

Automatic Motion Tracking with the ADV611

By Hakan Çivi and Alex Zatsman

OVERVIEW

Analog Devices has built and tested a two-board motion tracking system. This applications note explains how the two-board motion tracking demo system works. The system, particularly relevant for surveillance purposes, is based on the two previous demo boards, namely the Videopipe and CCTVpipe, along with new firmware in EPROMs. When the two boards are driven by the same video input source, the Videopipe board performs the motion tracking by computing the quality box coordinates, and passes this information to the CCTVpipe. The CCTVpipe, in turn, positions the quality box, as computed, on the compressed/decompressed input video. As you will recall, the CCTVpipe allows quality box size and location control through a mouse. The idea, instead, is to automatically control the quality box so it closely follows the moving regions in a scene. To accomplish this goal, a CCTVpipe board alone would be sufficient. To ease the software development the detection of motion in the video field was done in the motion tracker board, and display of the detected motion was assigned to a fairly standard CCTVpipe board. It is quite practical to do motion tracking in a single board.

Using this document and the referred algorithms you can do the following:

- Run the demo system and evaluate it.
- Learn about the underlying algorithms and customize them for your purposes.
- Design automatic motion tracking systems as suited to your applications.

The software source code and hardware schematics mentioned in this note are available on the Analog Devices web site at www.analog.com.

HOW TO RUN THE SYSTEM?

The system (seen in Figure 1) includes a Videopipe board, a CCTVpipe board, the associated new EPROMs, and a flat cable for communicating the motion tracking

information. The cable should be twisted and plugged into JP17 on each board, so that Pin 1 mates with Pin 7, Pin 2 mates with Pin 6 and so on. For a meaningful demo, the boards should be driven by the same video input source, typically a video camera. A simple tee connection works well. Finally, by connecting the output of the CCTVpipe to a video display, the performance of the motion tracking system can be evaluated. Also note that, except the quality box control and motion detection (they are now controlled externally by the motion tracker) functions, the CCTVpipe works identically. Thus, the CCTVpipe user interface can be used as before to control the rest of the functions.




Figure 1. Two-Board Motion Tracking System

TUNING THE MOTION TRACKER

The motion tracking system depends on three control parameters: threshold for motion sensitivity, filter length, and delay for motion smoothing. The parameters, explained in the following sections, are part of the DSP code, and changing them requires a source code modification. Motion sensitivity threshold can easily be changed by modifying the corresponding constant in the source code. A change in delay or filter length requires two things: changing the corresponding constants, and recomputing and modifying the filter coefficients. A C program is available for filter coefficient computation.

ALGORITHM AND COMPUTATIONAL COMPLEXITY

Two unique properties of the ADV601LC/ADV611 video compression chips make it particularly suitable for motion tracking and detection for the CCTV applications:

- Multiscale representation of images inside ADV601/ADV611: This makes it possible to compute motion on a scaled version of a field of video (see next section).
- Quality Box feature of the ADV611: This allows allocating more bits (and thus video quality) to the areas with motion and reduces bits/quality/contrast in static areas.

Initial Motion Computation

The motion computation is done through analyzing the difference between the consecutive video fields. This could be done on fields themselves, but the computational burden would be prohibitive. Fortunately, as a side effect of the compression algorithm we receive a scaled version of the image. The advantages of using the scaled version are:

- Noise Filtering: Because each sample in the scaled version of the image is built by averaging many pixels of the original image, the noise contribution cancels out.
- Reduced computational burden: The scaled version of NTSC image is only 23 × 16 samples (= 368) and 24 × 18 (= 432) for PAL.

After the differences between individual samples are computed, those samples where differences exceed the user-specified motion sensitivity threshold are considered to contain motion. The motion area is defined as the minimal rectangle encompassing all the samples containing motion.

Motion Filtering and Prediction

The motion area computed as described above, is not particularly smooth. If it is directly used as the quality box it will exhibit an unstable motion. The first reason for this is the granularity of the result, because each sample in the scaled version roughly corresponds to a

32 × 32 area in the original image. The second reason is possible spurious differences in the image. Another reason for filtering is that there is always a delay (typically 2 to 3 fields long) between the last field used in the motion computation and the first field to which the computed quality box can be applied.

A special filter was designed to solve these problems. The motion tracker uses a linear approximation of motion for the duration of the filter and predicts where the motion will be after 2 or 3 fields of delay. Because the filter output is computed over multiple fields, the overall effect is some motion smoothing.

As explained above, the filter has two parameters: delay and filter length. Delay specifies for how many fields ahead the quality box positioning is estimated, while filter length is the number of previous positioning samples used for the estimation. Delay is typically 2 or 3 (2 in the current version), and the filter length is 8 to 60 (15 in the current version).

Complexity

Table I shows the number of cycles spent by ADSP-2185 on motion tracking, including the time spent on the decoding on the scaled image (Block 39 of ADV601). This time does not include the cycles spent on searching for Block 39 (this is minimal as the system uses fixed bin widths, thus zeroing out all blocks except Block 39) in the bit stream and other small housekeeping operations, which are also not significant.

Table I.

	Per Field	Per Second	Percentage
Motion Computation	14,768	886,080	12.26
Decoding Block 39	105,504	6,330,240	87.74
Total	120,272	7,216,320	100.00