriented architecture feature, it is important to establish high confidence that it does provide the intended protections, and that cannot be done with conventional engineering techniques. In this paper we put the Morello architecture on a solid mathematical footing from the outset. We define the fundamental security property that Morello aims to provide, reachable capability monotonicity, and prove that the architecture definition satisfies it. This proof is mechanised in Isabelle/HOL, and applies to a translation of the official Arm specification of the Morello instruction-set architecture (ISA) into Isabelle. The main challenge is handling the complexity and scale of a production architecture: 62,000 lines of specification, translated to 210,000 lines of Isabelle. We do so by factoring the proof via a narrow abstraction capturing essential properties of arbitrary CHERI ISAs, expressed above a monadic intra-instruction semantics. We also develop a model-based test generator, which generates instruction-sequence tests that give good specification coverage, used in early testing of the Morello implementation and in Morello QEMU development, and we use Arm’s internal test suite to validate our model.