THOMAS STERLING

Dr. Sterling is a Professor at the IU School of Informatics and Computing Department of Intelligent Systems Engineering and serves as Director of CREST. Since receiving his Ph.D., he has engaged in applied research in parallel computing system structures, semantics, and operation in industry, government labs, and academia. Dr. Sterling is best known as the "father of Beowulf" for his pioneering research in commodity/Linux cluster computing for which he shared the Gordon Bell Prize in 1997. Sterling is currently involved in research associated with the innovative ParalleX execution model for extreme scale computing to establish the foundation principles guiding the development of future generation Exascale computing systems. ParalleX is currently the conceptual centerpiece of the XPRESS project as part of the DOE X-stack program and has been demonstrated via the proof-of-concept HPX-5 runtime system software.

High Performance Computing for Artificial Intelligence – not just for Neural Nets

THOMAS STERLING

Data Analytics and Machine Learning are disciplines of intense interest, almost defining current trends in artificial intelligence with significant accomplishments at hand. The US 2015 Presidential Executive Order, the National Strategic Computing Initiative (NSCI), calls for the unification of High Performance Computing (HPC) and big data applications, including machine learning, on the path to future exascale computing by the early part of the next decade. Even more innovative concepts include neuro-morphic processing that is brain-inspired and takes many forms with the presumption that if we can make the equivalent of a neuron, we can make the equivalence of the human brain and therefore the property of intelligence. With great excitement, such approaches are being pursued in the US and internationally again with interesting accomplishments. It is the speaker's opinion that none of these strategies will lead to the ultimate achievement of machine intelligence any more than building a chess-playing machine – or, for that matter, a theorem-proving algorithm - will yield a system that can derive general relativity, recognize a metaphor, or prepare a soufflé. It is asserted that intelligence is an algorithm, not the mimicking individual idioms of a lemur's mental acuity. Fundamentally, machine intelligence requires the ability to understand and manipulate knowledge in response to newfound objective functions unless of course we're prepared to accept the logical capabilities of a grasshopper. The intelligence algorithm has not been derived; indeed, the metrics of intelligence are ill formed. But prior art does expose many of the functional primitive that machine intelligence will have to be able to perform expeditiously and therefore capabilities required of such a future machine. This presentation will give a hypothesized high level architecture and workflow of an intelligent machine and from these identify the key primitive operations that are likely to be key to ultimately achieving the emergent property of intelligence from non-biological systems.

High Performance Computing for Artificial Intelligence – not just for Neural Nets