How it Works

An LED with a narrow field of view is used to illuminate the filament. A shadow of the filament is projected onto a linear sensor array (400 pixels/inch). An 8 bit processor is used to detect the edges of the shadow image and calculate the width. The image sample and width measurement occurs every .01 sec. The sensor board outputs a voltage that is related to the measured width. The sensor board firmware produces two types of output, ratiometric and absolute. Ratiometric output is used for interfacing to the controllers of 3D printers or Filament extruders. Absolute output is used when the user wants to read the filament diameter directly with a voltmeter.

Ratiometric output:
The output voltage from the sensor board is interpreted as a ratio. An output matching the full supply voltage indicates a 5.00mm measurement. This way variations in supply will still provide accurate measurement. This ensures that if the supply voltage is not exactly 5.00, or varies over time, the attached controller's A to D input which is also ratiometric will get an accurate reading from the sensor board. The controller and sensor board need to be powered by the same 5v supply for this to work.

Absolute output:

The sensor board can output the absolute filament width as a voltage (i.e. 1.0v is 1.0 mm) independent of variations of supply voltage (4.9v to 5.1v). This way, the user can read the filament diameter directly with a voltmeter. The accuracy is slightly reduced in this mode.

The sensor has some variability depending on how close the filament is to the linear sensor array, because the size of the shadow will change with distance. Maintaining alignment while the filament is travelling is critical and the case is designed to keep the filament close to the sensor. There is some wiggle room, maybe .4 mm. The illuminating LED is about 45mm away, so a .4 mm movement causes roughly a 1% change in shadow size (45-.4)/45, so +.03mm deviation. One way to reduce this is to move the LED even further away, or have some sort of spring-loaded roller that holds the filament in place. In practice I have not seen much variation when I wiggle the filament in the sensor. I wonder if moving the filament away from the sensor also diffuses the shadow edges, thereby effectively compensating for an increase in shadow size.

The sensor firmware uses a subpixel image processing approach to improves apon a more basic approach. The basic approach is to essentially count the number of 'dark' pixels in the array to estimate the shadow size (proportional to the filament diameter). This requires thresholding the pixel values to come up with a definition of 'dark'. Its a little more complicated, because you want to search for an edge - the transition between a 'illuminated' pixel and a 'dark' pixel. And then search for the other edge dark to illuminated.
The subpixel approach takes the basic approach (count of the dark pixels) and adds corrections for each edge. Since the filament edge shadow does not fall exactly on a pixel boundary, it requires looking at the 3 pixels on the edge (definitely dark, grey, definitely illuminated). The algorithm uses a quadratic fit of the 3 pixel values to estimate a more precise edge. This is done for both edges, yielding a subpixel correction amount for each edge (something between 0-1 pixel width). Then add up the pixels counted

in the basic approach and the subpixel amounts from each edge and you get a better width approximation.

Here is a link to the paper I used for the math:
dev.ipol.im/~morel/Dossier_MVA_2011_Cours_Transparents_Documents/2011_Cours1_Document1_1995-devernay--a-non-maxima-suppression-method-for-edge-detection-with-sub-pixel-accuracy.pdf

Pictures in the paper do a good job showing whats happening.