

Unikernel: Library Operating System for the Cloud

Ruihan Wang

12/11/2018



Unikernels

“A unikernel is a specialized, single address space machine image constructed by using library operating systems. A developer selects, from a modular stack, the minimal set of libraries which correspond to the OS constructs required for their application to run. These libraries are then compiled with the application and configuration code to build sealed, fixed-purpose images (unikernels) which run directly on a hypervisor or hardware without an intervening OS such as Linux or Windows.”

— Wikipedia



Unikernels


“Unikernels are specialized, single-address-space machine images constructed by using library operating systems.”

— Unikernel.org



Unikernels

- **Specialized:**
A unikernel holds only one single application
- **Single-address Space:**
It doesn't separate user/kernel space inside the kernel
- **Library OS:**
Infrastructures are provided as libraries (and linked statically). Similar with compiling a user level program



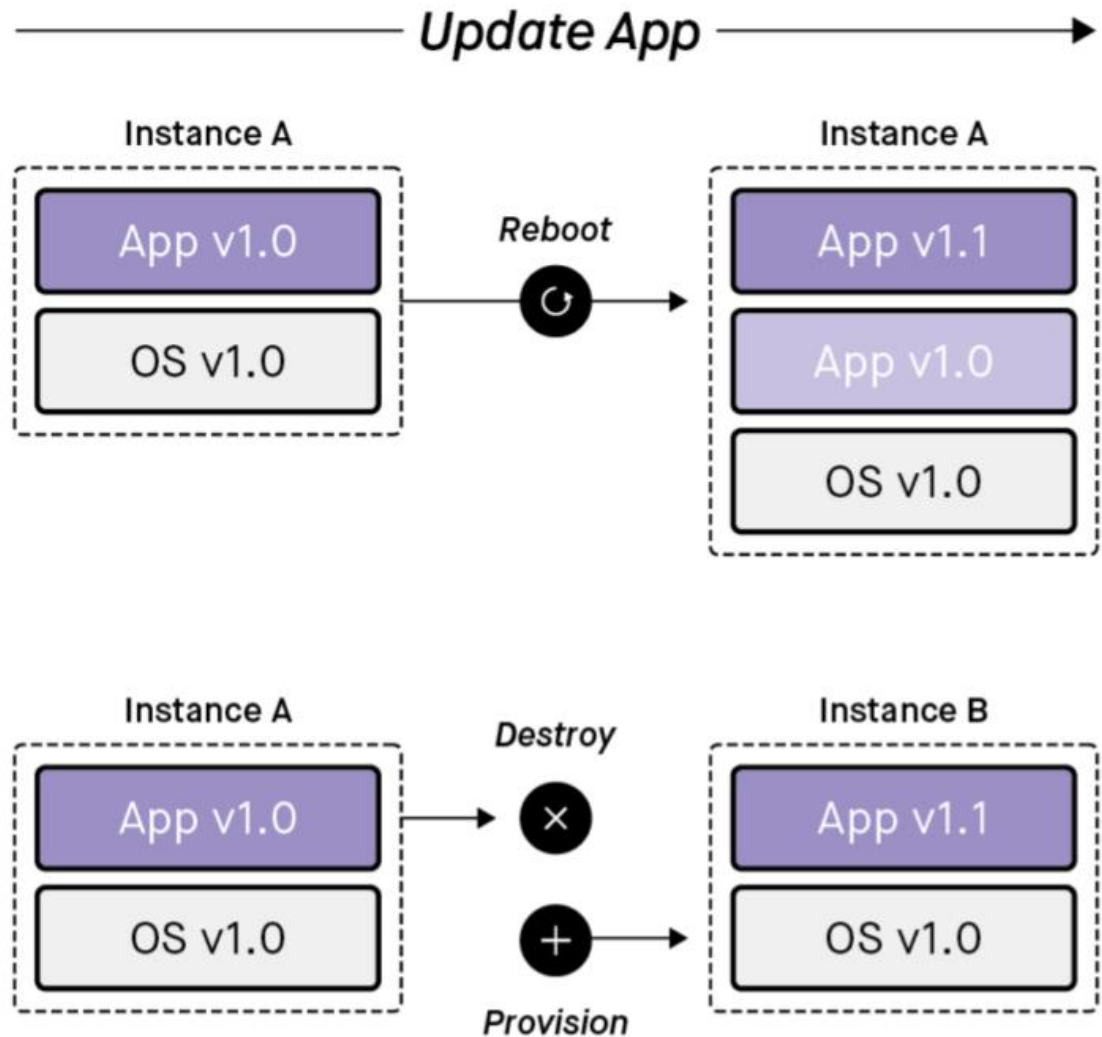
Why Unikernel matters?

- **Lightweight:** disk space cost \$\$!
- **Fast:** boot faster than a VM (but still slower than containers).
- **Isolation:** complete isolation from host.
- **Security:** no terminal, no ssh, nothing but application. Plus type-safe codes (more on this later).

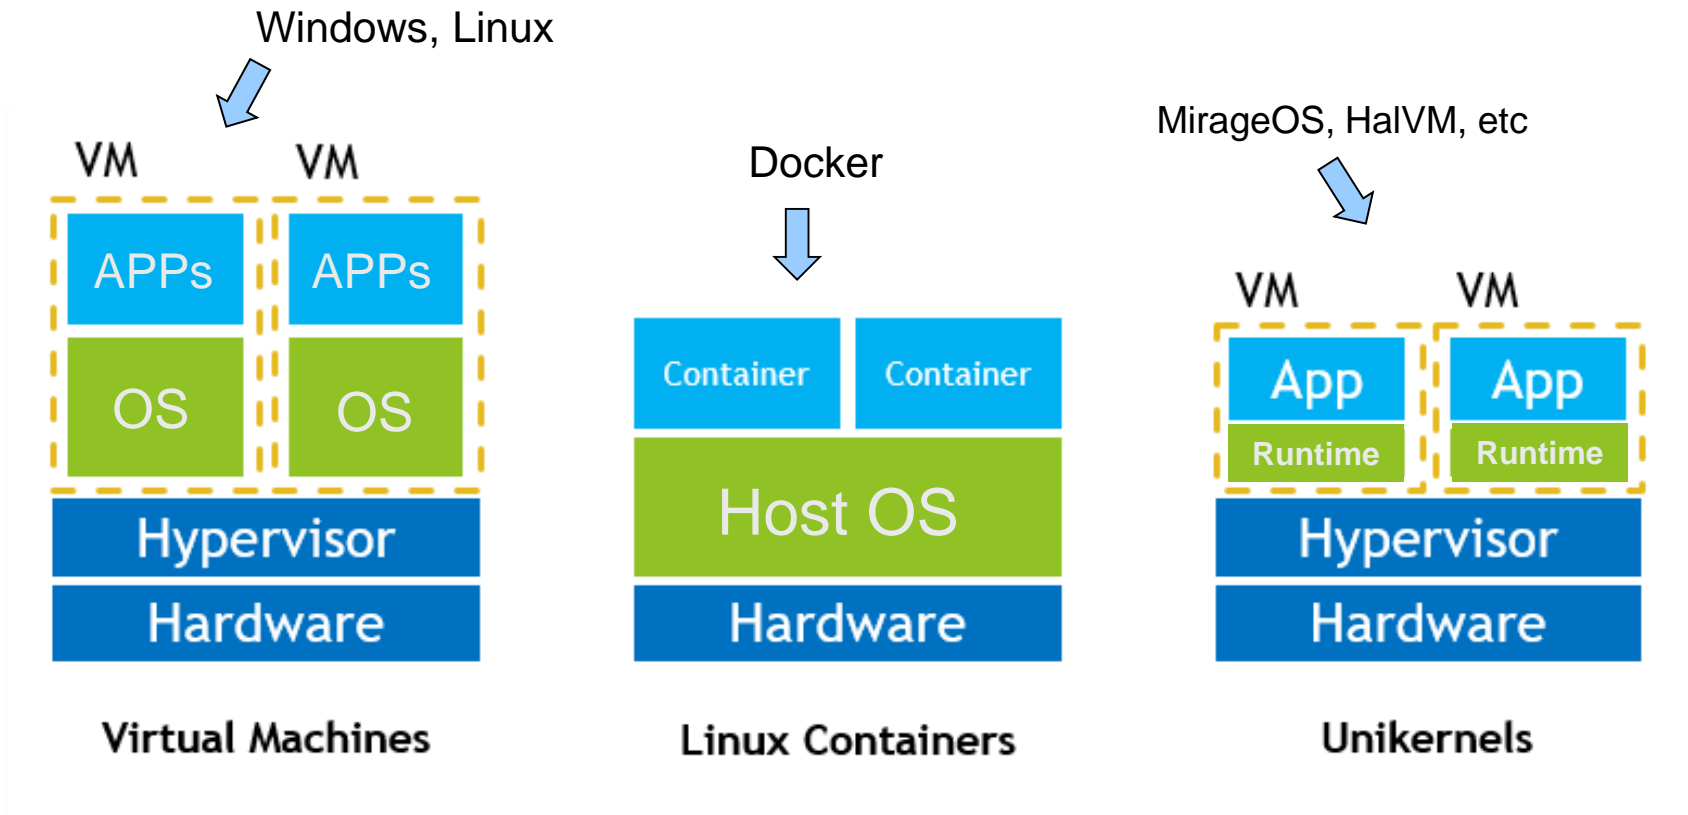


Immutable Infrastructures

Since in Unikernels everything was compiled statically, deploy or update an application is fairly simple: instead of creating a new version of dynamic library and reboot the application to adapt the change, one can destroy and create a new version of the whole instance. In other word, Unikernels are **immutable**.



VM vs. Container vs. Unikernel



Container **shares**, Unikernel **shrinks**.





They write the paper

+



OCaml



+



“Haskell Lightweight Virtual Machine”



I love Haskell

- HaLVM: <https://github.com/GaloisInc/HaLVM>
- MirageOS: <https://mirage.io/>




Being Pure Functional

- There are many Unikernels written in different programming languages... you can even find one written in Javascript (runtime.js)!
- High-level language features (first class functions, module system, threads, GC, etc)
- Property based test
- Type-safety!
- Pure (explicit side effects) Monad!



Monad Wat?



Managed Side Effects

loadFromDb :: (MonadError e m, MonadReader r m, AsDbError e, HasDbConfig r, MonadIO m) => m MyData

loadFromDb = ...

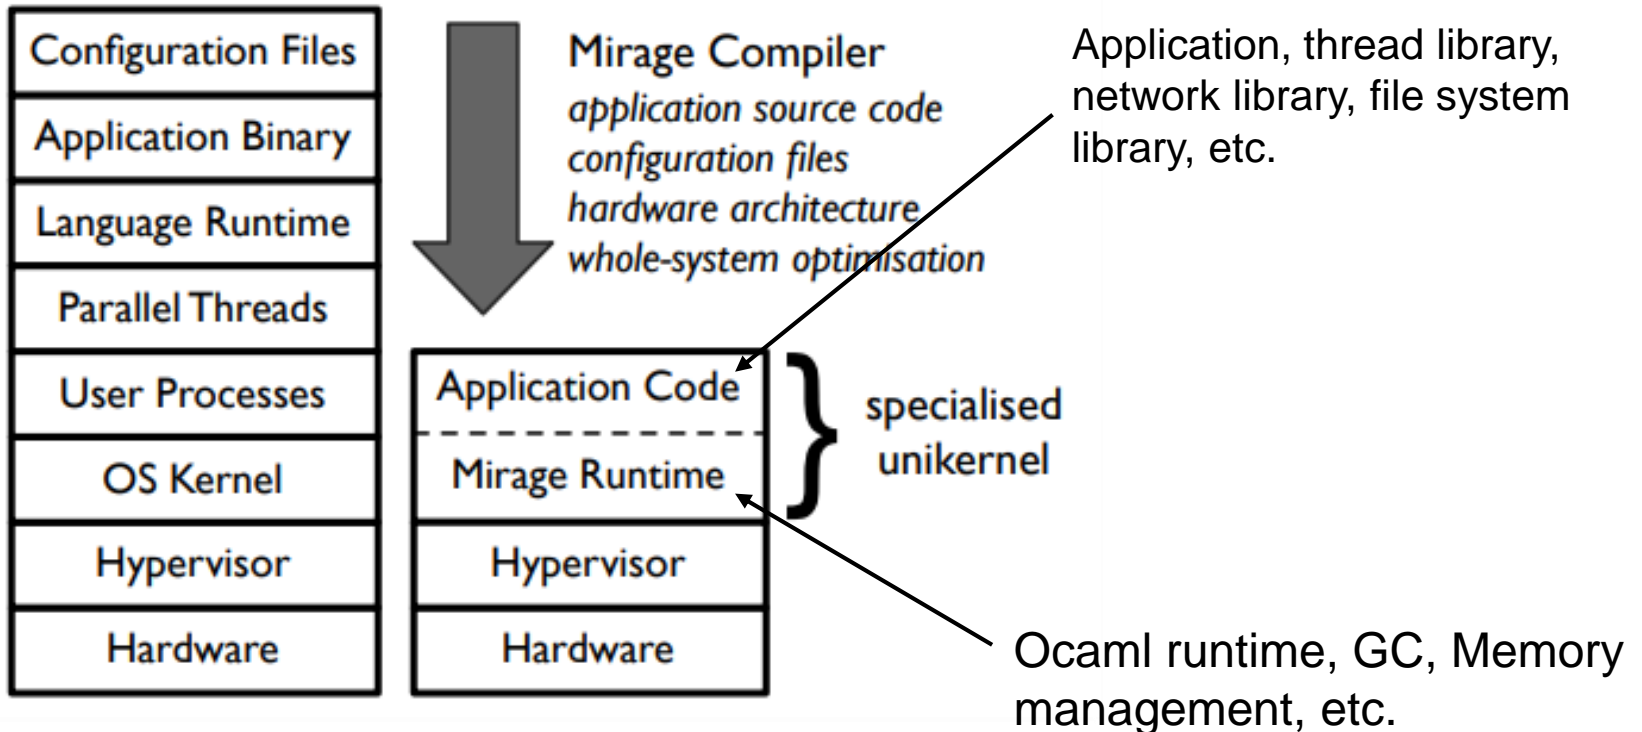
sendOverNet :: (MonadError e m, MonadReader r m, AsNetworkError e, HasNetworkConfig r, MonadIO m) => MyData -> m ()

sendOverNet = ...

loadAndSend = loadFromDb >>= sendOverNet



Architecture of Unikernels

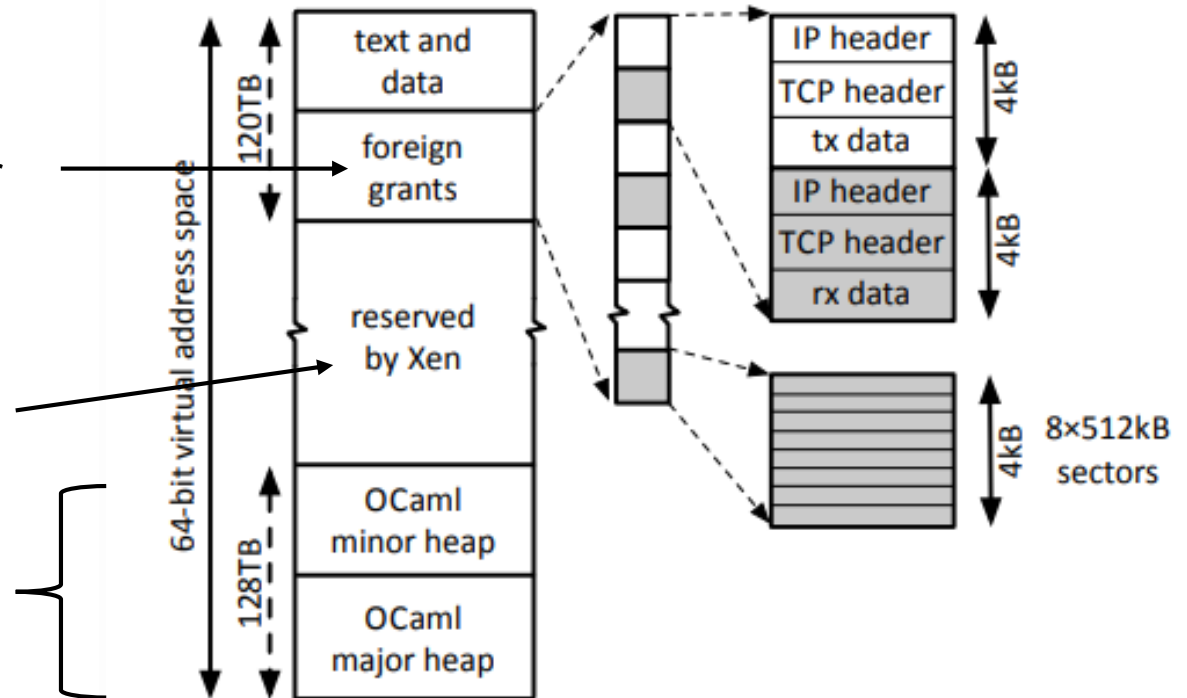


Memory Layout

External IO pages for communicating with other VMs (can be directly accessed by other VMs)

Store the physical-to-machine page table, etc.

Ocaml heaps (generational GC)



Evaluation




Figure 6: Boot time using an asynchronous Xen toolstack.

?

	Binary size (MB)	
	Standard	Dead code elimination
DNS	0.449	0.184
Web Server	0.673	0.172
OpenFlow switch	0.393	0.164
OpenFlow controller	0.392	0.168



Typical Linux image size is ~400MB



Evaluation (Cont'd)

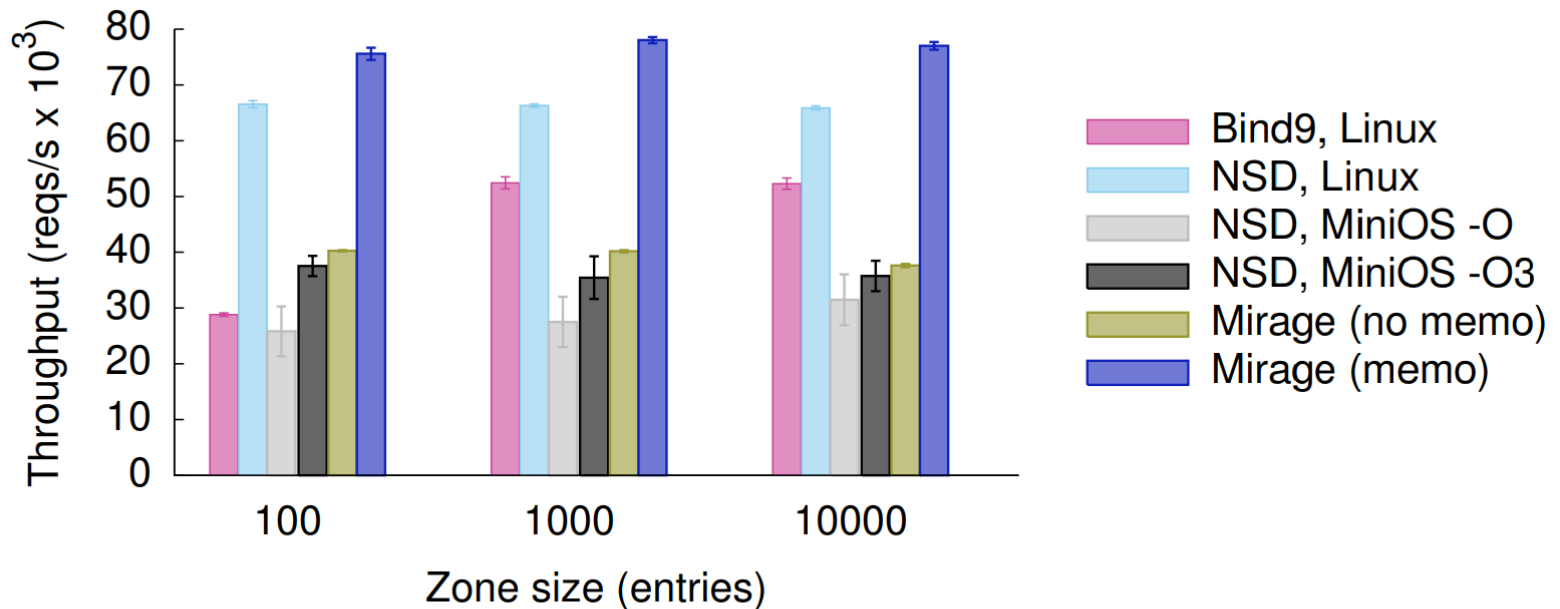


Figure 10: DNS performance with increasing zone size.



Downsides of Unikernel

- Developing Unikernel is painful
 - Most libraries won't work (and you need to rewrite them all!)
 - Manually initialize hardware when startup
 - Lack of debugging tools
- Only works on top of hypervisors
- Static resource allocation



References & Question

- Anil Madhavapeddy, Richard Mortier, Charalampos Rotsos, David Scott, Balraj Singh, Thomas Gazagnaire, Steven Smith, Steven Hand, and Jon Crowcroft. 2013. Unikernels: library operating systems for the cloud. SIGPLAN Not. 48, 4 (March 2013), 461-472.
- Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. 2003. Xen and the art of virtualization. SIGOPS Oper. Syst. Rev. 37, 5 (October 2003), 164-177.
- <https://github.com/cetic/unikernels>
- <https://github.com/GaloisInc/HaLVM>
- <https://github.com/gwils/next-level-mtl-with-classy-optics>

