

The Kinect Sensor

It seems Microsoft may have got it right for once. The Kinect sensor for their xBox is, to use a word from the appropriate generation, awesome. With it, the xBox opens up a whole different world of game experience – no joysticks to waggle, no Wii sensors to accidentally throw at the screen, and (as all Wii users do once they know about it) you cannot cheat by using small, sharp movements to simulate large ones.

But there is more. The interface is USB and, although MS have not (to my knowledge) published the interface specification, it can plug straight into a Linux (or Mac or Windows) box.

A group called “OpenKinect” (openkinect.org) have reverse-engineered the interface and published an open source library which supports using the sensor; it is available for Linux, Mac OS X and Windows, and supports languages such as C, C++, Python, Java, JavaScript, and even Lisp. A library called `libfreenect` and several demonstration applications are readily available.

Let us start with the sensor.

It has its own little computer inside it, and firmware updates are possible.

It has position sensors within it so that it knows when it is level, and has a motor that can tilt it up or down by up to 31 degrees (in half-degree steps). The unusual thing is that you tell it the angle you want from the horizontal, and it automatically compensates for the base not being on a level surface.

It has a small colour video camera (15 or 30 frames per second, 640 x 480 pixels). Nothing surprising about that. Many xBox games use it to take pictures of the players (and upload them to Microsoft??)

It also has an infra-red camera and light source. This can be used to take infra-red pictures (it has 10 bits of resolution, but the noise level varies enormously from pixel to pixel) but is more often used in conjunction with the light source as a depth sensor. This means that you get a picture (again, about 640 x 480 pixels) whose pixel values represent the distance of the object from the sensor. This is how the xBox is able to determine the body pose of the player(s) – it can assume that pixels closer than the general background are part of the player, and separate the player from the background much more accurately than just with the colour image – a person can easily blend in to the background if they are wearing a jumper of a similar colour, for example, but the distance image does not have this problem.

Microsoft claim that the infrared cam is 320x240 and the RGB cam is 640x480.

Even more interesting is that it has four microphones! (`libfreenect` cannot yet support these). The point about having four is that it is possible to “beamform” the sound (as is done in a submarine's sonar) so that it listens in a particular direction. For example, if there were three people in front of it (they would need to be about 2M apart) then you could identify the positions (using the depth sensor) of the heads, form four beams aimed at them, and determine which one is speaking because the signal would be strongest for that particular beam.

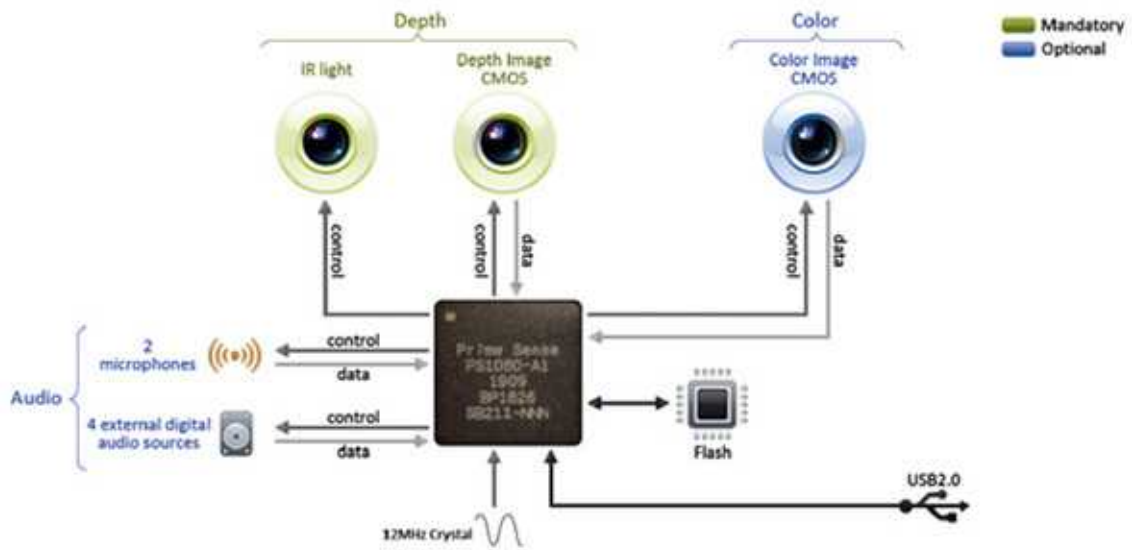
The following photos are from www.ifixit.com.













It has two stereo Analogue to Digital converters with preamplifiers, a flash memory, a 512Mbit Dram



Some more flash memory (1Mb)



An accelerometer used for inclination and tilt sensing



These are three of the microphones



And this is the disassembled sensor.

The software was easy to install. For Ubuntu Lucid and Maverick, adding their repository as a source allows you to apt-get install it; Deb packages are available; and for anything else you can install the source and build. The steps to do each of these are spelled out on their Getting_Started Wiki page. (http://openkinect.org/wiki/Getting_Started). Their web site suggests that most of their knowledge of the sensor is from reverse-engineering, rather than a published specification.

Once installed, there are several utilities that allow you to drive the device:

`freenect-glfwview` is a basic tool that lets you run the motors and shows a dynamic sensor display. You can also set the led mode (I am not sure if that does anything useful). It lists the commands on startup.



`freenect-glfwview` performs some basic 3D processing. It does not offer much help, but 'w' and 's' zoom in and out, the mouse can be used to tilt and rotate; 'c' turns off colour.



There are lots of sample videos available on the net, including:
<http://www.youtube.com/watch?v=HNkbG3KsY84>

Now lets have some fun with it.