

OpenBSD/sparc64

Mark Kettenis
kettenis@openbsd.org

2 December 2007



OpenBSD/sparc64

- OpenBSD port to the Sun UltraSPARC (sun4u) architecture
- Started by Jason Wright and Henric Jungheim
Significant contributions from Artur Grabowski, Mark Kettenis
- Based on NetBSD/sparc64



SPARC-V9

Scalable **P**rocessor **ARC**hitecture
Based on 32-bit SPARC-V8

- 32-bit instructions
- 8 global registers + 24 windowed registers
- 32 double precision FP registers
- 128-bit quad FP support
- Big-Endian or Little-Endian
- Unprivileged and Privileged mode
- Multiple levels of (fast) traps



SPARC-V9

Implementation defined:

- MMU
- Interrupts
- Additional traps
- Number of register windows

Implementatuons:

- Sun UltraSPARC
- Fujitsu/HAL SPARC64



UltraSPARC/sun4u

Architecture

- Superscalar
- Big-Endian
- MMU with context
- 8 register windows
- No hardware 128-bit quad FP support
- VIS multimedia instructions
- OpenFirmware



UltraSPARC/sun4u

CPUs and Machines

- **UltraSPARC** Ultra 1, Enterprise 4000
- **UltraSPARC-II** Ultra 30, Enterprise 450
- **UltraSPARC-IIi/IIe** Ultra 10, Fire V120
- **UltraSPARC-III/III+** Blade 1000, Fire V880
- **UltraSPARC-IIIi** Blade 1500, Fire V440/V445
- **UltraSPARC-IV/IV+** Fire V890



Busses

- **SBus** sbus(4)
- **UPA** upa(4)
- **PCI** psycho(4), schizo(4)
- **PCIe** pyro(4)

IOMMU: 32-bit devices can address >4G memory



OpenBSD/sparc64

- Fully 64-bit: no 32-bit compat code
- Runs on almost any UltraSPARC-I/II/III machine
- Supports almost all onboard devices; exceptions:
 - VGA frame buffers
 - CD-ROM



Driver Reference Architecture

“Your driver isn’t finished until it runs on sparc64!”

- 64-bit
- Big-Endian
- Strict Alignment
- IOMMU
- Built-in Diagnostics



Security Features

- W ^ X (in kernel too)
- IOMMU
- StackGhost



StackGhost

Buffer overflows modify return address on stack

Register Window Architecture helps

StackGhost: Mike Frantzen, Mike Shuey (for sparc)

“Encrypt” return address spilled on the stack

“Decrypt” upon filling register window

Encryption: XOR with random per-process cookie

Decryption of overwritten address results in (unaligned) random address.



New features

- StackGhost
- UltraSPARC-III support
- Sensors
- Lots of little platform-specific drivers
- Cassini+ network cards
- GENERIC.MP



UltraSPARC-III

- MMU is different from UltraSPARC-I/II
- New PCI hostbridge: Schizo
- Insufficient documentation
- CPU bugs?

Most work done by Jason, Art and Henric

Running with caches disable after 2006 Hackathon

Fully functional in OpenBSD 4.1



GENERIC.MP

Architecture supports SMP very well:

- Synchronization instructions
- Inter-Processor Interrupts
- Cache Coherent
- Firmware support

New code:

- Mutex implementation
- TLB flushes
- FPU context saving

Builds your sparc64 snapshots now!



Future

- Make it faster!
- UltraSPARC/sun4v support
- OpenBSD on UltraSPARC-IV machines
- OpenBSD on Sun Enterprise 10000?
- Support for Fujitsu SPARC64 CPUs?



UltraSPARC/sun4v

- Virtualization
- Hyperprivileged mode: Hypervisor (BSD-licensed)
- Abstraction layer for MMU, hostbridge, interrupts

CPUs and Machines:

- **UltraSPARC T1** Sun Fire T2000
- **UltraSPARC T2** Sun Fire T5220



Bonus: ExpressCard support

Hotplug PCIe support: Surprise Removal

Interrupt handling for bridges

Removal: PCI drivers need to be made detachable

Insertion: assign PCI resources (normally done by BIOS)

In -current: msk(4), sili(4)

Done: puc(4), behind PCIe-PCI bridge

