

Re: MMX speedup for Floyd Steinberg error diffusion

Source: <http://coding.derkeiler.com/Archive/Assembler/comp.lang.asm.x86/2008-05/msg00042.html>

- *From:* Phil Carmody <thefatphil_demunged@xxxxxxxxxxxx>
 - *Date:* Fri, 09 May 2008 17:18:44 +0300
-

Terje Mathisen <spamtrap@xxxxxxxxxxxx> writes:

rep_movsd wrote:

On May 8, 7:00 pm, Terje Mathisen <spamt...@xxxxxxxxxxxx> wrote:

How does it handle the corner/edge cases?

The edges are handled by ignoring all pixels which do not have the neighbours, the simplest way I found was to keep my x and y loops ranging over [1, width-1) and [0, height-1).

OK. Quick & dirty indeed.

It seems to me like the `getNearestPalColor()` function could easily dominate the processing time here!

Well Initially the code used the standard distance matching algorithm as seen in zillions of color reduction libraries :

For each color in the palette
Compute the RGB colorspace distance between the palette entry and the target color $\rightarrow (r - rp)^2 + (g - gp)^2 + (b - bp)^2$
Choose the index with the minimum distance.

This was pathetically slow... It has to loop 256 times for each pixel doing several arithmetic ops!
One optimization is to use `abs` and get rid of the multiplications.

But that messes up the distance calculation because it gives the wrong answer for many edge cases. OTOH, the error diffusion would tend to

Re: MMX speedup for Floyd Steinberg error diffusion

distribute this error across the neighboring pixels, making the visual impact much less.

It doesn't "mess up the distance calculations" at all, it simply selects a different metric, namely the L₁ metric rather than the L₂ metric. However, both are "wrong" in the strict sense, as a metric which models human vision with any attempt at fidelity would weight the three different chroma components differently, giving more weight to the green component, and less to the red component. (Or even take place in a different colour-space entirely.) The volume of space for which the unweighted L₂ metric chooses the wrong palate colour (not the nearest according to the model of human visual perception) will most probably be larger than the volume of the corner cases where L₂ selects the correct colour but L₁ doesn't. (And of course, there will even be times where the L₁ selects a better palate colour than L₂.)

Phil

--

Dear aunt, let's set so double the killer delete select all.

-- Microsoft voice recognition live demonstration

.