M. Davies et al., "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning," in IEEE Micro, vol. 38, no. 1, pp. 82-99, January/February 2018, doi: 10.1109/MM.2018.112130359.

The paper's **general topic** is neuromorphic computing and spiking neural networks in hardware. As the authors mention: "Neuroscience offers a bountiful source of inspiration for novel hardware architectures and algorithms," and "The neuromorphic-computing field of research spans a range of different neuron models and levels of abstraction." The **specific behavior or activity** being studied is spike-based computation and on-chip learning. As noted in the paper, "Loihi is a 60-mm^2 chip fabricated in Intel's 14-nm process that advances the state-of-the-art modeling of spiking neural networks in silicon," and "Running a spiking convolutional form of the Locally Competitive Algorithm, Loihi can solve LASSO optimization problems with over three orders of magnitude superior energy-delay product compared to conventional solvers running on a CPU isoprocess/voltage/area." The paper does not pose a single explicit **research question**, but the closest formulation of their aim is: "to develop algorithms and hardware in a principled way as much as possible," and "We present a result that unambiguously demonstrates the value of spike-based computation for one foundational problem." So the implicit research question is whether spike-based computation, supported by a new neuromorphic architecture, can provide real computational value beyond conventional architectures.

The main **problem** identified is that conventional hardware does not serve SNNs well: "Spiking neural network (SNN) models, on the other hand, are exceedingly poorly served by conventional architectures," and that "Flexible and well-provisioned SNN connectivity features are crucial for supporting a broad range of workloads." The main **claims** are that Loihi is a novel digital SNN chip with programmable learning and that it demonstrates the value of spike-based computation. The **evidence** is quantitative and architectural; the authors describe the chip design, report energy and performance measurements, and compare Loihi's sparse-coding performance against CPU-based LARS and FISTA solvers. The paper's **evaluation** relies on benchmarking tables and comparative metrics such as energy, delay, and EDP.