

SACRED-MA: Safe And seCure REmote Direct Memory Access

Gregory Chockler
University of Surrey

SACRED-MA: General Info

- Funded by EPSRC under the auspices of VeTSS
- Investigators: Brijesh Dongol, Surrey (PI), Azalea Raad, Imperial (PI), Gregory Chockler, Surrey (Co-I)
- Post-docs: Guillaume Ambal (Imperial), Milad Ketabi (Surrey)
- Partners: NVIDIA, Tel-Aviv University, University of Colorado Boulder, Cornell, MPI


RDMA (Remote Direct Memory Access)



- A new, fast, network protocol, off-loaded to the network card...

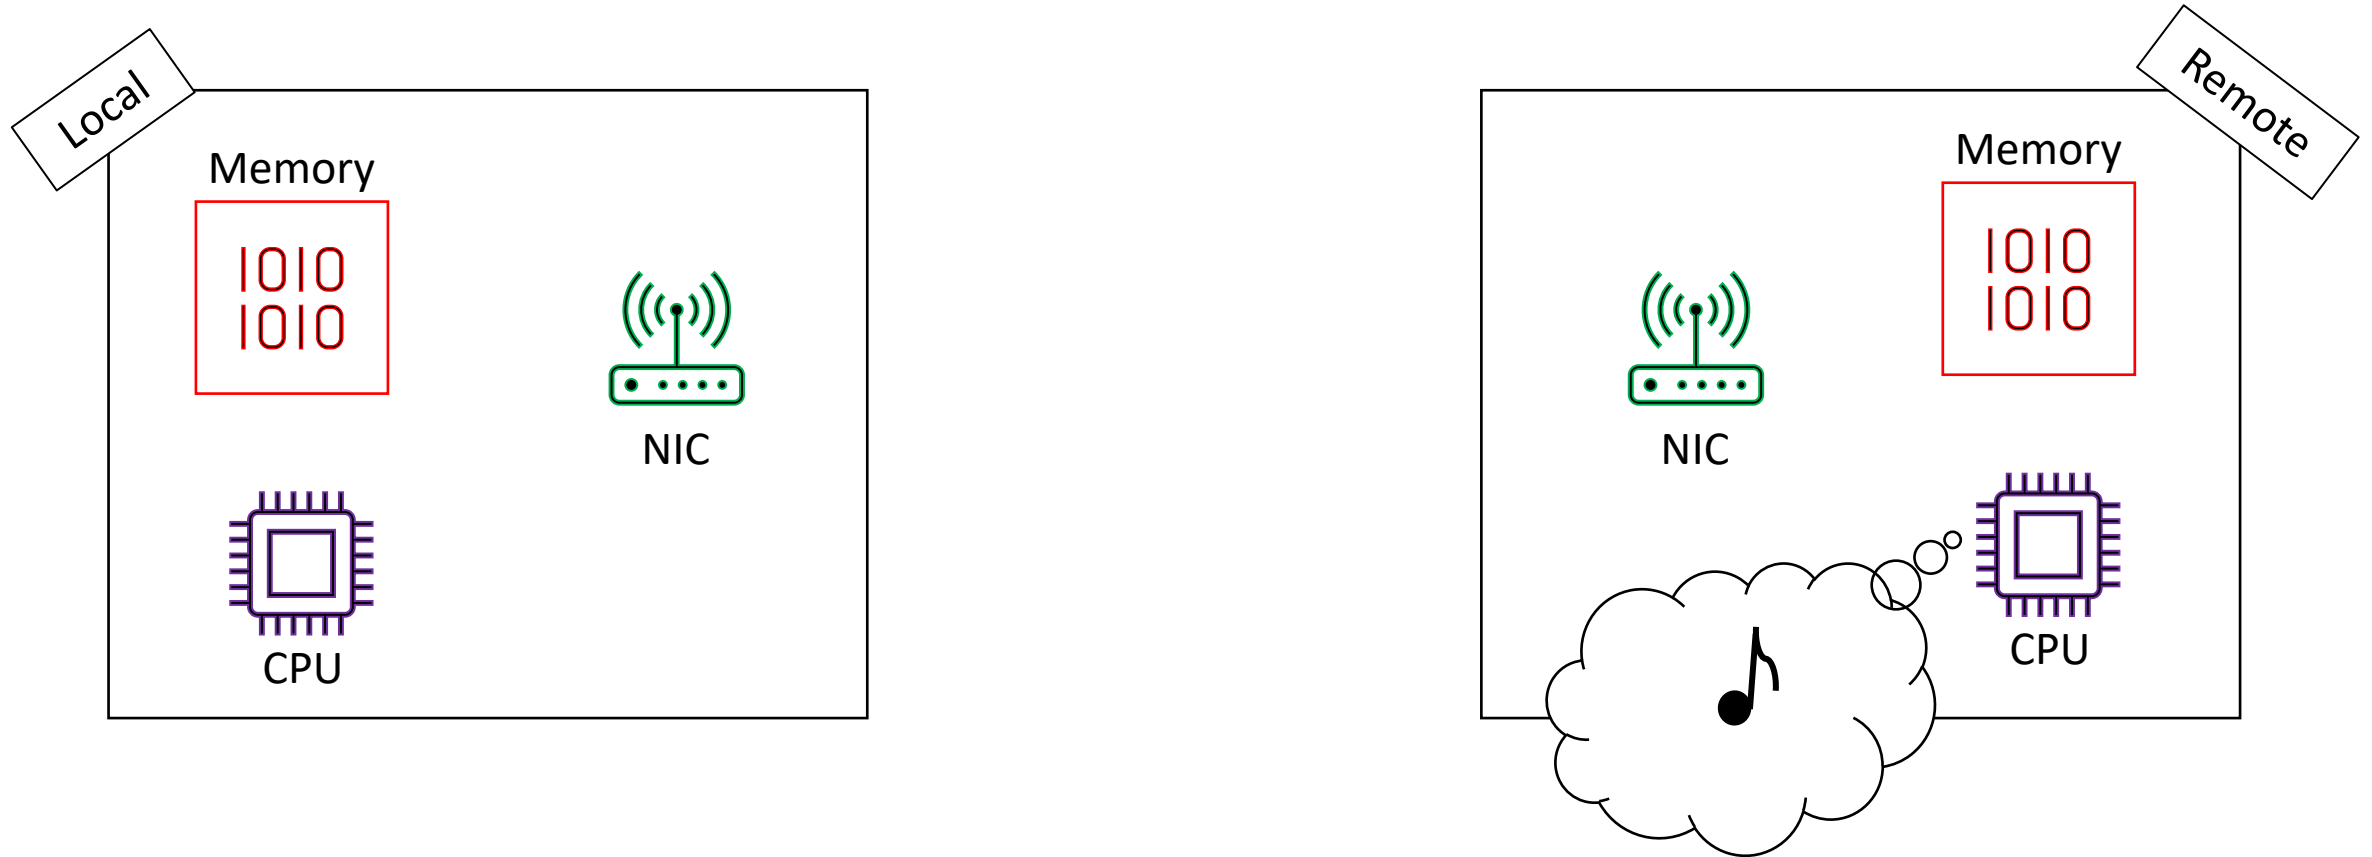
RDMA (Remote Direct Memory Access)

A new, fast, network protocol, off-loaded to the network card...



RDMA (Remote Direct Memory Access)

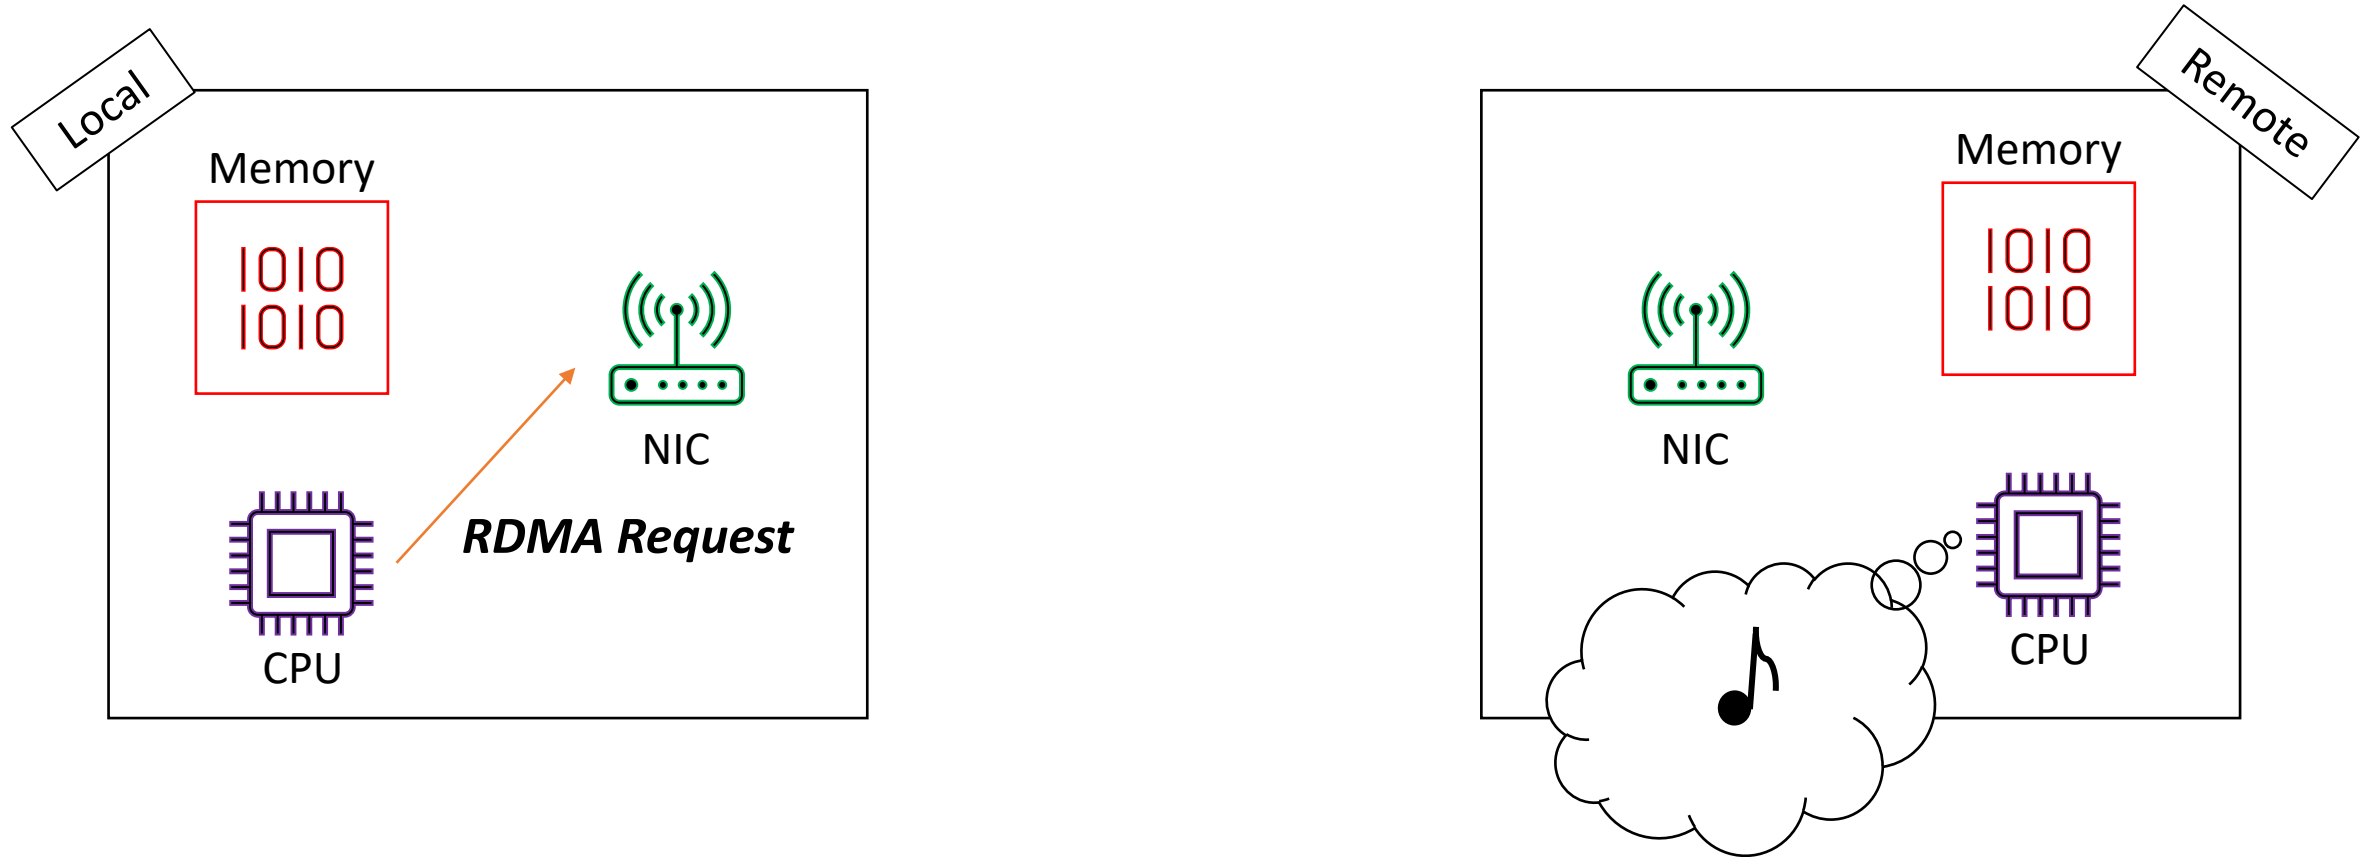
A new, fast, network protocol, off-loaded to the network card...



Bypasses CPU and Kernel at the remote node allowing us to have a faster network layer

RDMA (Remote Direct Memory Access)

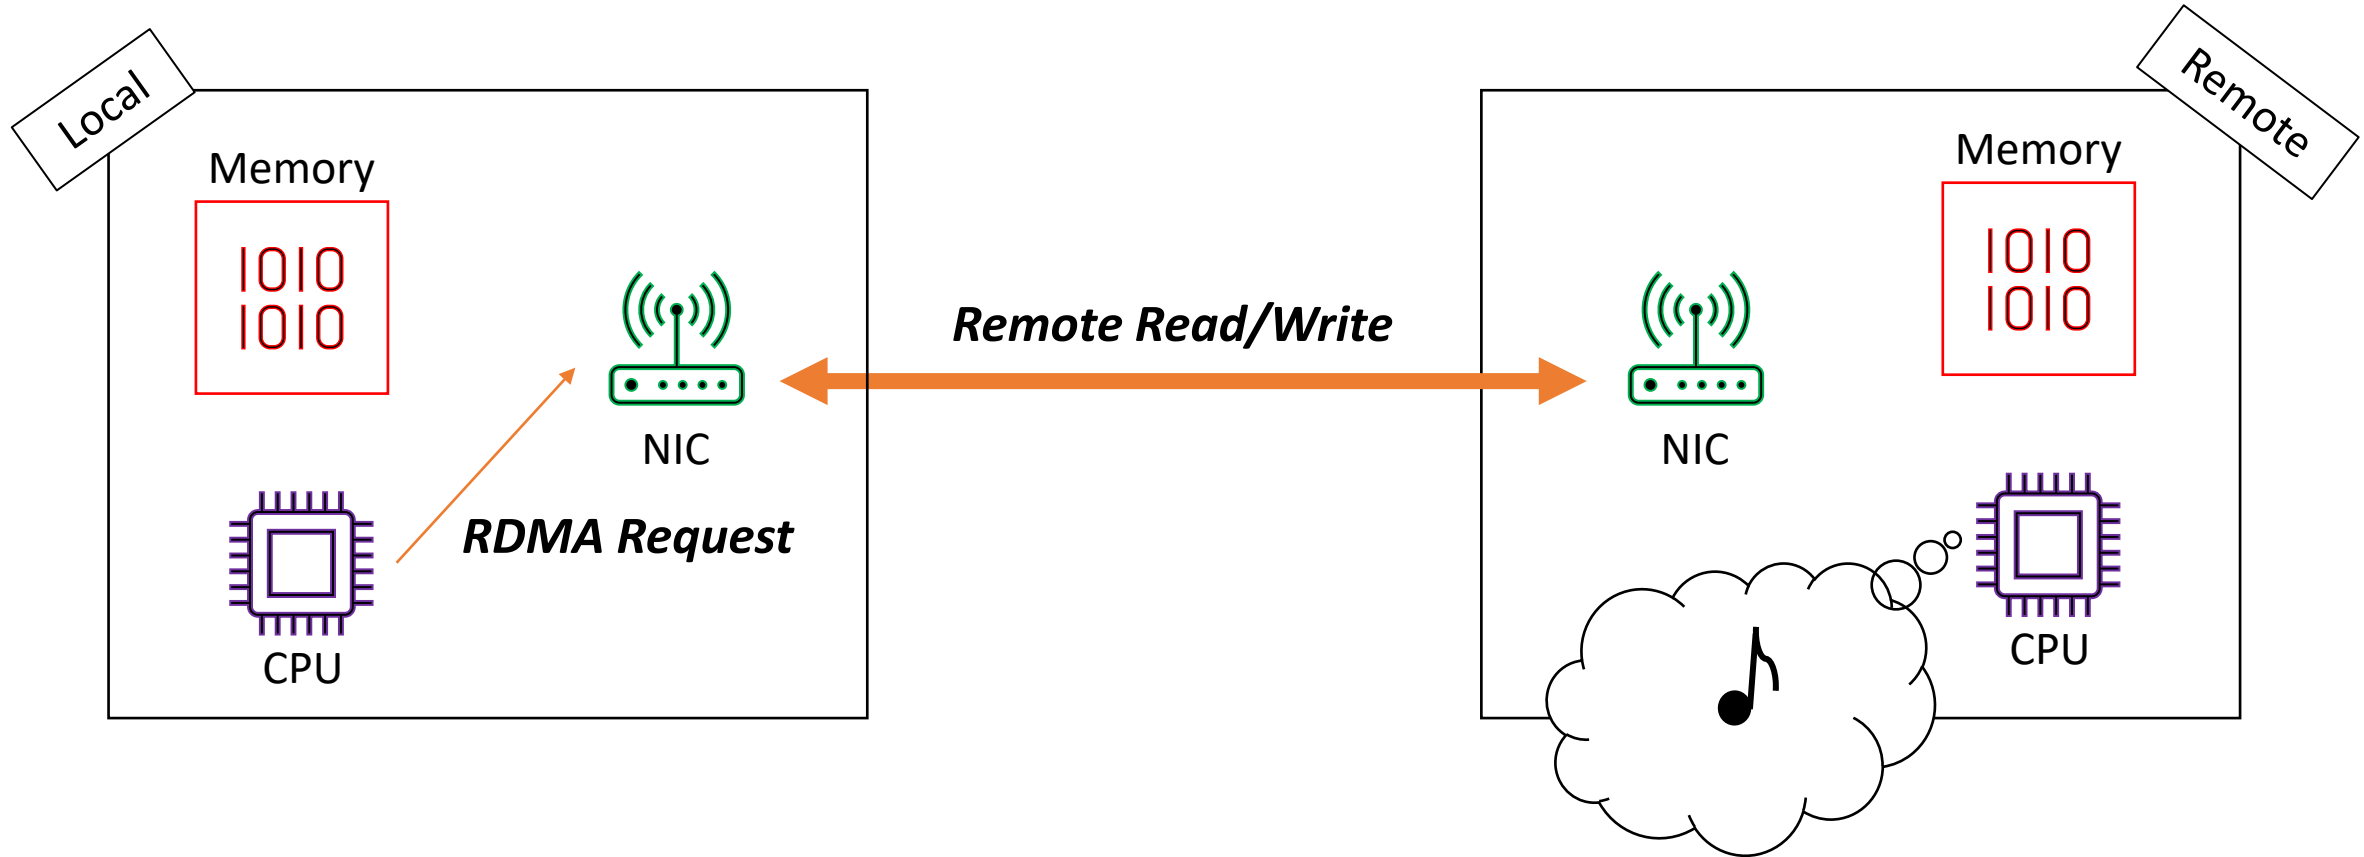
A new, fast, network protocol, off-loaded to the network card...



Bypasses CPU and Kernel at the remote node allowing us to have a faster network layer

RDMA (Remote Direct Memory Access)

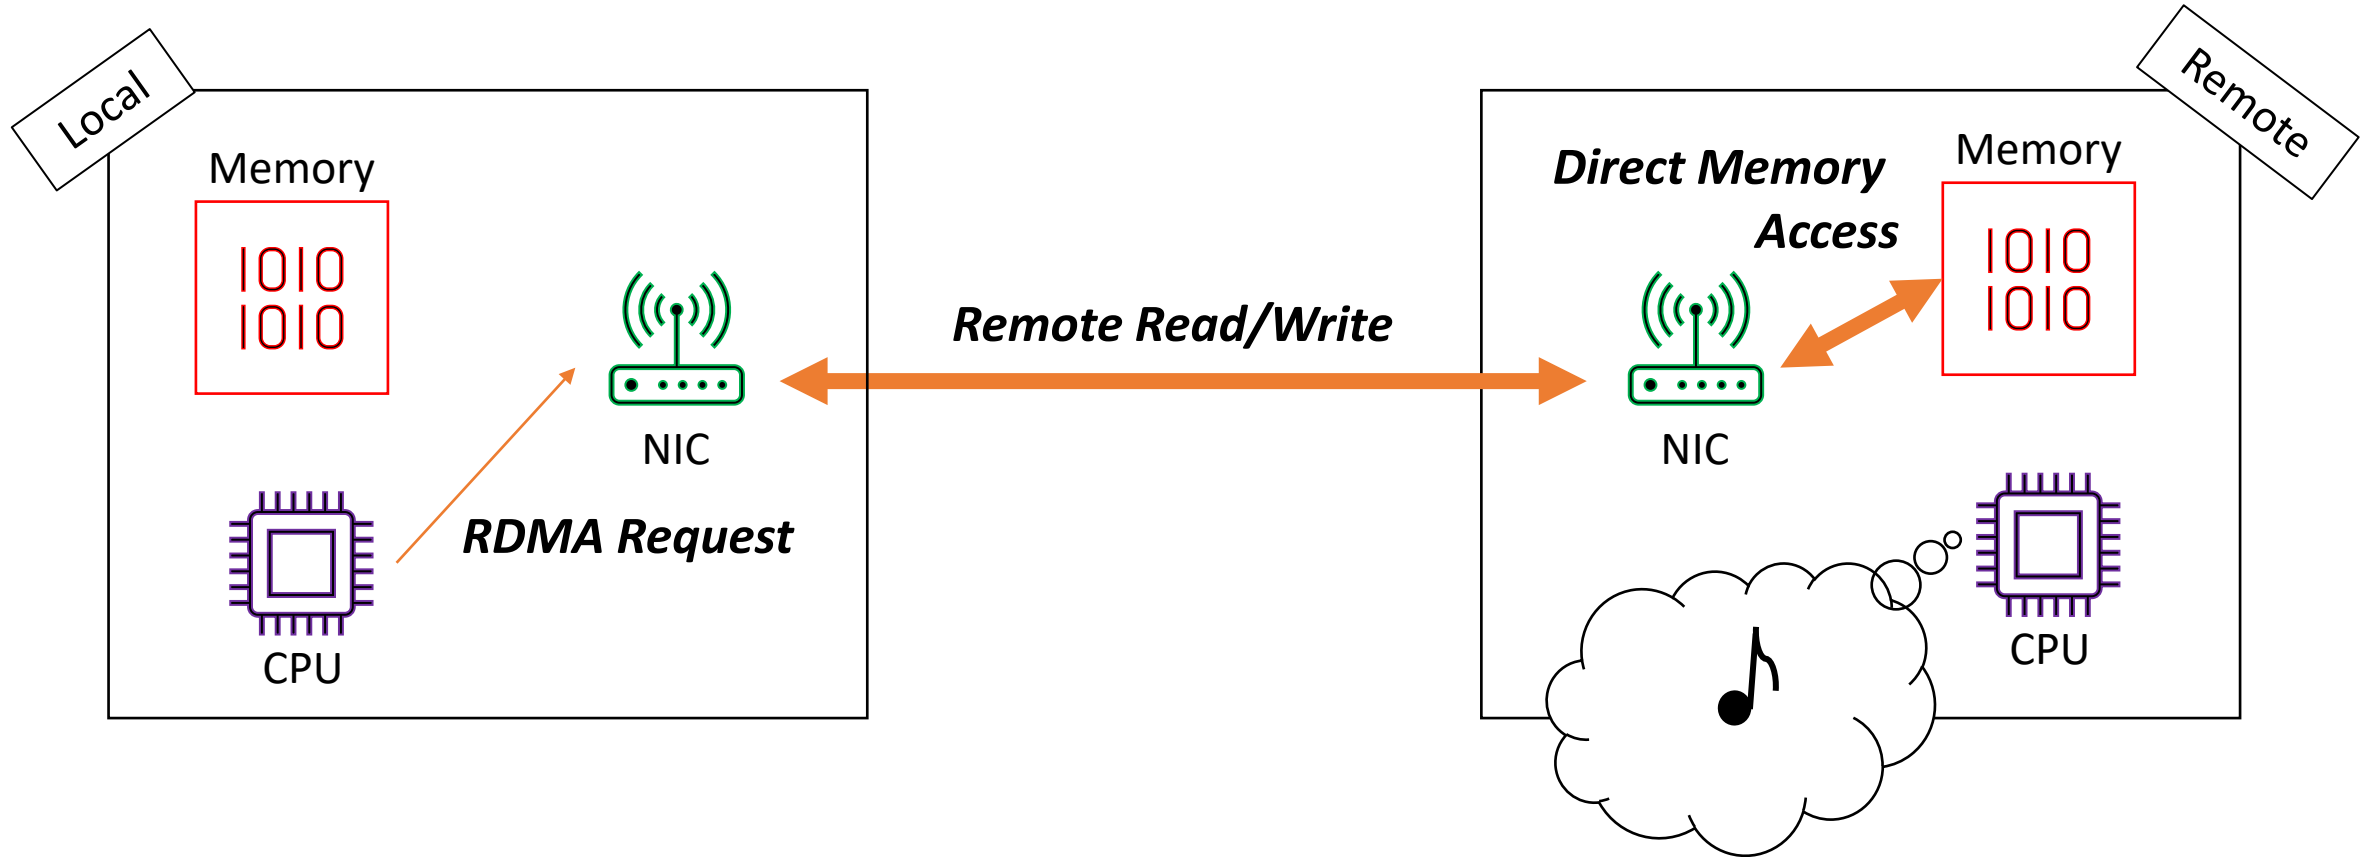
A new, fast, network protocol, off-loaded to the network card...



Bypasses CPU and Kernel at the remote node allowing us to have a faster network layer

RDMA (Remote Direct Memory Access)

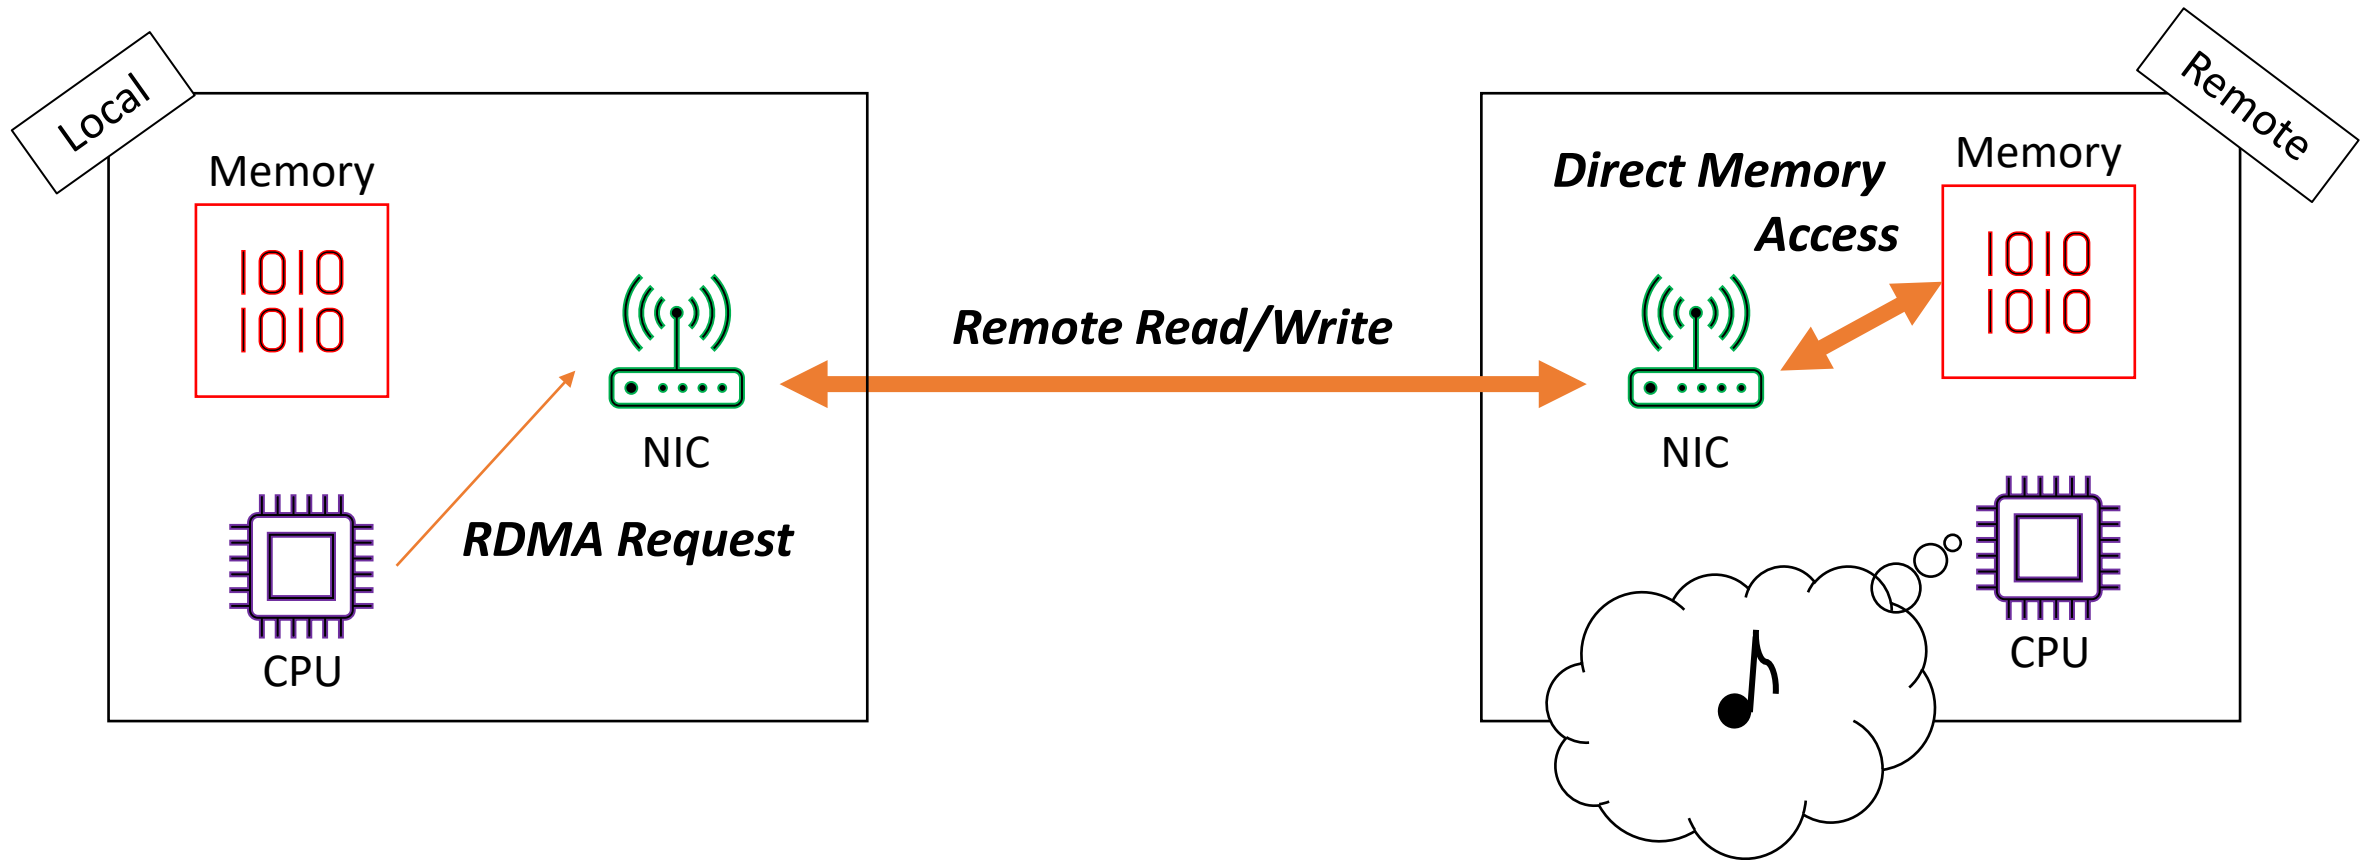
A new, fast, network protocol, off-loaded to the network card...



Bypasses CPU and Kernel at the remote node allowing us to have a faster network layer

RDMA (Remote Direct Memory Access)

A new, fast, network protocol, off-loaded to the network card...



*RDMA is fast, but the interface and properties are different ... how can we leverage for **distributed computing**?*

RDMA for Distributed Computing

- **Low latency** point-to-point communication
 - One-sided reads/writes: comparable to shared memory
- **FIFO** memory updates
 - Acknowledging message receipt acknowledges all priors
- Efficient **catch-up** for slow replicas
 - Only need to acknowledge recent messages as they implicitly acknowledge all priors
- **Completion** mechanism to avoid explicit acknowledgments
- **Permissions** to prevent nodes from accessing unauthorized regions

Gaps

- Most current work is focused on **systems engineering** and performance improvements
- Correctness/security issues are often neglected/overlooked
- Existing specifications are ambiguous or incomplete
- Semantics of interaction with local memory models is not well understood

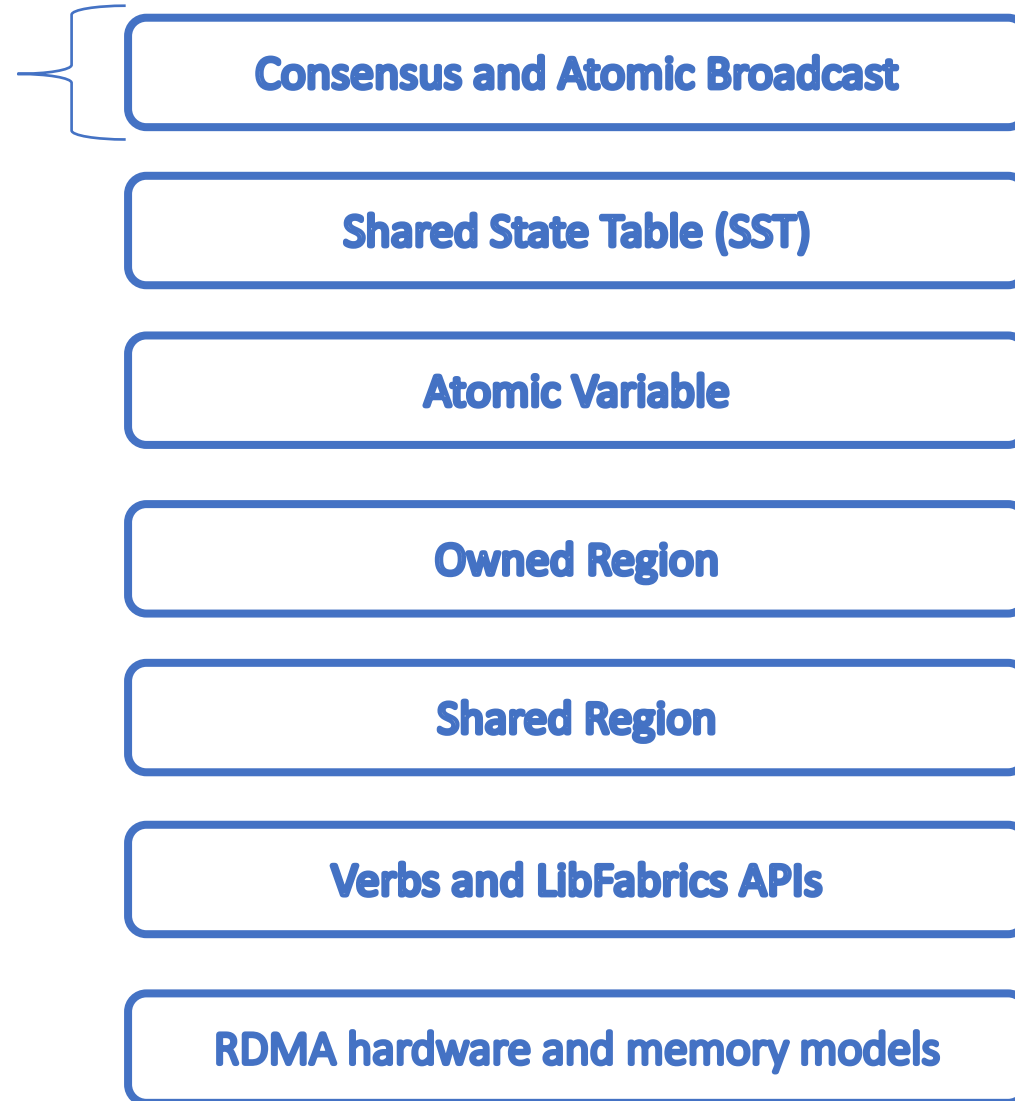
SACRED-MA: Safe And seCure REmote Direct Memory Access



- Develop robust RDMA models suitable for formal correctness reasoning and verification of distributed protocols
- Full formally verified RDMA stack
 - From formal hardware models
 - To high-level abstractions useful for safe and secure distributed computing

Formally verified RDMA stack

Izraelevitz, Wang,
Hanscom, Silvers, Lehman,
Chockler, Gotsman,
**Acuerdo: Fast Atomic
Broadcast over RDMA,**
ICPP'22



Reusable compositional
libraries of safe and secure
abstractions

Hodgkins & Izraelevitz,
**LOCO: Rethinking Protocols as
Objects in Network Memory**

Current Status

Consensus and Atomic Broadcast

Shared State Table (SST)

Atomic Variable

Owned Region

Shared Region

Verbs and LibFabrics APIs

RDMA hardware and memory models

In-progress

Completed