

Towards Replicated-kernel OS Support for Task Migration on Heterogeneous-ISA Platforms

Marina Sadini (student), David Katz*, Antonio Barbalace, Alastair Murray, Binoy Ravindran

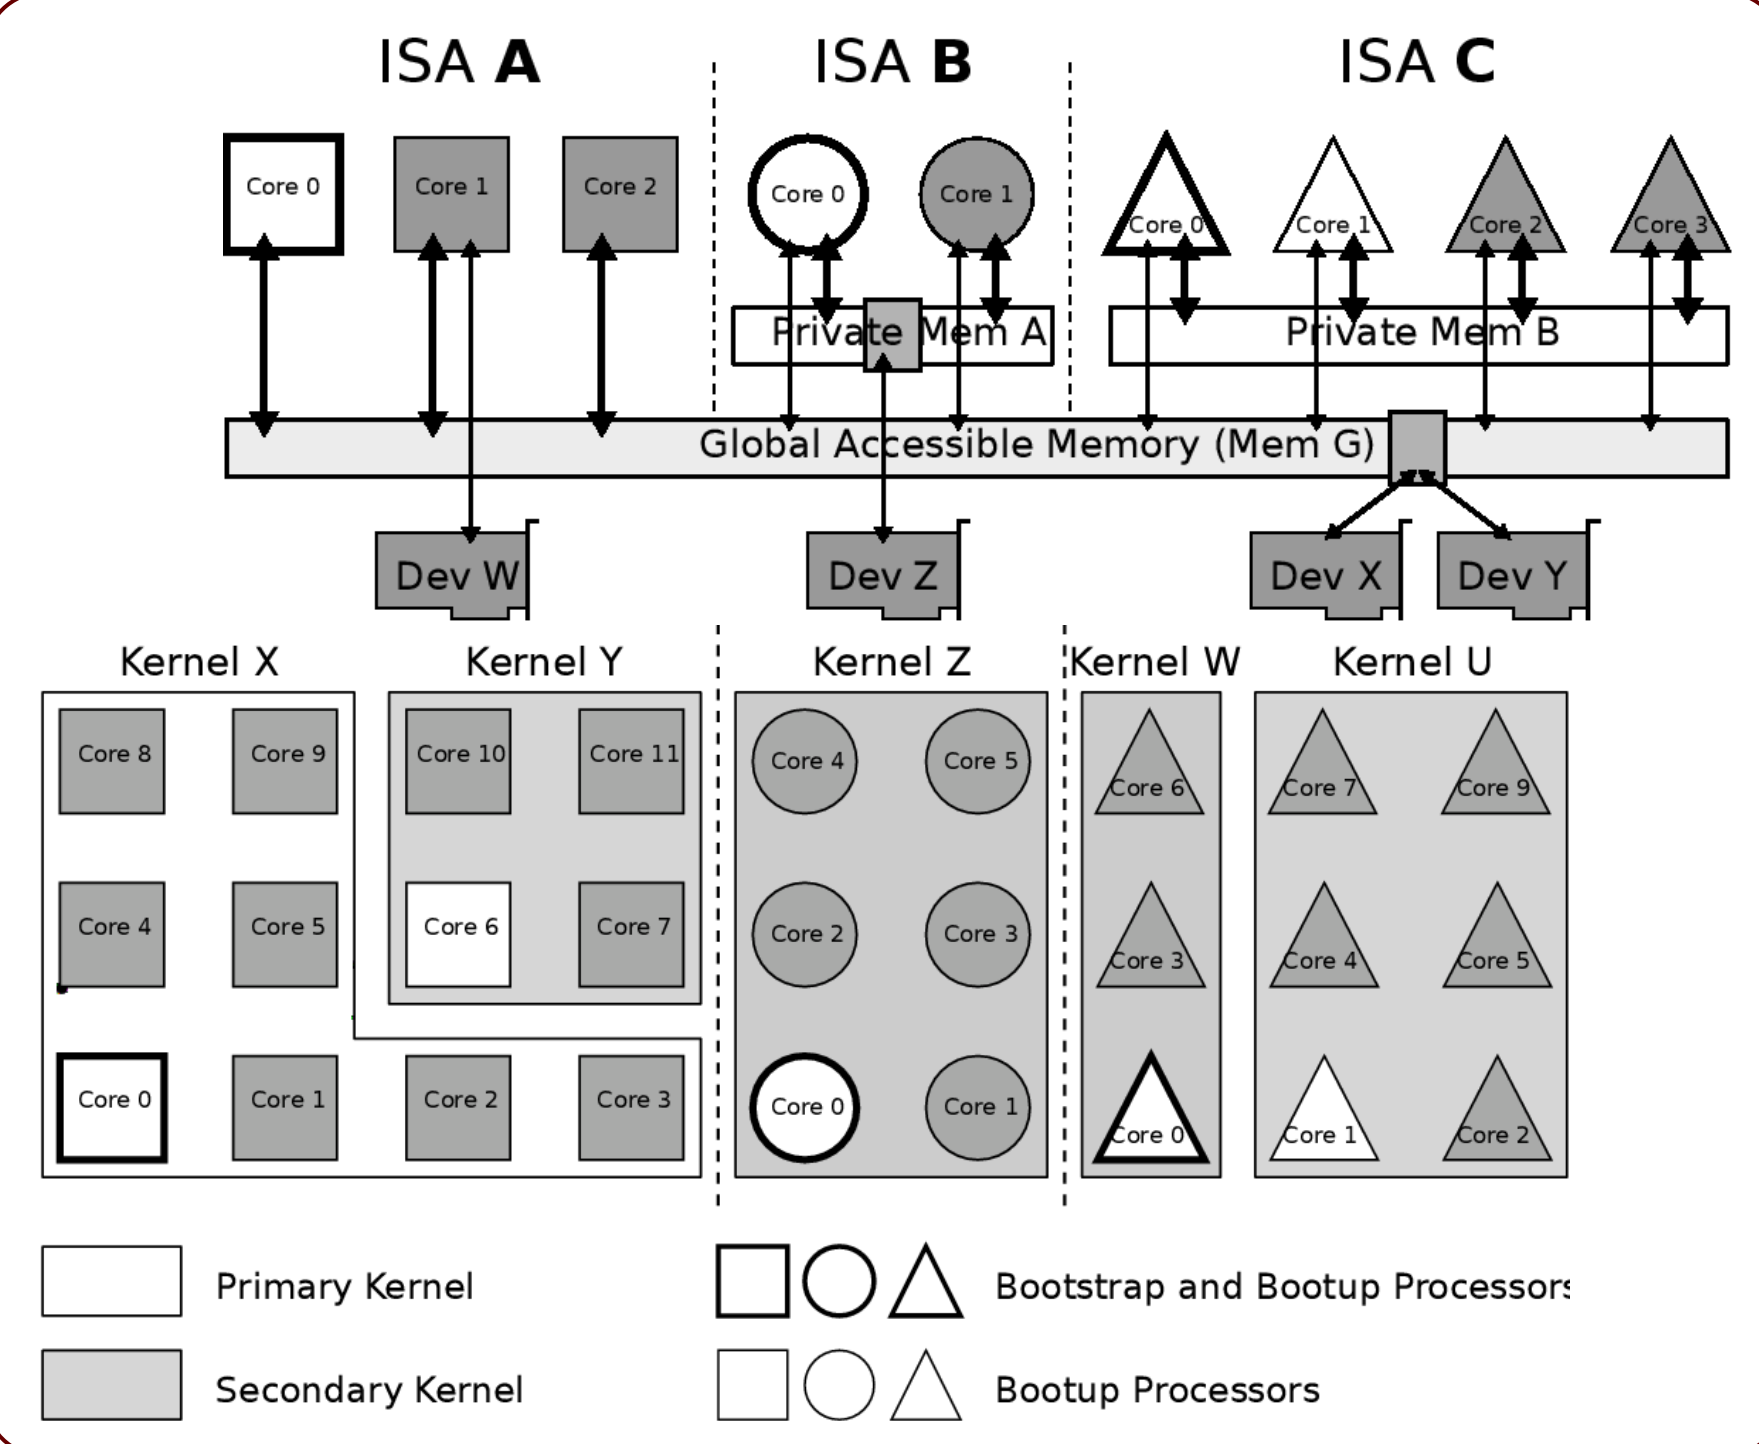
Department of Electrical and Computer Engineering, Virginia Tech, Virginia, USA

{sadini, davidk, antoniob, alastair, binoy@vt.edu}

*David Katz is with John Hopkins University Applied Physics Laboratory

Future computing platforms are increasingly ISA-diverse and parallel

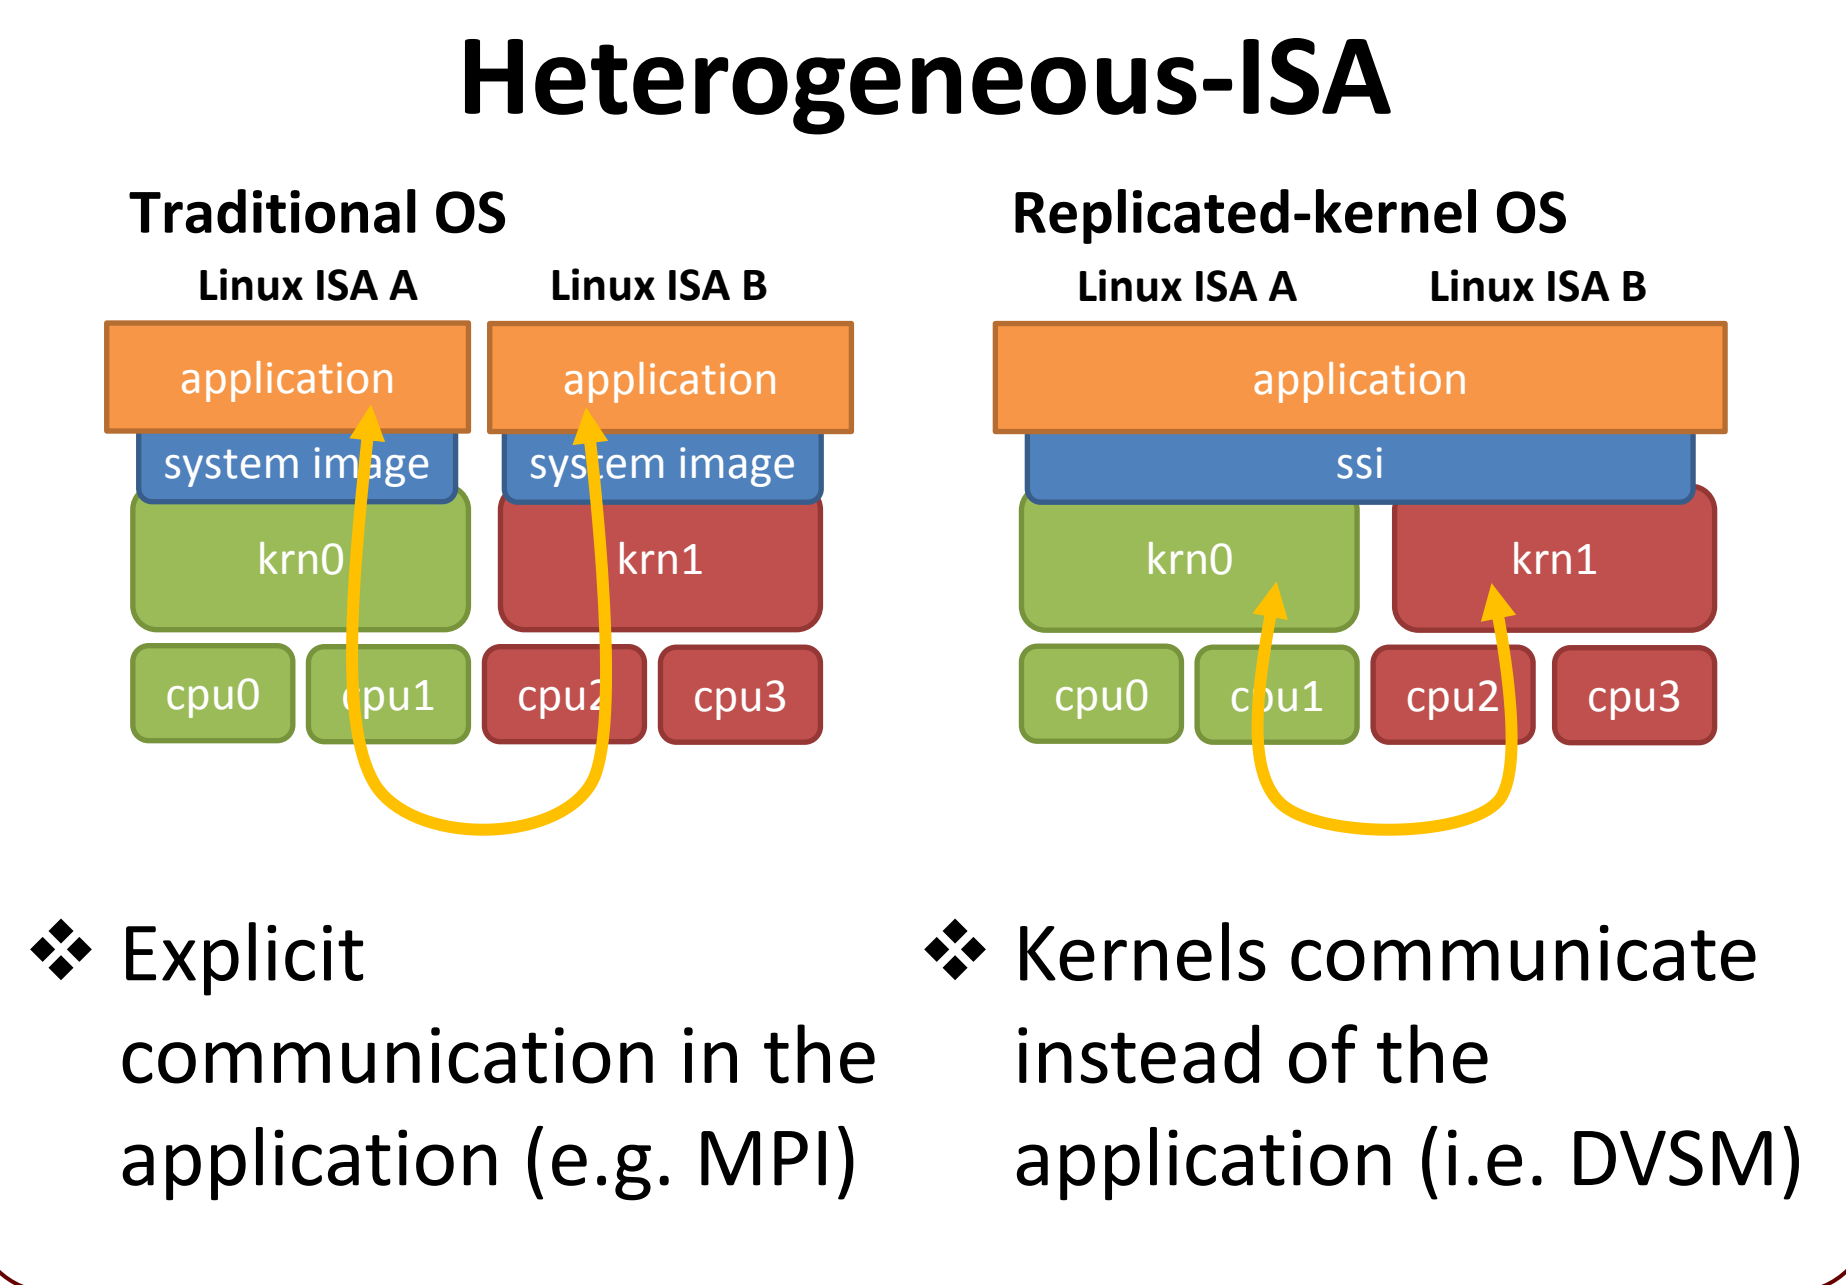
Replicated-kernel OS designs fully exploit new hardware



- ❖ Basic kernel-to-kernel communication is provided by a messaging layer (using shared memory and IPIs on x86)

- ❖ Traditional SMP single-image OSs (Linux, UNIX, Windows, etc.) cannot run on heterogeneous-ISA platforms as-is
- ❖ ISA diversity prevents the single-image per machine paradigm, one kernel (or more) per ISA group is needed
- ❖ In a replicated-kernel, different kernel instances are running in parallel on the same hardware
- ❖ Applications run transparently, like in a single-image operating system (constraints on the ABI exists)
- ❖ Kernels communicate to maintain a common OS state that is (partially) replicated over each individual kernel instance

SMP Applications run Transparently



- ❖ Explicit communication in the application (e.g. MPI)
- ❖ Kernels communicate instead of the application (i.e. DVSM)

Goal: To exploit all hardware resources

Task (Processes and Threads) inter-kernel migration, benefits:

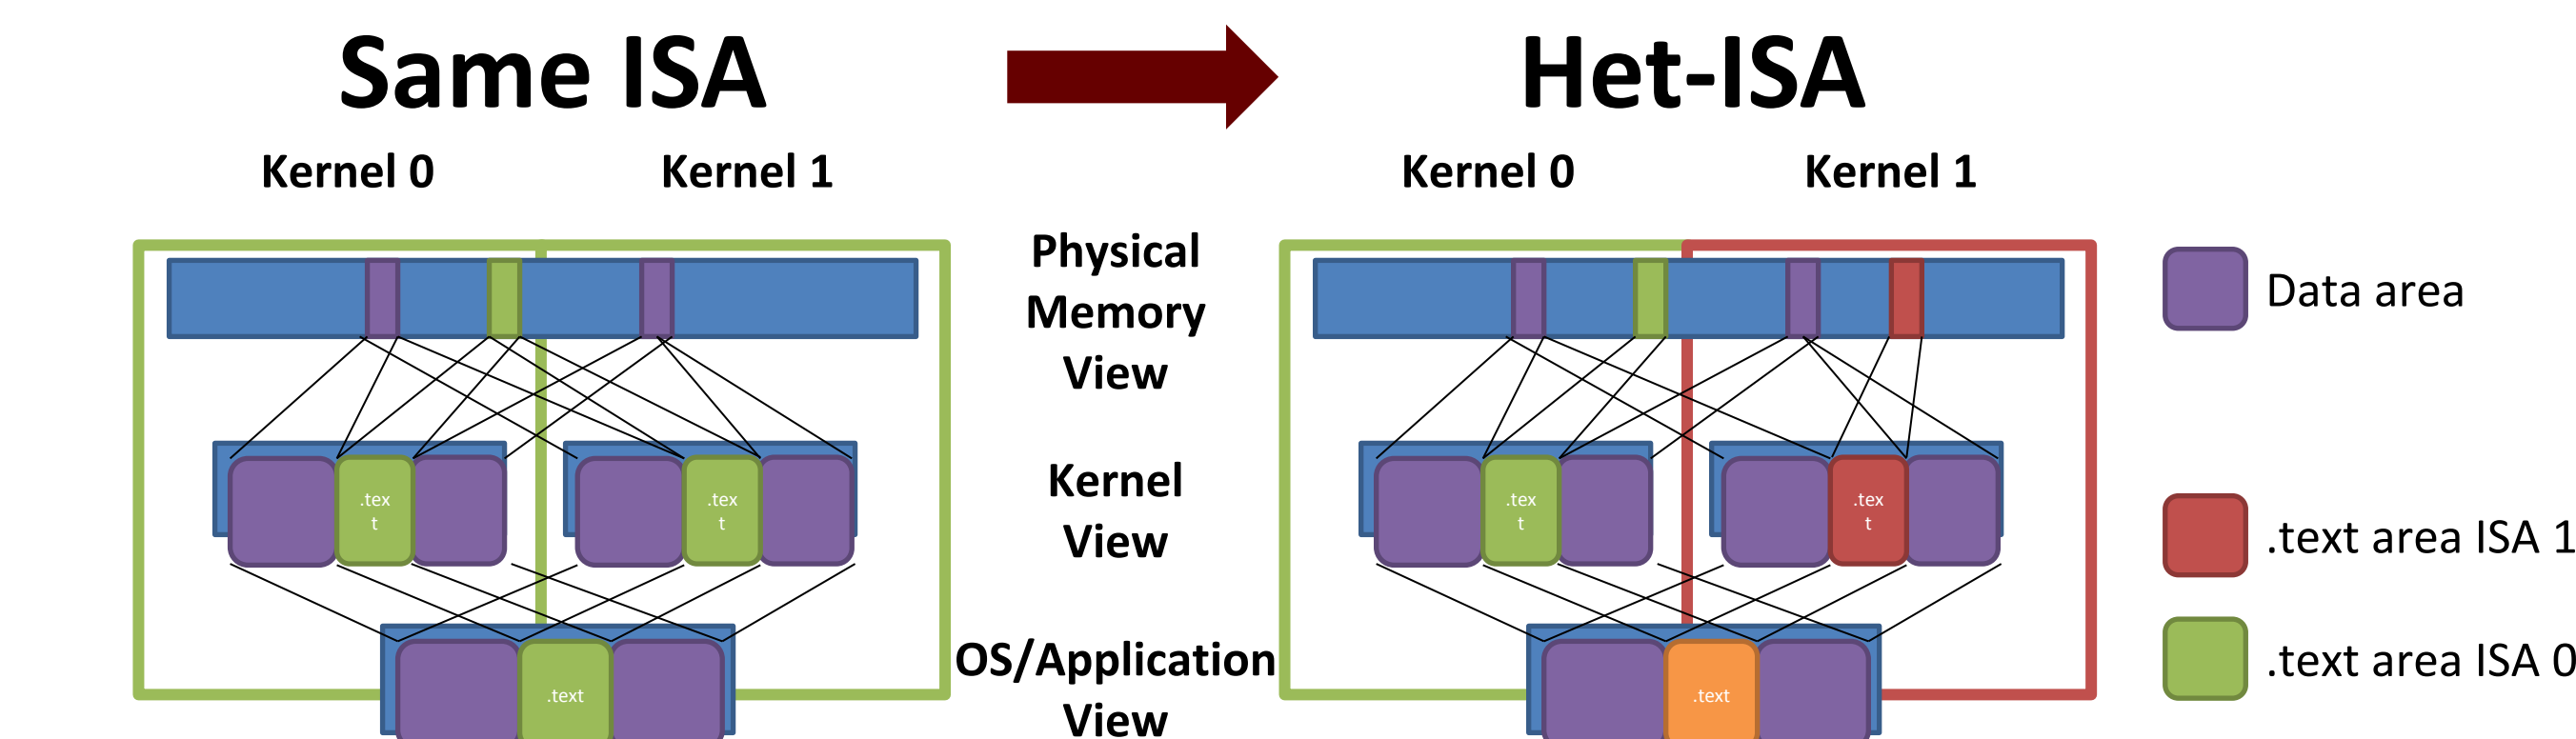
- ❖ A process can migrate to/from any kernel

Load balancing, power saving, memory/device locality, performance, expanded memory space, greater device count, etc.


We implemented this mechanism in Popcorn Linux, and tested on x86

Cache Coherent Shared Memory

- ❖ Replicated page table per kernel (same application)
- ❖ Sharing the same physical pages amongst kernels
- ❖ On-demand protocol to maintain a coherent page state mapping



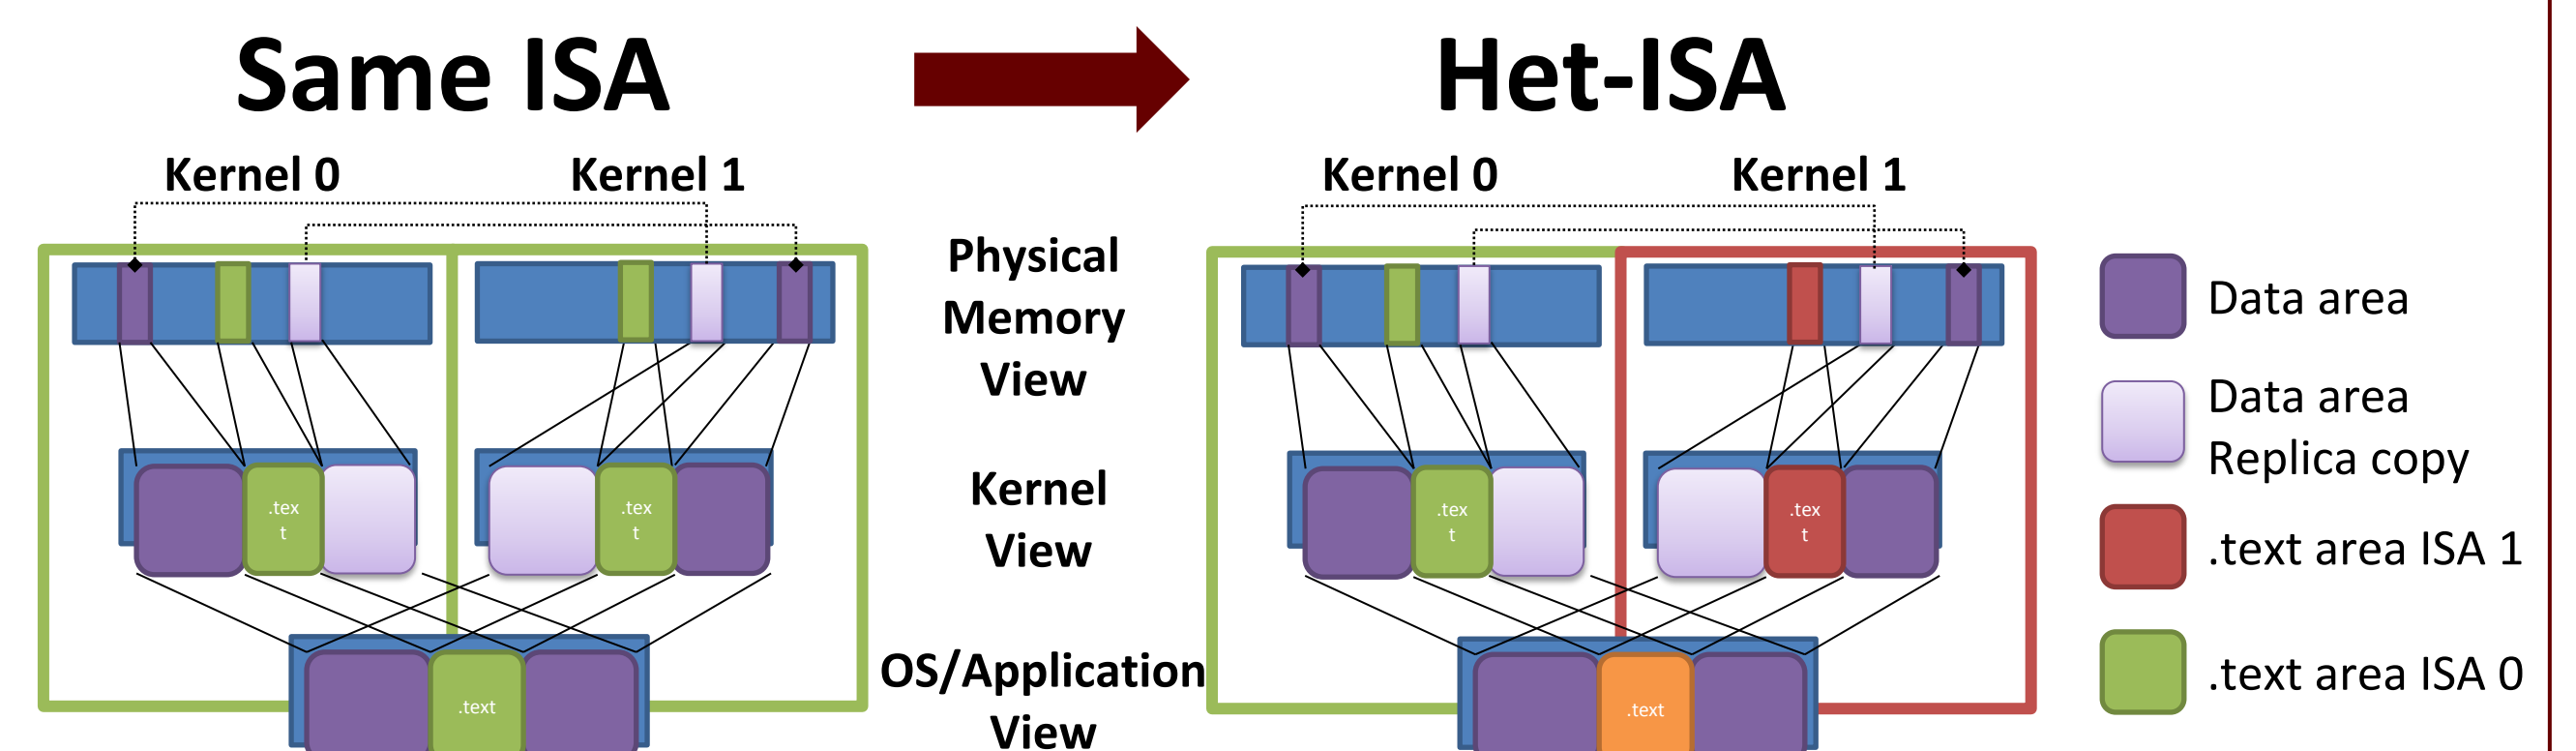
Initial Results (migration overhead)




- ❖ Bulk page migration
- ❖ On-demand page migration

Per ISA Memory

- ❖ Replicated page table per kernel (same application)
- ❖ Replicated page content on each kernel
- ❖ Page level protocol to guarantee a coherent memory view to the app



Page Coherency Protocol Overview



- ❖ We extended the MSI coherency protocol with *Unmapped* and *NotReplicated* states
- ❖ Linux pages are not present at initialization (*Unmapped*)
- ❖ On-demand pages become present and access any of the other states (N, M, S and I)

Non cache-coherent SHM architectures are handled by sharing the physical pages

while running a coherency protocol to synchronize the cache flushes on different kernels.

Conclusions

A traditional SMP OS, can be crafted to run on heterogeneous-ISA platforms. Task migration is a critical component.

We propose solutions for runtime task migration, in a replicated-kernel OS, on different het-ISA memory architectures.

www.popcornlinux.org

ACKNOWLEDGMENTS

This work is supported by US Office of Naval Research under Contract N00014-12-1-0880.

