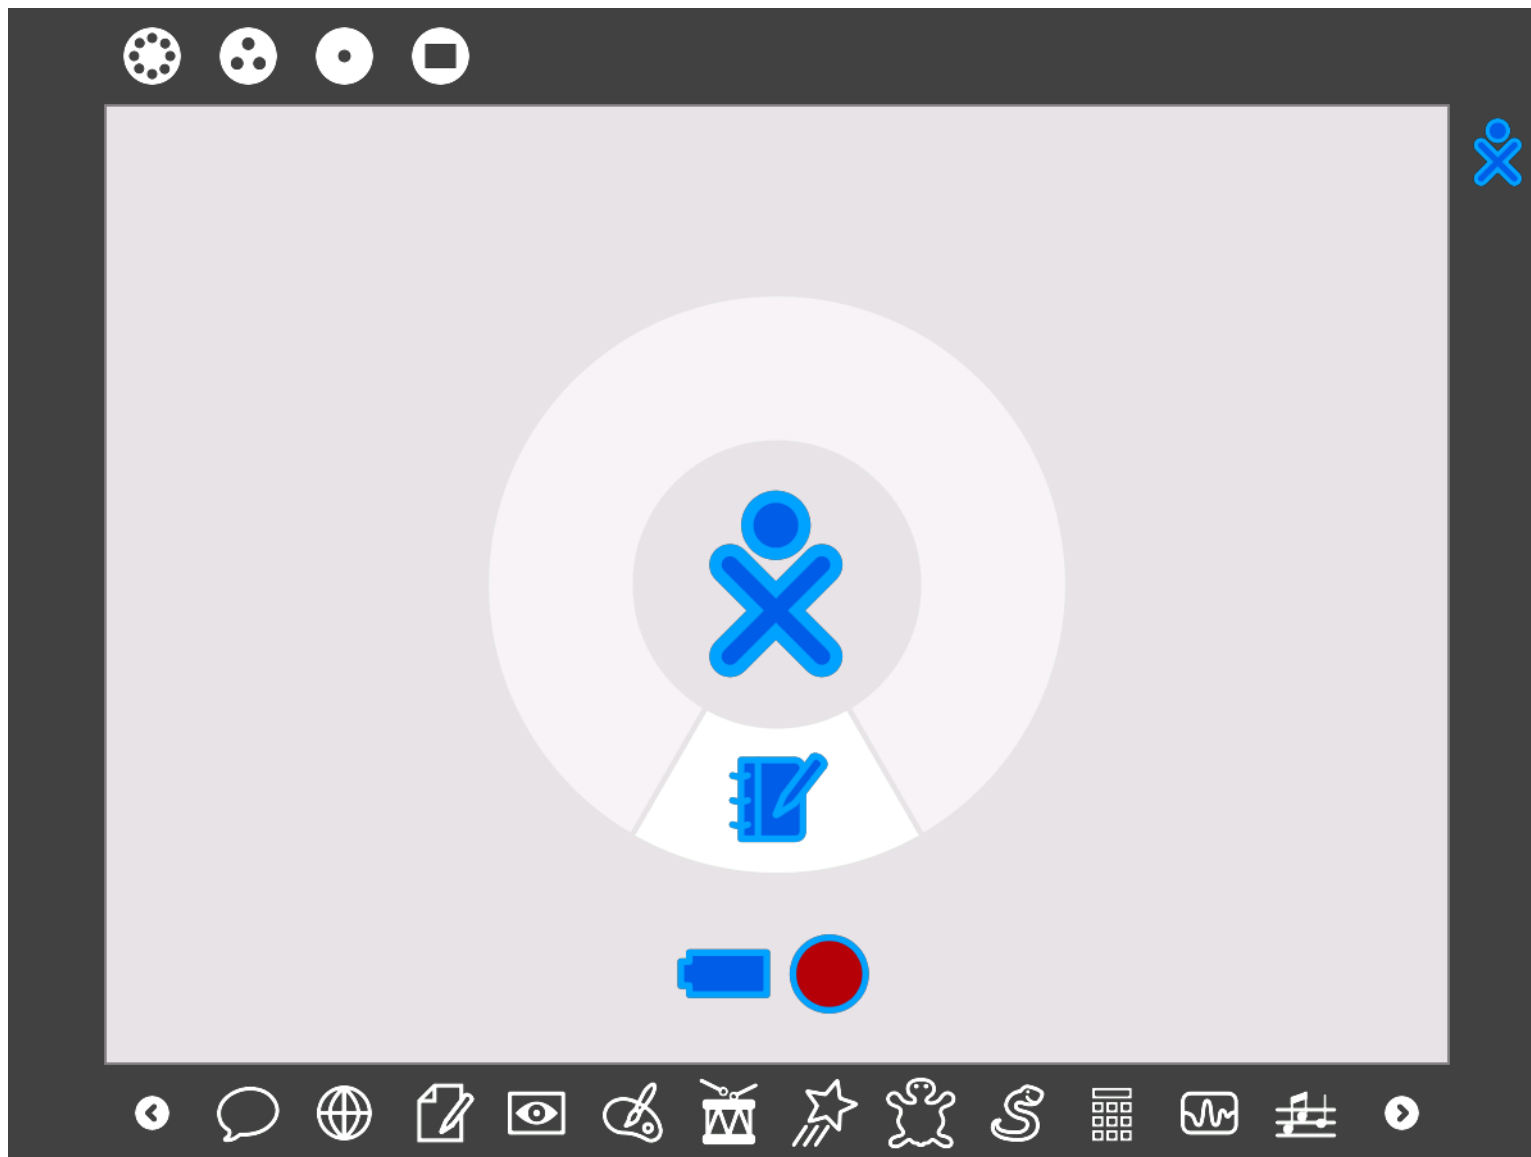




Under the XO Hood


A slightly different look at the OLPC

Jordan Crouse
Linux System Software Developer
April 10, 2008



<http://wiki.laptop.org/go/Image:Sugar.png>






http://en.wikipedia.org/wiki/Image:Nicholas_Negroponte.jpg





[http://wiki.laptop.org/go/Image:Green and_white_machine.jpg](http://wiki.laptop.org/go/Image:Green_and_white_machine.jpg)







·“Tux” logo by Larry Ewing




AMD Geode LX CPU




AMD CS5536 Southbridge



Wireless Dongle



Embedded Controller



BIOS? We don't need no stinkin' BIOS!

- What we don't need
 - Legacy calls
 - ACPI
 - VGA
 - 15 or 20 second boot times
 - Costly proprietary (and closed) code
- But surely, we need something, right?
 - Setup the hardware (memory!)
 - Load and run a Linux kernel
 - Debugging
 - People friendly features



Requiem for a BIOS

- First, there was Coreboot (LinuxBIOS)
 - Set up the hardware and then loaded a payload
 - Very fast
 - First payload was a “Linux as Bootloader” solution
- Then, there was OpenFirmware (OFW)
 - Unexpectedly and fortuitously released to the public by Sun Microsystems
 - OLPC acquired the services of Mitch Bradley, the “father of OFW”
 - Originally ran with Coreboot as a payload
 - Replaced Coreboot starting with the B-3 board revision



March of the Penguins

- Areas of interest for the XO platform
 - Display / DCON
 - Camera
 - NAND Flash
 - SD Card
 - USB Wireless
 - Battery
 - Power Management



Framebuffer Driver

- Provides a text console in lieu of VGA support
 - Never underestimate the power of a good dmesg!
- Can shut down the GPU when the DCON is frozen, saving power
- Understands how to save and restore the entire graphics context on suspend
 - Useful for userspace applications that don't know that a suspend is happening



Speaking of Userspace.....

- Why use X11?
 - Acceleration, acceleration, acceleration
 - Its the sad truth - all the good GUI toolkits are X based
 - It is familiar and well supported
- The Geode video driver
 - Full 2D acceleration (including hardware accelerated compositing!)
 - Hardware accelerated rotation
 - Xv support (more on this later)



The DCON

- The DCON is the most special of all the special sauce on the XO
 - Has on-board memory for “freezing” the frame
 - Controls the backlight, color swizzling and antialiasing features of the LCD
- Sits in the data stream between the CPU and the and the TFT
 - The Geode doesn't know any better; it thinks its talking directly to a panel
 - The graphics source never changes always 1200x900 with 16 bit color
- DCON control is provided through a side band SMBus connection




Dance, Dance, DCON!

- The DCON driver controls all of the DCON features
 - Backlight level
 - Monochrome or full color
 - Antialiasing
 - Source; CPU (not frozen) or DCON (frozen)
- Full userland control
 - `/sys/devices/platform/dcon`
 - `/sys/class/backlight`



The Camera

- The camera controller is located on the CaFE chip
- Driver written by our own Jon Corbet
- Remember the Xv driver I mentioned before?
 - Uses the hardware to do YUV to RGB colorspace conversion
 - 90% of the CPU utilization in software-only decoding is spent in color space conversion
 - Also handles scaling



Various other gadgets...

- NAND Flash Controller
 - Written by David Woodhouse
- SDHCI (Secure Digital Host Controller Interface)
 - Uses standard class driver by Pierre Ossman
- Battery Interface
 - Written by David Woodhouse
- Touchpad input driver
 - Written and re-written by a wide number of folks



USB Wireless

- The concept: A USB attached 802.11N wireless controller with an onboard processor for offloading networking and mesh support
- The setup: The dongle is powered externally (**not** through USB), and the processor remains powered while the system is suspended. A side band signal to the embedded controller serves as a wakeup signal when a packet arrives. A small microcontroller serves as the bridge between the USB interface and the on-dongle processor, making the dongle appear to the system as a USB.



USB Wireless (continued)

- That doesn't sound so hard...
 - USB doesn't always behave under the best of times, let alone when suspending and resuming
 - A USB slave device that is externally powered and doesn't lose its state during suspend and that can actually wake up the processor again is probably completely unique in the world of USB gadgets
 - The firmware turned out to be a fully qualified operating system in its own right - which was just as hard to debug and fix as the regular kernel



(Less) power to the people

- The primary goal of the OLPC project is worldwide education
 - This just in: The world is really big!
- If the primary goal is worldwide acceptance, then the biggest blocker is efficient power management
- The math is simple:

$$N \propto 1/P$$



How the hardware helps...

- At the core...
 - AMD Geode LX700 consumes 1.3W typical
- And the supporting cast...
 - Power supply design is tolerant of different voltages and power qualities
 - LCD uses 1/10 the power of a standard TFT panel
 - The entire board is optimized to save power
- The end result?
 - 7W typical consumption, up to about 9W with everything turned on and cranked
 - That's 2 - 3 hours on a 22 watt hour battery



Turn out that light!

- Remember how your mother always told you to turn off the lights when you weren't in the room?
 - The best way to save power is to not consume it in the first place
- Our goal is to turn off as many components as we can when we're not using them
- Our efforts fall into two classes
 - Runtime power management
 - Suspending the processor and components



Runtime Power Management

- The Geode design helps with something called Automated Hardware Clock Gating (AHCG)
 - The clock trees turn themselves off when not in use
 - The software must make an effort not to use things needlessly
- But what about the things you can't avoid using?
 - The kernel drivers need to provide knobs for turning things off
 - Userspace needs to take an active role in understanding how the system works and how to turn the knobs for full effect
 - Example: video display is fetching 2 megabytes of data 50 times a second
 - Use the DCON to freeze the screen and turn off the GPU



- The processor has a special halt (x86 opcode *hlt*) state that it can go into if nothing is happening
 - Stops clocking instructions until an interrupt is received
 - Default Linux idle behavior is to be in the halt state as much as possible
- On an older Linux kernel the clock ticks at 100Hz if we need it or not
 - Clearly inefficient; ticks we don't need rob us from precious time spent in halt
 - The solution is a tickless system; only tick when we need it
- On modern systems, tickless systems use the HPET or local APIC to tick the clock. The Geode is relegated to the legacy PIT. Or is it?
 - The southbridge has special timers we can use in place of the legacy PIT.
 - The system went from 100-110 interrupts per second down to 20

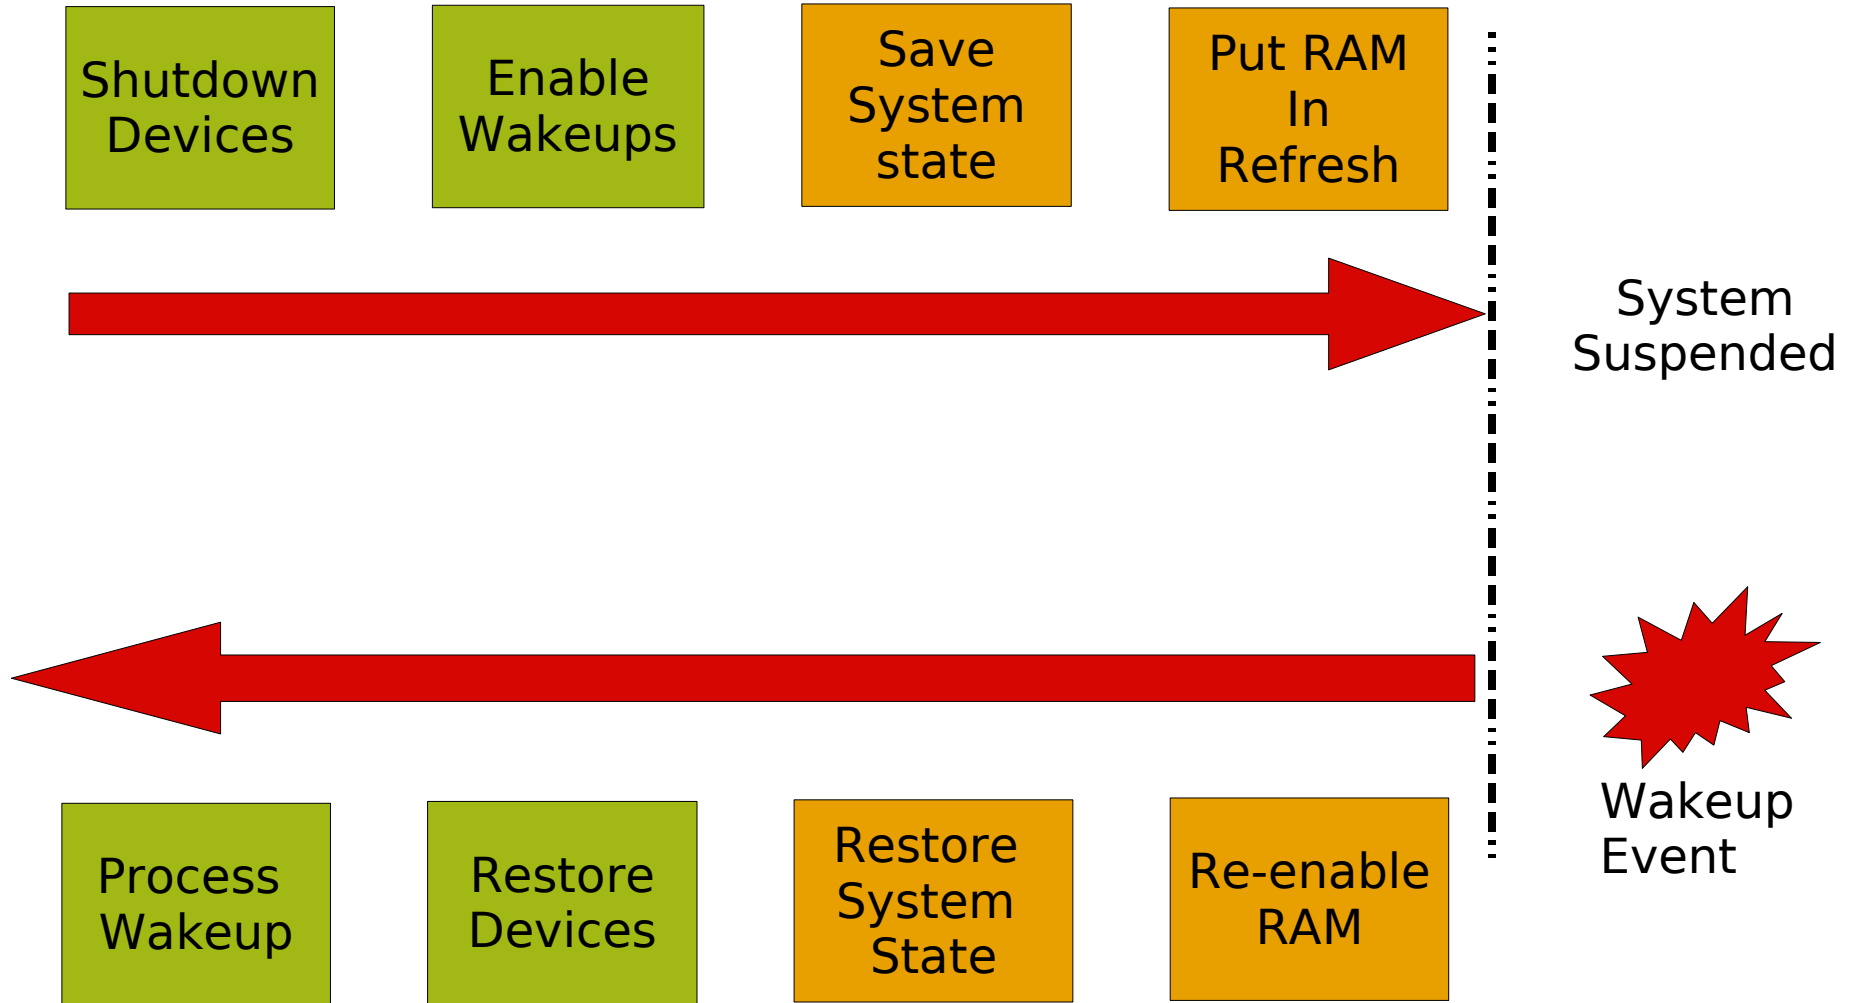


Suspend to RAM

- Suspend to RAM (STR) is the process of turning off most devices on the system, putting the memory into self-refresh and physically turning off the processor.
 - A few devices remain powered to handle wakeup events
 - These still consume power, but not nearly as much as an active system



How Suspend to RAM works...



Wake up!

- Wakeup events are any external event that causes the processor to resume
 - power button
 - Lid event
 - real time clock alarm
 - mouse
 - keyboard
 - network
 - battery



Breaking new ground


- On a traditional x86 system suspend to RAM support is handled through ACPI
- The XO platform has no ACPI!
- Challenges of a custom implementation
 - ACPI performs a lot of business “behind the scenes”; we need to reimplement that behavior in the kernel
 - It is difficult to negotiate the interface between the firmware and the kernel
 - We must handle processing wakeup events ourselves (and we have a lot of them!)



Speedy Resume

- Because suspend to RAM saves so much power, it is the primary power saving method for the XO platform
- The faster the suspend/resume process then the more occasions we'll have to use it
 - The new 'cpuidle' interface in the kernel knows how long it will be until the next event that the kernel needs to process
- The goal for the XO resume is 250ms
 - With tickless enabled, the time between events on an an idle system can easily be 250ms between events
- Challenges
 - USB Wireless





IF YOU KEEP SAYING "BEAR WITH ME FOR A MOMENT", PEOPLE TAKE A WHILE TO FIGURE OUT THAT YOU'RE JUST SHOWING THEM RANDOM SLIDES.



Questions?

Trademark Attribution

AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners.

©2008 Advanced Micro Devices, Inc. All rights reserved.