

Monday, June 9, 1987

The following test is the total dump of all messages posted to message base in the developer support website as of Friday June 6, 1997. I used autoformat in MS word to help the readability of the messages, but I understand that it is not perfect the next web dump will improve substantially.

Please be mindful to check the date of these messages as some may become inaccurate after time.

-Peter

Forum: General Messages

Current Message System vs. Usenet

Written by Greg LaBrec on Thu May 08 07:41:26 1997

We are at the stage where we need to make a decision. We can either move to the next stage of developing the current message base system (search all forums, mark read/unread, cookies, etc...) or we can deep 6 the current system and switch to Usenet Newsgroups.

I welcome your thoughts on this topic.

Re: Current Message System vs. Usenet
Written by Jeremy Gordon on Tue May 20 21:07:10 1997

To Greg LaBrec:

Yeap, ditch this system. Usenet newsgroups on a private, passworded system is the way to go. All the features of this one, including HTML and binary attachments would still work great. This system makes it really painful to read and post messages.

- jeremy

Re: Current Message System vs. Usenet
Written by Dave Scheele on Mon May 12 13:44:10 1997

To Greg LaBrec:

I advocated using UseNet Newsgroups 'way back when the BBS was still up. Check there for my thoughts on the matter.

Nice to see you're thinking of dumping this system ... it's too cumbersome compared to newsgroups, imho. I think the poster who mentioned the difference in volume of postings here vs. the BBS is dead on. I also like the fact that I can archive usenet postings that interest me locally using my favorite newsgroup reader.

I don't think security is an issue ... my ISP has had local-user-only newsgroups since day 1. I don't know much about such setups, but it seems to me a separate password-protected server would do the trick. However, please DON'T copy over the password file from here ... let us set up our own please! or allow the ability for us to change them

Dave

Re: Current Message System vs. Usenet
Written by Christer Ericson on Fri May 09 22:22:44 1997

Usenet newsgroups are far superior to this system in all aspects. However, considering the fears that feed the NDA system, I don't see how it would be possible (even though it is highly desired) to take all discussions into a usenet newsgroup, even were it to be moderated. Do you plan to have a specific server that we would connect to? How would you solve the logging in problem? Newsreaders do not have password capability built in AFAIK.

Going to some other BBS-like system I don't know about. What makes usenet newsgroups good is just the ability to use a full-featured newsreader of your choice. A BBS system will most likely be a fraction as powerful as your average newsreader, which pretty much would make it on par with this system, thus making the question moot.

All IMHO, of course.

Re: Re: Current Message System vs. Usenet
Written by Tom Plunket on Thu May 22 02:39:15 1997

To Christer Ericson:

>Newsreaders do not have password capability built in AFAIK.<

Actually, the "good" ones do... I know that Netscape supports it, and I use FreeAgent (a freeware newsreader), and it supports password logging-in to servers. In fact, I use it that way every day.

Another note on this, on the password scheme, if each company had their own logins, then your user name and email address would be yours on the posting, but you'd log in using your company's login. Make sense? It would work, since the .newsrsrc (or equivalent) file is stored locally, so you'd only get messages that you personally hadn't read.

- Tom

Re: Re: Current Message System vs. Usenet
Written by Mike Fulton on Fri May 09 23:04:21 1997

To Christer Ericson:

I don't know if the password protection thing is a recent addition to newsgroup server software, but we do have it setup for our Net Yaroze pages. Basically, we would set up a separate news server for commercial developers, and configure it for password access.

I think we may be able to just copy over the web site password list, or maybe even use the same password database. I don't know how that part works.

Mike @ SCEA

Re: Re: Re: Current Message System vs. Usenet
Written by Greg LaBrec on Tue May 13 16:18:43 1997

To Mike Fulton, Christer Ericson:

The only way we could keep it from becoming an administrative nightmare is to copy over the existing logon/password file.

A user would set up their newsreader to access our private server using their company logon. Then, in the "identity" section of their newsreader they would set their own "username" and e-mail address.

Re: Re: Re: Re: Current Message System vs. Usenet
Written by Dave Marshall on Thu May 15 17:21:48 1997

To Greg LaBrec:

The fact that the original message has been posted for a week and only 3 people responded, proves my point that fewer people are using this system versus the old BBS system.

Re: Re: Re: Re: Re: Current Message System vs. Usenet
Written by Rob Vawter on Thu May 22 18:47:57 1997

To Dave Marshall:

Actually, it doesn't prove anything. There is no way to check here how many people have read a message, (unlike the BBS) which was a really good way to tell how much traffic was on the BBS. BBS traffic had its ebbs and flows; as the documentation has improved, the libraries have been improved (via bug fixes and additions of new functionality), and developer skill has improved through experience and training sessions, the BBS traffic decreased markedly. Also, we now have far

more messages coming into the devtech_support address than before. So, the perceived drop-off may be due to other things rather than the current web-site system.

I'm actually for implementing a new system, but the new system certainly does not guarantee that more developers will read the messages.

Re: Re: Re: Re: Re: Re: Current Message System vs. Usenet
Written by Dave Scheele on Fri May 30 18:17:32 1997

I think Dave Marshall's point was that there is less public interaction amongst developers here than there was on the BBS, or would be on a better newsgroup-based system. Your comment that you are now getting far more email at devtech_support kind of points this way too—rather than post in public easily and get input from other developers, it's now easier to just email you folks

A less-cumbersome system would hopefully encourage informal information sharing again. I, for one, miss the "good ol' d days" of gleaning tidbits and chit-chatting on the BBS that just don't happen any more ... for us poor folks located out here in the wilds of PA, it's almost the only way to get to know you Californians!

Dave

Re: Current Message System vs. Usenet
Written by Dave Marshall on Thu May 08 17:37:40 1997

To Greg LaBrec:

Junk this crappy message system.

It takes too much time to hunt thru this system looking for new messages. You need to be able to look at one screen and tell which conferences have new messages in them (like the BBS).

I'd like to know what the average number of postings per day is for the Web site vs. the average number of daily postings for the old BBS. It seemed like the BBS had a lot more message activity on it. If this is true, has anyone there thought about the possible reasons for the decline in message postings?

dave

LOADEXEC

Written by Joseph Wilkerson on Tue Apr 22 22:33:13 1997

Can you Loadexec a program on top of your current program

Re: LOADEXEC
Written by frank on Wed Apr 23 18:32:18 1997

To Joseph Wilkerson:

Yes. But we don't really recommend using this method. If an error occurs while reading the new program, you're hosed. You may have noticed that there's no return value for LoadExec(). When you call it, you're putting yourself at the mercy of the O/S, and the control you relinquished (your program) is overwritten (or ignored if it's still in memory). If the new, expected program isn't delivered as promised, the O/S just sits there and pouts.

It's much safer to use Load() followed by Exec(). It's true you won't have quite as much memory available for the new executable, but the dependability gained is worth sacrificing a few K for a launcher program.

Frank Sheperson
SCEA

Re: Re: LOADEXEC
Written by Tom Plunket on Thu Apr 24 00:17:11 1997

To frank:

As a note on this topic, and from experience:

You can LoadExec() over your current program. The problem with being hosed on a failure, however, can be overcome. LoadExec() _will_ return if the load fails, the problems occur when you've loaded far enough to overwrite where your LoadExec was called from.

Bubsys 3D still uses LoadExec(s), but it's done in a "safe" way to trap errors. We got to submission, and couldn't pass the "overnight" test due to our LoadExecs. It was my task to come up with a solution, without reengineering the game.

So, LoadExec will work, but you have to be clever. Probably best to take Sony's recommendation against it, even though I like my particular solution.

Tom

Re: Re: Re: LOADEXEC
Written by Robert W. Calfee on Mon May 12 23:37:49 1997

To Tom Plunket:

"It was my task to come up with a solution, without reengineering the game. So, LoadExec will work, but you have to be clever. Probably best to take Sony's recommendation against it, even though I like my particular solution."

So, what was your clever solution?
rwc

Re: Re: Re: Re: LOADEXEC
Written by Tom Plunket on Thu May 22 02:43:23 1997

To Robert W. Calfee:

The solution to LoadExec failing was to memcpy the needed routines to some place in memory that wouldn't get stomped on by the new load, then execute it there... In that fashion, you can have fault-tolerant code...

- Tom

Posting messages
Written by Tom Rolfs on Fri Apr 04 20:01:00 1997

Is it possible that when we post a message or reply to one that the message is marked as read for the person who posted it. It is annoying to have to display the message you have just posted. (I prefer to use the "show unread messages" option so that I can quickly see new messages since my last visit.

Thanks,

Tom

Re: Posting messages
Written by Greg LaBrec on Thu Apr 10 17:11:09 1997

The ability to mark messages read and unread is on the "to do" list.

New hardware?

Written by Christine McGavran on Thu Mar 20 20:19:20 1997

We've heard rumors that Tekken 3 will be shipping in Japan later this year with a "hardware upgrade", and Warzard with a RAM card. What is the truth to these rumors, and is there any way we can get more details?

Re: New hardware?

Written by Peter Alau on Fri Mar 21 19:48:20 1997

To Christine McGavran:

Namco is using the PlayStation ARCADE boards for the ARCADE version of TEKKEN3. This may be where the source of confusion came from.....

For the record, our developers will be the first to hear about any changes, additions or technical notes with regards to the hardware via the Red E reports sent weekly.

- Peter Alau

Title Safe Area

Written by John Alvarado on Thu Mar 20 19:36:39 1997

Hello,

Does Sony have a recommended region of the display (in 320x240 mode) that is considered safe for important display information? I'm referring to the problem that a consumer television set will obscure part of the edges of the display, so if you put a small icon at (0,0) it is not likely to be seen on any television set.

I can guess based on the few TV models we have around the office but I was hoping someone might have done more extensive research on this that I could rely on.

- John Alvarado

Re: Title Safe Area
Written by Brian Dawson on Mon Mar 24 04:38:55 1997

To John Alvarado:

In the SCEA Sourcebook - Guidelines section, it is suggested that all text and critical graphics be placed a minimum of 20 pixels away from any edge in 640 x 240 display mode. This is to compensate for overscan differences between televisions. Therefore at 640 x 240:

X Y

min 20 20

max 620 220

These values should be adjusted with respect to the display mode.

Re: Title Safe Area
Written by Ian Schmidt on Mon Mar 24 04:31:31 1997

To John Alvarado

We've used a 16 pixel "safe frame" on all 4 sides in our shipping PSX games; this works ok on every TV I've seen.

str/mdec

Written by dan on Thu Mar 13 21:23:39 1997

Does anyone have some code to read movies from the hard drive and play them (not from the emulator or cd). I would like to know that my str conversion is correct before trying to write code to display them. my phone is 408-374-8300 my e-mail is regalpro@ix.netcom.com

Re: str/mdec
Written by Mike Fulton on Fri Mar 14 02:04:24 1997

To dan:

I'm afraid Sony doesn't have sample code like that. However, I would think you could modify one of our CD-based movie samples to read from the PC using PCread without any big problems.

Mike

Re: Can't Read Messages
Written by Webmaster on Thu Mar 13 18:16:23 1997

The problem should be fixed. You should now have a scroll bar to allow you to move over.

Let me know if my fix creates any new problems.

Re: Re: Can't Read Messages
Written by Dave Scheele on Thu Mar 20 15:26:10 1997

To Webmaster:

- The problem should be fixed.

And it is! I'm using Netscape 3.01, if it matters, and I can now scroll far enough to the right to view ALL of the text on all messages. Thanks!!

- You should now have a scroll bar to allow you to move

I always did have a scroll bar in Netscape .. it just didn't move things far enough to the right. I don't see anything new, but can scroll all the way to the right now.

- Let me know if my fix creates any new problems.

None that I can see. Thanks again!!

Dave

Re: Re: Can't Read Messages
Written by Tom Plunket on Thu Mar 13 20:39:08 1997

To Webmaster:

If you've done what I seem to see, under Internet Explorer 3 I don't get a scrollbar for that, I get the checkboxes wrapped automatically down to the next line. However, my window size is more than sufficient to see it on the same line anyways.

This isn't a problem, just telling you what I see here...

Also, is there any way to "Mark selections as read"? That would certainly be handy.

Tom

Re: Can't Read Messages
Written by Webmaster on Mon Mar 17 18:43:24 1997

To Tom Plunket:

Sorry about that, usually I write for Netscape first and then see what I have to fudge to get it to work with Explorer. In this case, I didn't have to check on the effect with Explorer but thanks for doing it for me. I'll poke around and see what I can do but I usually have little luck getting table definitions to work properly with Explorer.

"Also, is there any way to "Mark selections as read"? That would certainly be handy."

It's on the list of things we would like to see happen. I submitted a list of changes to Proxicom, the company who wrote the cose for the message base system. I should be hearing something this week.

Can't read messages

Written by Dave Scheele on Tue Mar 11 15:23:41 1997

I still can't read several of the messages on this web site because their subject line, author and date are so long that the little check box I need to click to display the message is pushed off the right of the screen. Could this box be moved to immediately after the subject? I really don't mind if the author and date written get cut off, but I need that little box!!

Dave

Re: Can't read messages
Written by Mike Fulton on Tue Mar 11 23:20:42 1997

To Dave Scheele:

Have you tried going to your browser preferences and changing the font size to something smaller?

What browser are you using?

Mike

Re: Can't read messages
Written by Jerry Jessop on Tue Mar 11 19:58:45 1997

To Dave Scheele:

I had the same problem a few weeks ago Dave and it has been passed on to the Webmaster.

Unfortunately you have to be set up in 1024 by 768 to view all threads....

My 800 x 600 would not cut it.

Jerry

Re: Can't read messages
Written by frank on Tue Mar 11 17:25:15 1997

To Dave Scheele:

We'll pass along your request to the webmaster, but in the meantime...have you tried 1024x768?

Frank

Bad Floating Point libs
Written by David Ross on Mon Mar 03 21:27:21 1997

We are currently having problems with the floating point library. If one adds (or subtracts) two very small floating point numbers, the result can be minus infinity (=0xff800000). Presumably, if the add caused underflow, the result would be 0.0 (and NOT minus infinity).

But the REAL problem we are having is that this variable, after it has been set to minus infinity, will then cause an infinite loop in __floatsisf if it is used in an operation (such as multiply or add). The infinite loop is at __floatsisf+0x40 and is the following 3 instructions:

```
800E93C4 sll a0,a0,$4 800E93C8 slt v0,a1,a0 800E93CC beq v0,zero,$800e93c4
```

This is becoming quite frustrating as it happens frequently and there is no way to detect it (except for testing writing out floating point expressions as a series of intermediate values and always checking each value. I believe that if the minus infinity DIDN'T happen and instead a zero resulted from underflow that this wouldn't be a problem.

Is there anyway that this can be fixed soon, or at least a work-around developed?

PS: We are using lib 3.6, with compiler switch -fshort-double. Code that exhibits the problem follows:

```
{  
  
float ftemp, ftemp2, ftemp3;  
ftemp = -1.25386E-038;  
ftemp2 = 6.2693E-039;  
ftemp3 = ftemp + ftemp2;//this will result in ftemp3 being set to -inf  
ftemp = ftemp3+ftemp2;//this will result in an infinite loop in __floatsisf  
}
```

- Dave

Re: Bad Floating Point libs
Written by Mike Fulton on Tue Mar 04 02:04:53 1997

To David Ross:

I don't really know what the standard is for handling underflow conditions like this, but have you forwarded information on this problem to SN Systems? They are the ones who work on the tools like the C compiler.

Their EMAIL address is "support@snsys.com".

Mike

Re: Re: Bad Floating Point libs
Written by David Ross on Wed Mar 05 17:49:30 1997

To Mike Fulton:

In researching this bug on the SONY BBS I find that a similar bug was reported over a year ago in the floating point multiply: multiplying a negative number times 0.0 resulted in a -0.0 (0x80000000) which, when converted to an INT, would result in the INT set to 0x80000000 instead of 0. I find that this year+ old bug still exists in the latest version of libmath (if you convert the bogus INT0x80000000 to a float, then __floatsisf will go into an infinite loop).

This leads to several questions:

If that bug was never fixed, can I expect ANY action on the bug I recently reported?

The reporter of the original bug ended up re-writing __mulsf3() (BBS MSG From: Tim L.

Lowery

Subject: Re(3): Floating point multiplication bug) to fix the bug:
has his fix been posted
anywhere that I might get a copy of it, so that I might use it as a
template for fixing the floating point add/subtract bug?

Does SONY have the source to the floating point routines? And can I get a copy so that I can fix the bug, or will I be forced to take the time to figure out how to rewrite it (INSTEAD of programming on the game I'm working on)?

Re: Re: Re: Bad Floating Point libs
Written by Mike Fulton on Thu Mar 06 02:34:24 1997

To David Ross:

Sorry I didn't connect your situation with the "negative zero" problem at first, but that was resolved many months ago and it just didn't occur to me.

In researching the "negative zero" problem some months back, I reported it to SCEI (i.e. Sony Computer Entertainment Japan) and was told by their engineers that this was in fact common behaviour among floating point libraries and that they didn't think it should be changed.

When I heard this, I did some tests and was able to confirm that the same results did in fact occur with other compilers (non-Playstation and non-MIPS compilers, such as Visual C++ for MSDOS/Windows).

The opinion was that since this behaviour was common to other FP libraries besides Sony's, changing it may result in a different set of problems occurring. So the decision was made to not change it in the library.

Since we didn't want to change the library, I made a fix to that one particular routine. This was separate from whatever Dave Howell did and in fact I think he switched over to what I gave him. (Dave was the one actually using the "Tim Lowery" account on the BBS.)

After I made a fix, I sent him a replacement object module for that one particular routine. This can be found on the BBS (a message dated July 23, 1996), but I'll attach the object module to this message so you can download it here.

I'm not sure if the same fix will apply to your situation, but please give it a try and let me know.

Mike

Re: Re: Re: Re: Bad Floating Point libs
Written by David Ross on Tue Mar 11 17:35:55 1997

To Mike Fulton:

Indeed, the new mul3 fixes the negative zero bug, and I thank you for that. But that is NOT my main problem: my main problem is that the libmath routines, when they achieve a result so close to zero as to be unrepresentable, chooses NOT to return zero, but to return THE LARGEST POSSIBLE VALUE, the value the FARTHEST from zero that they can get. This is BAD. What is WORSE is that if you convert one of these negative results to an INT, you get 0x80000000; and if you try to convert that to a float, the libmath routine LOCKS UP in an infinite loop. You and your engineers may think that behaviour is OK, but any game programmer (indeed, any embedded system programmer) will disagree.

I have modified addsf3 and mul3 to return a ZERO on an underflow and modified fixfsi to return 0x80000001 (instead of 0x80000000) which prevents the conversion to float from locking up. All is working merrily: no more lock-ups in LIBMATH!!!!

Re: Re: Re: Re: Bad Floating Point libs
Written by Mike Fulton on Tue Mar 11 23:30:51 1997

To David Ross:

I'm glad you've fixed the problem for now.

Actually, just so you know, I don't especially think it's OK for it to not be changed. I understand the concern from SCEI that a change might cause other problems, but personally I'd prefer to try it and see. But the decision isn't in my hands. All I was doing was repeating back what they told me.

However, the more I think about it, the more I'm not sure that the problem with it locking up when it tries to convert negative zero was included when we originally reported the negative zero problem to SCEI. So I'm going to write up a new report to them that includes that information and we'll see what they say this time around.

Mike

Re: Re: Bad Floating Point libs
Written by David Ross on Wed Mar 05 16:41:28 1997

To Mike Fulton:

SN System's reply:

>I'm afraid it'll sound like you've caught us on Buck-Passing Day, >but the floatsisf function is implemented by Sony, in FLTSISF.OBJ >in LIBMATH.LIB, and we don't have the source they used to build it. >We can take a look at a possible patch function for it if you like, but I >can't promise anything immediate.

>

>If it's any encouragement, the compilers I have for other MIPS >systems (such as the Irix) handle your example fine, the point being >that it's definitely down to the implementation of floatsisf, rather than >being a problem in general.

So, the ball is back in your court. Echoing my other reply to this message: Is there any way we could get the source to the floating point modules of libmath so we could fix the problem SOON. Has any developer used third party floating point routines?

- Dave, dross@sculptured.com

Re: Re: Bad Floating Point libs
Written by David Ross on Tue Mar 04 21:55:58 1997

To Mike Fulton: I will email SN Systems, but it seems that this is a floating point LIBRARY issues (as opposed to a compiler issue). Does SN Systems maintain libmath? Is it possible for me to get a copy of the source to the floating point modules of libmath so that I might make a fix? Are you aware of any PSX developers using third-party floating point modules? -Dave

New Libs
Written by Jeff Beckham on Wed Feb 19 03:54:36 1997

Support,

When will the 3.7 metrowerks libs be available? Can somebody do a quick in house conversion and post them?

Jeff

Re: New Libs
Written by Peter Alau on Fri Feb 21 17:18:24 1997

To Jeff Beckham:

Hi Jeff,

We just posted those tools yesterday under the programmertools/beta section. Hope this helps,

- Peter

Experienced developers available
Written by Jay Stelly on Tue Feb 18 18:33:21 1997

An experienced Sony PlayStation development team is available for contract work immediately.

The team consists of both artists and programmers who are familiar with all aspects of PlayStation development, including direct GTE/GPU programming. We have recently completed NanoTek Warrior (a 30fps arcade shooter using an in-house developed 3D engine) to be released by Virgin Interactive in March.

Our artists are capable of producing high quality in-game graphics, as well as exceptional pre-rendered sequences. Our programmers are well versed in 3D rendering techniques, especially as they relate to the Sony PlayStation.

Our team is able to create challenging and innovative titles, on budget and on time. We are available to begin a new title immediately.

Contact : Mark Skaggs

Tetragon Inc.

(972) 680-8101 ext 11

MSTetragon@aol.com

Bug free MOVCONV and MOVPACK please...

Written by Simon Freeman on Thu Feb 06 01:37:23 1997

Hi,

The latest versions of MOVPACK and MOVCONV (Dev CD 1.8) are a liability.

Crashing is a frequent occurrence, files are not displayed in selection boxes, having to be typed in rather than selecting from a list, and there is no indication of progress, indeed it is possible to exit the tools while they are still processing, resulting in an eventual, though by no means immediate crash.

Do you have a policy regarding which version of these tools should be used?

Best,

#6

Re: Bug free MOVCONV and MOVPACK please...
Written by Mike Fulton on Wed Feb 19 04:01:55 1997

To Simon Freeman:

I haven't been able to replicate the problems you mentioned, with the exception of being able to quit while a movie is being converted (which causes a crash).

Can you provide more details and/or a sample file that demonstrates the problem?

Re: Bug free MOVCONV and MOVPACK please...
Written by Chia-Ming Wang on Mon Feb 10 19:57:35 1997

To Simon Freeman:

By the way, when you say the "latest" movconv and movpack on the CD's are a liability, to which versions do you refer?

movconv 2.1beta1

movpack 1.5e-beta1

OR

movconv 1.98e

movpack 1.4e

You can tell the version number by reading the title bar of the application's main window.

Thanks,

Chia-Ming

Re: Bug free MOVCONV and MOVPACK please...
Written by Chia-Ming Wang on Mon Feb 10 18:48:43 1997

To Simon Freeman:

I'm sorry they are causing you so many problems.

Briefly, we have two new versions that you can download from the WebSite: MovConv 2.1e and MovPack 1.5e. If you have the proper permissions,

1. Go back to the Developer Support Homepage.
2. Select "files" on the left-hand side of the screen.
3. Select "Programmer Tools" and fill in your password.
4. Scroll down to the heading "GRARAPHIC TOOLS" and

select movconv.zip.

These also exist in Beta form on the Dev CD 1.8 in the \beta\movconv directory if you can't get to the Web right away.

Please let us know if you have any problems, and whether or not this helps,

Chia-Ming Wang

Changing Forums w/out Entering Password
Written by Webmaster on Fri Jan 10 22:54:50 1997

If you use the "Other Forum" button to switch between message bases rather than using the "Back" button on your browser, you will not have to re-enter your username and password.

Re: Changing Forums w/out Entering Password
Written by Borden Covell on Wed Jan 08 01:22:13 1997

To Webmaster:

It would be nice if the Change Forum option was available at all times. As it is now, that option is not available while you are reading a message.

Thanks for the quick response on the other issues.

Borden

Re: Re: Changing Forums w/out Entering Password
Written by Steve King on Fri Jan 17 14:13:23 1997

To Webmaster:

It would also be nice to have a "Next Message" button somewhere that I didn't have to scroll to get to. The little message tree display at the bottom of the message is nice, but when I'm reading a long message that I really don't care about it'd be better to be able to go to the next message without scrolling down to find the link.

Also, when reading the last message in a reply chain the tree fails to display the first message of the next topic, making me have to go back to the "Contents" page to pick up the next thread.

Re: Changing Forums w/out Entering Password
Written by Webmaster on Wed Jan 08 17:25:42 1997

To Borden Covell:

DONE!

Good suggestion (...and one of the few that we've received that could actually be implemented)

All SCEA Engineers

Written by Greg LaBrec on Wed Jan 15 16:48:52 1997

I have to add your message base logon name to the SCEA Engineers group so you will have access to the special icons when posting or replying to messages. I will do this automatically when I detect that you have set up your message base account.

Posting a reply to this message will help me notice faster.

Re: All SCEA Engineers
Written by Eric W. Brown on Tue May 20 18:26:17 1997

To Greg LaBrec:

Ok, I've set up a logon name of eric.

You can reach me at eric@virtualmusic.com, if you need to.

Eric W. Brown

eric@virtualmusic.com

Re: All SCEA Engineers
Written by Kelly Walker on Thu May 08 04:35:39 1997

To Greg LaBrec:

I do reply.

Re: All SCEA Engineers
Written by Robert W. Calfee on Thu Feb 20 22:50:08 1997

To Greg LaBrec:

I'm not sure if I'm supposed to respond here, but I am anyway.

Robert

Forum: LIBGPU/GTE

VSynCallback
Written by MYamada on Wed Jun 04 22:58:44 1997

I wanted to install a callback function through VSynCallback which would enable me to update small stuff on the screen while I was loading files from the CD. The callback function, however, does not get called unless an explicit call to VSynC() is made, which really defeats the purpose of the callback function. Is there something else I need to do? Currently, I've implemented the function by

creating an Event but the Event gets pre-empted when a sound interrupt occurs and I can see hiccups in the movement of the graphics when the interrupt occurs.

I'd appreciate any suggestions anyone might have. Thanks!

Re: VSyncCallback
Written by Mike Fulton on Thu Jun 05 03:20:10 1997

To MYamada:

First of all, you should take a real hard look at the advisability of doing this sort of stuff in the vertical blank period. As a general rule, the best stuff to do in the vertical blank is synchronize your graphics buffer switching to the display.

Since CD access can be done asynchronously, you should have no problem updating the screen in your mainstream code, and then you can deal with the CD stuff from callbacks.

You're mistaken about the way VSyncCallback works. It does not require that VSync be called for you to get an interrupt each VB period.

Mike

Perspective transformation, Z
Written by Scott Hartin on Fri May 23 00:19:23 1997

Someone correct me if I'm wrong, but it appears to me that the GTE "pins" rotated/translated Z values to 16bit before performing perspective calculation.

Here is my problem:

My world coordinate system uses 8 bits of fractional. I am directly programming the GTE, and when I RotTransPers a point and the ROTATED Z is greater than can fit in a 16bit SIGNED value, I can see my polys being "pinned" to some plane in the distance.

Since I'm using 8bits of fractional, I can only draw polys up to 128 integer units away before they get "pinned".

Am I losing my brain, or is the GTE pinning rotated/translated Z before performing perspective calculation? Is there some way to get around this problem?

Thanks
Scott

Re: Perspective transformation, Z
Written by Michael Koziniak on Tue May 27 19:43:57 1997

To Scott Hartin:

Your right, the GTE will limit your world coordinates before doing a perspective calculation. I will look into ways of getting around this problem.

mk

Polygon Subdivision

Written by Mike Michaels on Fri May 09 01:23:09 1997

I'm trying to implement polygon subdivision with DivideFT3 and RCpolyFT3. The problem I'm having is in predicting the number of polygons a particular division is going to produce. The reference says that divisions are 2x2, 4x4, 8x8, 16x16 or 32x32. I just don't understand what that means. Will 32x32 generate 32 polygons from a single polygon or will it generate 32 recursions per polygon? I'm preallocating space for polygons so I need to be able to predict how many polys are going to be used when this function returns.

Any help would be appreciated,

Mike Michaels
Virgin Interactive Entertainment

Re: Polygon Subdivision
Written by Mike Fulton on Fri May 09 17:44:18 1997

To Mike Michaels:

Using 2x2 subdivision will divide a polygon into 4 polygons, 2 across and 2 down. Using 16x16 subdivision will create 256 polygons, 16 across and 16 down. Etc.

Mike

Fog

Written by Stephane Leblanc on Fri May 02 20:48:41 1997

Is there an easy way to fade a textured polygone to any colors? I've always faded my polys to black since today and it work fine but when I tried different colors such has white, red ..., my textured polys seems to saturate before reaching the FarColor (RFC,GFC,BFC), meaning that they will be whitish but never so that we can't see the texture details. I don't want to change the clut as shown in the Sony samples cause it would be totally impossible to manage.

Any help would be appreciate.

Re: Fog
Written by Simon Everett on Wed May 07 19:32:23 1997

To Stephane Leblanc:

The GPU will not allow you to fog textured polygons to arbitrary colours using the method described by Sony, but there is a way to do this without using cluts. How many polygons are on screen, and how many of these need to be fogged?

Simon Everett.

Re: Re: Fog
Written by Stephane Leblanc on Thu May 08 13:31:48 1997

To Simon Everett:

It's really hard to know how many polygons need to be faded. There are between 3500 to 7000 polygons for the environment. After clipping, I have around 1700 to 5500 polygons and I would say that 10 percent of them as to be faded not forgetting to add about 1500 to 3000 polygons for characters and objects.

I thought of adding a flat transparent polygone over my textured one but my GPU drawing time is already reaching the red line as well as my packet buffer.

So I've resign myself to use cluts and it's working fine even though I thought it would have been a lot more difficult to implement but I'm running quite low on VRAM space so even if what you have in mind would not work for me, I would like very much to know about it anyway.

Thank.

Stephane Leblanc
Behaviour Interactive

Re: Fog
Written by Michael Koziniak on Fri May 09 19:30:12 1997

To Stephane Leblanc:

I'm sorry but there is not great way to fade textures to any color. The PlayStation's interpolation value will work according to the range of color values your working with. When you move away from black there is less to interpolate between.

michael koziniak

BLK_FILL

Written by Peter Lipson on Thu Apr 17 19:45:53 1997

If I remember correctly, this primitive type was removed since you can get the same results by drawing with an appropriately colored POLY_F4. Isn't this true? But it raised a larger question:

I couldn't find any reference to it either on this site or on the developer CD. What's the best way to search for old information like that?

Peter Lipson
Mindscape

Re: BLK_FILL
Written by Mike Fulton on Fri Apr 18 00:04:13 1997

To Peter Lipson:

The best thing to do is simply ask us what you need to know. It usually also helps if you give us a bit of background information about what you're doing.

Yes, you can do essentially the same thing as BLK_FILL by using POLY_F4. However, using the TILE primitive also does the same thing and is faster.

Mike @ SCEA

fading of sprite

Written by John Bryant on Wed Apr 16 01:00:59 1997

I have been trying to fade a sprite (graphically) from normal appearance to non-existence. I have tried the various transparency rates mixed with various texture shading techniques (ie shading the sprite to 0, etc) and have not been successful.

Any ideas?

John J Bryant

Realtime Associates, Seattle Division

Re: fading of sprite
Written by Mike Fulton on Fri Apr 18 23:52:48 1997

To John Bryant:

David Coombes has put together a sample which demonstrates this technique. You'll find it in the attached file.

Name: Fadedemo.zip
Content-type: application/x-zip-compressed
Size: 65558 bytes

Re: fading of sprite
Written by Mike Fulton on Fri Apr 18 18:42:49 1997

To John Bryant:

David Coombes has come up with a sample that should demonstrate this technique. It's attached to this message. Have a look, and if this doesn't do what you had in mind, please let us know more precisely what you'd like to do.

Mike

Re: Re: fading of sprite
Written by Jeff Beckham on Wed May 21 18:19:46 1997

To Mike Fulton:

Mike, the attached file describing the fading of a sprite is no longer available. Would it be possible to get a copy of this?

Thanks,
jeff

Re: Re: Re: fading of sprite
Written by Mike Fulton on Wed May 21 23:55:41 1997

To Jeff Beckham:

The file with the sample code is still available and in fact I just downloaded it to make sure.

I suspect you've got yourself looking at only the most recent messages and the message with the file is just not being shown. Try setting the message base to display "All" messages within the category and then it should show up.

Mike

Re: fading of sprite

Written by John Bryant on Mon Apr 21 21:40:34 1997

To Mike Fulton:

Thank you, this is exactly what I was looking for!

John J Bryant

Using the DR_TWING primitive

Written by Mike Michaels on Thu Apr 10 18:19:18 1997

I'm attempting to implement semi-general texture tiling by inserting texture window primitives into the ordering table around the drawing primitives that they affect. Therefore, each tiled drawing primitive in the ordering table looks something like this:

```
TexWin (x:16, y:48, w:32, h:32)
PolyFT3
TexWin (x:0, y:0, w:256, h:256)
```

This should enable a texture window, draw the primitive and then disable the texture window (I think).

My problem occurs when I'm trying to use multiple tiled textures within a single model. I dump the ordering table and see the various different texture window primitives, but the rendered model only shows the first texture. The texture is tiled everywhere that I would expect to see tiling.

My question is: is there a problem inserting multiple different texture window primitives into the ordering table? Alternatively, is there a specific alignment requirement when I load textures into vram for tiling?

Thanks,

Mike Michaels
Virgin Interactive Entertainment

Re: Using the DR_TWING primitive
Written by Mike Fulton on Fri Apr 11 21:21:51 1997

To Mike Michaels:

There is no restriction regarding having multiple texture windows within your OT. And I'm not aware of any alignment requirements for windowing that don't apply to textures in general.

Can you confirm that you are adding the primitives to your OT in the right order? Remember that primitives at the same depth are rendered in the reverse order of how they were added to the OT.

If you are adding the polygon and DR_TWING to the same depth position in your OT, you must add your reset-window primitive first, then your polygon, and then the DR_TWING primitive.

Basically, what you want is this:

```
AddPrim( ot, my_reset_twindow_prim );
AddPrim( ot, my_poly_ft4 );
```

```
AddPrim( ot, my_set_twindow_prim );
```

The first one will reset the texture window for any polygons you've already added to your OT previous to this point. Then you have the polygon that will use the texture window, and then you have the actual set-twindow primitive.

You may also consider using DR_MODE instead of DR_TWIN.

If you're going to be changing the texture window, it's a good idea to reset it at both ends of the group of polygons, so you're certain you won't affect anything you don't want t-windowed.

Have you tried using DR_MODE instead of DR_TWIN?

Also, you may consider creating a small sub-OT that will contain the primitives for the texture windowing stuff, and then link that into your main OT. That may make it much easier to deal with, like this:

```
/* declare and init sub-OT */  
  
long subot[3];  
ClearOTag( subot, 3 );  
  
/* Add texture windowing primitives to sub-OT */  
  
AddPrim( subot, my_reset_twindow_prim );  
AddPrim( subot, my_poly_ft4 );  
AddPrim( subot, my_set_twindow_prim );  
  
/* link in sub-OT */  
  
AddPrims( ot, subot, &subot[2] );
```

Hope this helps!

Mike

<p>Re: Using the DR_TWIN primitive Written by Mike Michaels on Fri Apr 11 22:34:53 1997</p>

To Mike Fulton:

I have confirmed that I'm adding the primitives in the correct order. I'm dumping the ordering table to the debug screen and everything looks fine.

I haven't tried using DR_MODE. I'm already concerned about the performance hit we're going to take submitting multiple DR_TWIN primitives. I'm worried that since DR_MODE sets everything, it's going to take even longer. But if tiling isn't going to work for me otherwise, I'll go ahead and give it a try.

Thanks,

Mike

Virgin Interactive Entertainment

Re: Re: Using the DR_TWING primitive
Written by Mike Fulton on Sun Apr 13 20:14:47 1997

To Mike Michaels:

The DR_MODE and DR_TWING primitives take the same amount of time to process in the GPU.

Also, since these primitives do not write to VRAM, they take noticeably less time to process than does something like a polygon primitive.

Mike

Re: Using the DR_TWING primitive
Written by Mike Michaels on Mon Apr 14 21:45:32 1997

To Mike Fulton:

After more documentation perusal, I found out what was causing my problem. In the library ref, at the bottom of the description of the DRAWENV structure, there's a table of texture alignment requirements based upon texture window size.

Problem is, the table is incorrect. After some experimentation, it looks like the correct table is:

tw.w 16 32 64 128
tw.x multiple multiple multiple multiple

of 32 of 64 of 128 of 256

and likewise for the vertical direction.

Concerning the performance of the DR_MODE primitives, I remember seeing in the conference slides a warning against using DR_MODE because it was slow. Do you have access to any timing numbers on execution of this primitive. I assume that the DR_TWING primitive is just promoted to a DR_MODE primitive before it's passed to the GPU for execution.

Thanks,
Mike
Virgin Interactive Entertainment

Re: Re: Using the DR_TWING primitive
Written by Mike Fulton on Tue Apr 15 01:59:45 1997

To Mike Michaels:

I'll look into the information about the texture alignment thing.

As far as the DR_MODE and DR_TWING primitives go, I do not have access to timing information. They are not quite the same thing, but they are pretty close.

If you just want to change the texture page, then doing something like a POLY_FT3 is faster than a DR_MODE. This is what was being referred to in the slides you mentioned. (Keep in mind that the slides were support for the presentation, and didn't necessarily contain all the details.)

Mainly, I was wondering if you'd had different results with trying DR_MODE instead of DR_TWIN.

Mike

Re: Re: Re: Using the DR_TWIN primitive
Written by Mike Michaels on Fri May 09 01:14:28 1997

To Mike Fulton:

I didn't find any difference between using DR_MODE and DR_TWIN once I aligned the textures properly. I ended up using DR_TWIN in the end because its effect was limited to what I wanted to accomplish.

Mike

GPU Timeout - decoding diags.

Written by David Minogue on Wed Apr 09 18:45:02 1997

Can anyone show us how to decode the following GPU timeout message?

GPU Timeout

Que = 0,

Stat = 40020407

CHCR = 01000401

MADR = 0019DA90

FUNC = (80024EBC) (800436FC), 00000000)

In particular, how do we tell what primitive caused the error?

David Minogue

Re: GPU Timeout - decoding diags.
Written by Mike Fulton on Wed Apr 09 21:34:43 1997

To David Minogue:

GPU timeouts are caused by your primitive list becoming corrupted.

Basically what's happening is that either the link address in a primitive was bogus, or else the primitive type/length code was bogus.

The main ways you can corrupt a list are:

- 1) Link in the same primitive more than once. This creates a set of circular references which the GPU can't break out of.
- 2) Randomly overwrite memory. Accidental corruption is almost certainly due to something stepping on memory it doesn't own.
- 3) Start reusing the memory for a primitive list before the GPU is finished processing it.

In the message you saw, "Que = 0" means there are no items in the LIBGPU execution queue after this one.

The other values for "Stat = 40020407" and "CHCR = 01000401" are for things like the GPU rendering and DMA operation status flags. The important thing here is that it says the GPU is busy.

The only potentially useful thing in this message is the MADR value. This is the last address where it was trying to do DMA. This may be the address of the corrupted primitive. (Or maybe not.)

The best way to fix a corrupted primitive list bug is to dump it at certain positions to insure it's still OK. This would be done immediately prior to the "DrawOTAG" call, and again immediately after your "DrawSync" call or within your "DrawSyncCallback" routine. You might also do it at other places along the way depending on how your code is done.

Mike

Light Matrix

Written by Matt Fairweather on Thu Apr 03 00:24:45 1997

What do the elements of the LightMatrix and the ColorMatrix represent? I.e. are these matrices composed of 3 light/color vector couples? If so, does each row need to be normalized? I'm really lost here. Please help.

Re: Light Matrix
Written by Brian Dawson on Thu Apr 03 21:44:26 1997

To Matt Fairweather:

The light matrix is:

X1 Y1 Z1

X2 Y2 Z2

X3 Y3 Z3

The color matrix is:

R1 R2 R3

G1 G2 G3

B1 B2 B3

Each of these values is expressed in fixed-point (0=0,4096=1). Though the gte will perform these calculations on non-normalized (>4096) values, looking at the GTE reference it appears that under certain circumstances you'll get undesirable results. Therefore to be safe you may want to normalize (make them into unit vectors) before passing them to the GTE. Note that the GTE does not limit these values until after they have been applied to the appropriate vectors. Please refer to the GTE Command Reference and Programming Guide for further details. These can be found on the 1.8 Tech. Ref. CD in the technotes directory and were also handed out at the last Dev. Conference.

Brian Dawson@SCEA Dev. Support

dmpsx 3.0+ documentation

Written by Jim Rogers on Tue Mar 25 00:37:00 1997

Where can the file dmpsx.txt or documentation on the new inline headers (inline_o.h inline_c.h and inline_a.h) and dmpsx 3.0 be found?? Both the web site and the lib 3.6 CD's refer to the dmpsx.txt file for updated info but neither seem to contain such a file.

in particular I am looking for description of the differences and advantages/disadvantages between the inline_a.h and inline_c.h header files.

thanks

Jim

Re: dmpsx 3.0+ documentation
Written by frank on Tue Mar 25 01:32:31 1997

To Jim Rogers:

Try looking in \Technote\Gte_func.pdf on the reference CD. This should have what you're looking for.

An Acrobat search for "inline_c.h" will find this, but searching for "inline_a.h" will not.

Interesting...

Frank

GetGraphType() - What is it?

Written by Mick West on Thu Mar 20 19:33:43 1997

What does GetGraphType() do? it seems to always return 0, and I can't find any mention of it in the documentation.

it is seemingly only used in getTPage(...)

```
#define getTPage(tp, abr, x, y) \  
((GetGraphType()==1||GetGraphType()==2)? \  
(((tp)&0x3)<<9)|((abr)&0x3)<<7)|((y)&0x300)>>3)|((x)&0x3ff)>>6):\  
(((tp)&0x3)<<7)|((abr)&0x3)<<5)|((y)&0x100)>>4)|((x)&0x3ff)>>6)| \  
((y)&0x200)<<2)))
```

I assume it is something like an interlace mode flag or something that does not usually change, so one could write a much faster version of getTPage(...) without it.

Does anyone know what it does, or where there is documentation on it?

Mick West
Neversoft

Re: GetGraphType() - What is it?
Written by Mike Fulton on Fri Mar 21 01:16:09 1997

To Mick West:

GetGraphType is an internal library function which is purposefully not documented.

It returns a GPU code which is required for getTPage to work properly. This code will be different under different circumstances, so the call cannot be removed.

Mike @ SCEA

Re: Re: GetGraphType() - What is it?
Written by Jeff Blazier on Tue May 06 00:01:52 1997

To Mike Fulton:

Is there a prototype for this function somewhere? My compiler always get a warning from any macro calling this function even when libgpu.h is included...

Also, I assume the function is not inline so will cause icache faults!?

Thanks,
Jeff Blazier
Lobotomy

Re: GetGraphType() - What is it?
Written by Mike Fulton on Tue May 06 20:00:03 1997

To Jeff Blazier:

There is no prototype for GetGraphType. Add the following to the list of prototypes at the bottom of LIBGPU.H:

```
int GetGraphType(void);
```

The function is not inline (none of Sony's library stuff is designed around C++, and C doesn't have inline functions).

Calling it will possibly result in an I-cache miss. But it is very short... my guess is no more than 8 instructions, maybe as few as 4.

Mike

DR_LOAD Primitive Usable?

Written by John Alvarado on Tue Mar 04 22:36:23 1997

Is it possible to use the DR_LOAD primitive in a primitive list to download new texture data to VRAM? I didn't see any macros or functions in LIBGPU.H on how to setup this primitive. The structure declaration for DR_LOAD is not very informative.

The reason I ask is because I would like to be able to draw more graphics than can fit in VRAM and I would like my high-level code to just call a draw routine that handles getting the data into VRAM if it is not already there by adding the DR_LOAD command to the primitive draw list.

Thanks

- John

Re: DR_LOAD Primitive Usable?
Written by Mike Fulton on Thu Mar 06 03:11:40 1997

To John Alvarado:

The SetDrawLoad function initializes a DR_LOAD primitive. It takes a pointer to a DR_LOAD structure and a RECT structure that defines where in VRAM you want the data to go.

After you've used SetDrawLoad, the actual bitmap data goes into the DR_LOAD->p array, in elements 0 to 11. Element 12 of this array is reserved.

Please note that you can transfer a maximum of 48 bytes per DR_LOAD primitive. Larger bitmaps will have to be broken down and transferred in segments.

Mike

PSYLINK errors with latest DMPSX

Written by Jose Villeta on Wed Feb 26 18:48:47 1997

After upgrading to use inline_o.h and latest dmplx, we got a psylink error with the gte in-line functions:

value(0000 00D4) out of range in instruction patch (expression involves: wetsuit)

We are using DMPLX.EXE version 3.01 and lib3.7

Any news, if we go back to version 2.06 and inline.h we get no errors.

Jose

Re: PSYLINK errors with latest DMPLX
Written by Mike Fulton on Tue Mar 04 02:19:25 1997

To Jose Villeta:

This linker error typically means there was a 16-bit offset that was the target of a larger than 16-bit value. I doubt that it had anything to do with the version of DMPLX you were using except to the extent that things were moving around in memory.

You need to isolate the portion of the source code where the error is happening. Try breaking down the source code file specified into two separate files, then move things from one to the other until the error moves too. Then you'll know what is causing the error.

Mike

On clearing the OT [with code]

Written by Christer Ericson on Mon Feb 24 20:03:17 1997

Some time ago, it was pointed out to me by Davide of Digital Dialect that ClearOTag() was extremely slow, taking almost a frame for a large 16000+ entries OT (not that you'd want an OT that size, but still). Looking at the code, I found that it was far from as efficient as it could be, and came up with the following patch that I thought I'd share.

Of course, it should be pointed out that ClearOTag() shouldn't really be used in the first place, as ClearOTagR() which uses the hardware to do the linking, is several magnitudes faster. Still, if one is low on GTE time but has CPU time to spend, then by all means use (the patched) ClearOTag().

Hope this code is of use to someone, and perhaps we'll see the changes incorporated in lib 3.8?

```

;-----
;
; Function: extern "C" int PatchClearOTag(void);
;
; Output:  0 - No error
;         -1 - Patch not installed (lib code doesn't match)
;
; Description: Patches library ClearOTag(), speeding it up
;             by a little more than a factor two.
;
; By: Christer Ericson, Neversoft Entertainment, 970125
```

;

```
xref ClearOTag
xdef PatchClearOTag
```

PatchClearOTag

```
    li    t0,14
    la    a0,ClearOTag ; Before installing patch, make
    la    a1,@cmatch  ; sure library code matches the
    move  a2,a0       ; one that patch was created for
@1     lw    t1,$50(a0)
        lw    t2,(a1)
        subu  t0,t0,1
        bne  t1,t2,@abort ; Doesn't match, abort!
        addu  a0,a0,4
        bnez t0,@1
        addu  a1,a1,4

@copy  li    t0,14      ; Copy new code on top of old
        la    a1,@cpatch ; code
@2     lw    t1,(a1)
        subu  t0,t0,1
        sw    t1,$50(a2)
        addu  a1,a1,4
        bnez t0,@2
        addu  a2,a2,4

        jr   ra
        li   v0,0      ; Error 0: "no error"

@abort jr   ra
        li   v0,-1     ; Error -1: "not installed"

@cpatch ori  v1,v1,$ffff ; New code to be copied on top of
@3     addiu s1,s1,-1    ; the old code that sucks asteroids
        addiu s0,s0,4
        and  v0,s0,v1    ; This is the loop that links
        bne s1,zero,@3  ; the OT list
        sw  v0,-4(s0)
        rept 8

nop
endr

@cmatch lui  a1,$00ff   ; This is part of the original
        ori  a1,a1,$ffff ; ClearOTag() code, duplicated
        lui  a2,$ff00   ; here in order to be able to
@666  addiu s1,s1,-1    ; match it before copying over
        addiu a0,s0,4   ; the new code
        sb   zero,3(s0)
        lw   v0,0(s0)
        and  v1,a0,a1
        and  v0,v0,a2
        or   v0,v0,v1
```

```
sw    v0,0(s0)
bne   s1,zero,@666
move  s0,a0
lui   v1,$00ff
```

Christer,
Neversoft

GTE MVMVA bug

Written by Christer Ericson on Wed Feb 12 21:14:24 1997

Calling MVMVA (MVMVA sf,mx,v,cv,lm) with cv = 2 is supposed to specify FC as the translation vector used, but this is not the case on our development boards. Instead what happens is that there is no translation, and seemingly the first column products are not included in the result. (I haven't verified this for all possible combinations of arguments, but at least this is the case for sf = 1, mx = 0, v = 0 or 1, lm = 0).

Is this a known GTE bug? Should we not use cv = 2, or can we rely on the faulty behaviour being consistent and thus use it when it would be advantageous to do so; or, in other words, has the bug been (or will it be) fixed in some newer version of the GTE?

Christer Ericson,
Neversoft Entertainment

Re: GTE MVMVA bug
Written by Rob Vawter on Mon Feb 24 17:14:52 1997

To Christer Ericson:

The following quote is from the lib 3.7 changes.txt file: "Following macros of inline_c.h and inline_o.h were deleted since it was found that the far color vectors could not be used for MVMVA, a standard matrix operation instruction;"

I have confirmed with the Japanese engineers that MVMVA, cv=2 does not function properly, and is a known GTE bug. The GTE hardware has not been replaced, nor is it likely to be replaced in the future.

Also, please note that a typo exists at least in the GTE Command Reference handed out at the conference. The row and column corresponding to cv=1 reads:

pBK

p =R/B/G

it should read:

pBK

p = R/G/B

Re: Re: GTE MVMVA bug
Written by Mike Biddlecombe on Wed Feb 26 01:30:08 1997

To Rob Vawter:

Where can I find this v3.7 changes.txt file? It does not appear in the zipped library archive.

Re: GTE MVMVA bug
Written by Christer Ericson on Thu Feb 27 23:30:22 1997

Mike, if you go to the Files:Programmer Tools section and click on the date link for the lib37.zip-file you'll get a text file which is presumably the changes.txt file. At least it contains the quote Rob gave (thanks for the info Rob!) and much more info as well.

Christer,
Neversoft

Interlaced Displays
Written by Bob Koon on Wed Feb 05 01:33:30 1997

Hi folks...

I'm having some problems getting a SPRT to display in an interlaced (512x480) vsync callback routine. I have other SPRTs showing just fine, but this one in particular doesn't and I don't know why. It's 128x128 (8 bit) and it's semitransparent bit is set (although changing this makes no difference.) The symptoms are it's not displaying at all when it's running. However, when I use a VRAM viewer tool, the sprite is there on every other scanline. The sprite is on another tpage than the one's that are displaying, but I'm changing that properly. I've checked and rechecked all of the SPRT's structure elements, and they all check out. It isn't my GPU overrunning a frame, because it still doesn't display when I don't draw everything else.

My suspicions lead me to the way I 'double buffer'. Sony interlaced examples don't use PutDisp/DrawEnv() calls at all saying they're unnecessary. Also, they use ResetGraph(1) instead of a DrawSync() (albeit after a VSync() which I can't do here).

I use a DrawSync(), PutDrawEnv() and PutDispEnv(), but the only one that seems to make a difference is PutDrawEnv(). If I take the PutDrawEnv() out, all my textures go out the window. Changing my DrawSync() to a ResetGraph() doesn't change a thing.

Any help would be greatly appreciated.

- Bob Koon Vortex Media Arts, Inc. (818)557-2922

Re: Interlaced Displays
Written by Unni Pillai on Tue Feb 11 00:15:53 1997

To Bob Koon:

I hope you have already solved your problem... If not here are a couple of suggestions.. 1) Please refer to `sample/graphics/tmd/tuto0.c` This is a Sony example of interlaced mode with `DrawSync(0)`, `PutDrawEnv`, `PutDispEnv` calls. This works, and also numerous games are up and running on interlaced mode, so we can be assured that the problem does not lie with your code running on interlace mode, and maybe has nothing to do with your double buffering scheme. Maybe this might help... 2) I experienced a similar problem to yours while working on a project utilizing Interlace mode 640x480. I was able to display my polygonal database (textured polys, poly etc) but unable to display my sprites. All my sprite structures were being initialized correctly, and linked to the ordering table at the right places, and the sprite texture pages were being correctly set, and my vram viewer indicated that sprites were loaded at the correct VRAM addresses etc, but no sprite objects were being displayed. Also my profiler indicated I wasn't exceeding 1/60th of a second. I spent lot of time attempting to get my sprites to display correctly, but to meet my deadline I did the following hack, I replaced my sprite structures with `PolyFT4` and mapped my sprites as textures onto the appropriate sized polygons, and viola everything works beautifully. Unni

Re: Re: Interlaced Displays
Written by Bob Koon on Tue Feb 11 22:48:47 1997

To Unni Pillai:

I looked at my code, and I am following method #1 (except for the `VSync(0)` of course). Also, what's more concerning is that there's no difference between me drawing a `POLY_FT4` or a `SPRT`. :-/

I've done some tinkering and discovered the following:

When I don't do the 2 `Put...()` calls, the sprite (or poly) shows up, when I keep them in, it doesn't show up. Also, the fact that it's a `vblank` callback routine doesn't change anything.

On another screen, (following the callback routine's screen), things go out the window when I `_don't_` put the 2 `Put...()` calls in the double buffer code. :-/

I'm in a real bind over this one because it's such an unual problem, so any help would be grateful.

- Bob Koon Vortex Media Arts, Inc. (818)557-2922

Re: Interlaced Displays
Written by Jay Stelly on Fri Feb 14 16:55:31 1997

To Bob Koon:

The fact that there's no difference between the behavior of POLY_FT4s and SPRTs eliminates some possible problems... It sounds to me as if your ordering table is getting corrupted. Are you sure that the GPU is executing the sprites in question, yet failing to put them in the buffer? Why don't you try drawing them directly with drawPrim() and see if that works - if it does, it's likely that your trashing the OT later. Or try walking the OT after it's built and find the SPRTs. Just a thought.

Jay Stelly
Tetragon

Re: Re: Interlaced Displays
Written by Bob Koon on Fri Feb 14 20:27:31 1997

To Jay Stelly:

When I use my VRAM viewer tool, the sprite is being displayed (although it is every other scan line of it) which leads me to believe that my OT isn't being trounced. I've not tried drawPrim() or OT walking yet, so I'll do that today I suppose.

Thanks for the reply.

- Bob Koon
Vortex Media Arts, Inc.

Re: Re: Re: Interlaced Displays
Written by Mike Fulton on Tue Feb 18 21:56:31 1997

To Bob Koon:

Using the PutDispEnv() and PutDrawEnv() functions IS unnecessary when in interlaced mode because these functions are used to switch the drawing rectangle and the display rectangle. In interlaced mode, you don't switch so these functions aren't necessary.

The problem you're seeing is due to the fact that you are attempting to draw information into the same scanlines which are currently being displayed. By default, this is disabled so you see nothing.

You can change this by setting the "dfe" flag of your DRAWENV structure to 1 when you first initialize your display. This will allow the GPU to draw into the scanlines being displayed.

However, this will demonstrate a potential problem. It will now be possible to get screen glitches if you're updating scanlines at the same time they're being displayed. This can result in a "tearing" appearance on screen. The only thing you can do about this is try to speed up or synchronize your drawing operations to the display.

Mike @ SCEA

Written by Jeff Beckham on Tue Feb 04 06:52:04 1997

I'm experiencing a problem in which the dithering feature, which is enabled at reset by default, is turning off after a few runs of the main draw loop. Has anyone had any problem with this happening to them, or do you know why the GPU dithering might be on at one time, and then get kicked off? Even if I explicitly set the dither feature on with the DRAWENV structure during the double buffer swap, it doesn't seem to have any effect. It is on one minute, and when the main game loop starts up, it kicks off and I have no idea why.

Jeff

Re:

Written by Mike Fulton on Tue Feb 18 21:47:06 1997

To Jeff Beckham:

As Mark Lyons pointed out, you need to set the dither flag in the primitives themselves. Look into the use of the DR_MODE primitive, which is initialized via the SetDrawMode function.

Mike @ SCEA

Re: Dither
Written by Mark Lyons on Tue Feb 04 18:32:29 1997

To Jeff Beckham:

If you are sending any DR_TPAGE primitives in your OT table they overwrite the dither mode.

The mode must also be set in these primitives when they are built.

Mark

Texture Cache

Written by Joel Dinolt on Fri Jan 31 20:47:00 1997

I was wondering about the limitations of the texture cache. For a given type of texture (16-bit, 8-bit, 4-bit), there is a U,V size limitation for caching, eg 8-bit textures have a 32*64 size. Does this apply to the total area of the texture? If I have a 16*128 8-bit texture, does this get cached, or do I have to lie within both the U and V boundaries? Thanks...

Re: Texture Cache
Written by frank on Mon Feb 03 20:14:01 1997

To Joel Dinolt: To stay on cache, you must stay within the specified boundaries.

Frank Sheperson

LoadImage
Written by Jeff Beckham on Fri Jan 31 20:14:08 1997

I'm sorry...over enthusiastic beginner mistake...I forgot to take into account the header.

Jeff

Re: LoadImage
Written by Mike Fulton on Tue Feb 18 21:41:12 1997

To Jeff Beckham:

See the other message I left you for information about USING the TIM header instead of just skipping past it.

Mike

LoadImage
Written by Jeff Beckham on Fri Jan 31 19:48:40 1997

Hello,

I have been programming just a few days for the PSX, so this maybe a stupid question, but here goes.

I am simply trying to display a TIM file straight from Main Ram to Video Ram with the LoadImage routine. When I do this, however, the origin of the TIM when being displayed is incorrect. Part of the right side of the TIM is being displayed on the left side of the screen as though the routine is starting at the incorrect memory address. I set up my display environment as 256x240, and the image is put into memory at 0x80100000 with the pqbload utility. I pass RECT offsets = 0, and the address of the image data to LoadImage. Even when I change the offsets to something other than 0, the image will still be displayed with part of the right hand side on the left side of the screen. I suppose this will be a long year of frustration, if this is any indication.

Jeff

Re: LoadImage

Written by Mike Fulton on Tue Feb 18 21:39:42 1997

To Jeff Beckham:

Don't forget that you can use the OpenTIM and ReadTIM functions to extract information from the TIM data and make it easier to deal with. The following code extracts the appropriate information from a chunk of TIM data necessary to download the TIM to the portion of VRAM specified within the TIM file:

```
long tim;
TIM_IMAGE theAnalogTim;

tim = OpenTIM( (u_long *)analog_tim );
ReadTIM( &theAnalogTim );
LoadImage( theAnalogTim.prect, theAnalogTim.paddr );
```

colour multiplied textures

Written by Michael Koziniak on Tue Jan 07 00:00:43 1997

From: Dylan Cuthbert
Note: Moved from FirstClass

Hi,

Is there anyone out there that has managed to match the texture shading function of the PSX with just regular non txmapped polygons?

For example, if I have a texture that's gray and shade it so that it's red, then put a red non-txmapped polygon next to it, how do I calculate the correct red value to give it to match it as closely as possible?

The main problem lies with Alias, it uses a different kind of colour mapping when displaying it's texture maps, does anyone know precisely what the difference is between Alias's shading function and the playstation's?

Dylan Cuthbert
Senior Programmer
Sony Interactive Studios America

Poly subdivision and PMDs

Written by Michael Koziniak on Mon Jan 06 23:51:31 1997

From: Scott Osborn
Note: Moved from FirstClass

We're running into some texture "warping" with our PMD-generated polys, and are wanting to divide them using the Sony Divide...() functions. At this stage, we're using the RotPMD..() functions for our rendering, which obviously will have to change to get around the warping.

A few questions:

1) Are there any references available for the RotPMD..() functions? Anything (including psuedo code) would be fine.

2) Are there any inherent problems with using the Divide...() functions with PMD data? We've been unable to get the Divide...() functions to work properly using our PMD data.

In advance, thanks for the input.

Scott Osborn

Collision Detection

Written by Michael Koziniak on Mon Jan 06 23:43:53 1997

From - Dane Emerson

Note: Forwarded from FirstClass

I am attempting to detect collisions between my hero and a maze of walls. There are no restrictions to the width, length, position, or direction the wall faces (there are no horizontal & vertical restrictions... that would be easy!) For ease of presentation, let my hero be a single point in space represented on a single Y plane as (x, z) or (Px,Pz). Points A=(Ax, Az) & B=(Bx, Bz) are the endpoints of each wall in the loop. 'distance' is the final value calculated to determine the hero's distance from each wall in the loop. When distance is 0, a collision has occurred. All formulas were derived from the well-known Graphics Gems book by Andrew S. Glassner (pp. 9,10).

I pose the following questions:

- Without regard to the overhead of such a process, is this completely over the top of what I should be doing? Perhaps using bounding rectangles and keeping it simple? (How is that possible?) -WITH regard to the overhead, how can I simplify this formula and optimize each operation? I am fairly new to most any optimization technique, meaning I am open to any suggestions.

```
int HeroCollision()
{
int Ax, Az, Bx, Bz;
int Px, Pz;
int Vx, Vz;
int Nx, Nz;
int Qx, Qz;
int t;
int distance, dist;
int i;
for(i=0;i<scn.nWalls;i++)
{
```

```

Px = salinas.cord[0].coord.t[0];
Pz = salinas.cord[0].coord.t[2];
Ax = scn.wall[i].x1;
Az = scn.wall[i].z1;
Bx = scn.wall[i+1].x1;
Bz = scn.wall[i+1].z1;
dist = SquareRoot0((Bx-Ax)*(Bx-Ax) + (Bz-Az)*(Bz-Az));
Vx = (Bx-Ax)/dist;
Vz = (Bz-Az)/dist;
Nx = -Vz;
Nz = Vx;
t = ( (Nx*Px+Nz*Pz)-(Nx*Ax+Nz*Az) ) / (Nx*Nx+Nz*Nz);
Qx = Px-Nx*t;
Qz = Pz-Nz*t;
distance = SquareRoot0( (Qx-Px)*(Qx-Px) + (Qz-Pz)*(Qz-Pz) );
if(distance == 0)
return 1;
    }
return 0;
    }

```

Re: Collision Detection
Written by Brett Bibby on Wed Jan 15 15:30:26 1997

To Michael Koziniak:

If your hero moves, you may encounter problems when it completely misses the wall by skipping over it? It would be easier to check a line penetrating a plane (line segment defined by start and end coords, wall defining the plane), which has the nice advantage of being able to tell you where the line penetrated, thereby telling you how far to move the hero. You can find the calc in most advanced math books. In a racing game we sorted all the walls by x,y,z making it easy to quickly find which walls needed checking. With a lot of walls, I would recommend some sort of BSP or other partitioning alg.

_96_init

Written by Michael Koziniak on Mon Jan 06 23:32:27 1997

From - Dave Ross

Note: This was forwarded over from first Class

What would cause _96_init to hang up and not return? We're using Load()and Exec(). All recommend support code has be installed and sometimes the call hangs up in the parent process. I get the feeling from all that I have read, that Sony recommends Load and Exec should not be used. If this is the case why even have it included as part of the Library?

Forum: LIBGS

Rotation

Written by Jeff Beckham on Fri May 23 17:51:46 1997

This is probably a hugely ignorant question, but I want to learn. Using Gs calls, how would i rotate an object about a runtime centerpoint within the object, and not around its original centerpoint? Because the rotation axis must change during runtime, I imagine I must translate the new centerpoint of the object to 0,0,0 world, rotate, and translate back, but I continue to fail in my attempts to do this. Any examples, or help?

Thanks,
jeff

Re: Rotation
Written by Michael Koziniak on Fri May 30 00:31:08 1997

To Jeff Beckham:

I believe it is possible to do what you want with the aid of another coordinate system. Make a parent coordinate system for the object and then translate between the two until you get the desired effect. What do you think?

Michael Koziniak

lib3.7 libgs

Written by Jose Villeta on Wed Feb 26 18:39:30 1997

While upgrading to libgs 3.7 from the web page, we ran to a problem with GsSortObject4. Supposedly, the new libraries were supporting backfacing. Nevertheless, we got nothing when backfacing?

Any news about it.

Jose

Re: lib3.7 libgs
Written by frank on Thu Feb 27 20:10:47 1997

To Jose Villeta:

Did you try using GsSortObject4J? That's the one you want for backfacing. Maybe your message is a typo? Either way, let us know.

Frank

Re: Re: lib3.7 libgs
Written by Jose Villeta on Mon Mar 03 19:06:51 1997

To frank:

We tried both GsSortObject4 and GsSortObject4J without success. We are using libgs 3.7 dated 12/20/96 from the psylib info. We are also using the same convention of setting the attribute to GsNBACKC to turn it on/off. Let us know if we have old modules or need extra settings.

Thanks,
Jose Villeta

Re: Re: Re: lib3.7 libgs
Written by frank on Tue Mar 11 17:39:22 1997

To Jose Villeta:

The GsNBACKC attribute is unsupported. In order to access the backfacing functions (GsTMDfastF3LB, et al) you must use GsSortObject4J() and register the functions you need in the jump table. Plain vanilla GsSortObject4() doesn't access the new functions.

However, even when used correctly, the new functions don't seem to work right. The backfacing polys are still drawn, but some of the FRONT facing polys may get clipped. Is this anything like you've experienced?

I'll E-mail Japan and see what info they have to offer on this. Stay tuned...

Frank

Re: Re: Re: Re: lib3.7 libgs
Written by Jose Villeta on Thu Mar 13 17:48:41 1997

To frank:

We got a running version of GsSortObject4J running with the according jumtable hooks. One thing that we run into is that you can't have automatic subdivision with backfacing. Also, you need to change hook function calls before GsSortObject4J to support backfacing or not. I'll be looking into the clipping problem you refer. I have not run into it. Let me know any news about it??

Jose

Re: Re: Re: Re: Re: lib3.7 libgs
Written by frank on Thu Mar 13 18:41:28 1997

To Jose Villeta:

The clipping problem was me being dense. The OT length in the code I was using was puny.

Bumping it up made everything work fine.

Frank

Re: Re: Re: lib3.7 libgs
Written by frank on Tue Mar 04 20:45:30 1997

To Jose Villeta:

I'll look into this some more and get back to you.

Frank

Welcome to GS
Written by Michael Koziniak on Wed Jan 08 18:55:56 1997

Please provide us with your questions regarding our Extended Graphics Library here.

Michael Koziniak
Development Support Engineer
Sony Computer Entertainment

Forum: Europe/Japan Development

Euro AV cable issues...

Written by romesh prakashpalan on Mon Mar 31 19:37:20 1997

Hello,

We are having one last problem with the port to PAL from NTSC for Virtual Pool. According to Interplay/UK when they hook use a Euro/AV cable there is a "2-pixel multicolored column of pixels" appearing on the right hand side of the screen. Any ideas of where this might be coming from? Our textures ARE stored at the 640th column, and we are in 640x256 screen mode could this be where the problem is? Anyone ever hear of this problem?

Thanks,
Romesh

Re: Euro AV cable issues...
Written by Mike Fulton on Wed Apr 02 00:05:08 1997

To romesh prakashpalan:

What parameters are being specified for the 'screen' rectangle of the DISPENV structure being set? The 'w' field fo the RECT should be no more than 256 (the values do not correspond to pixels). This value controls how long the PlayStation video hardware continues to grab data from VRAM and display it.

You might try setting screen.w to lower values to see if the problem goes away. Try 256 and work your way down.

Mike

Viewing Kanji on US Win-95

Written by David Crane on Thu Mar 27 18:44:07 1997

I am trying to view source files created under the Japanese language Windows 95. These files contain mixed text: English language routine names, standard 'C' structure, and shift-JIS Kanji strings (and comments). All I want to do is view the Kanji on screen (and possibly printer) within the English language version of Windows 95. What I have tried so far:

1. Installed Internet Explorer and added Explorer's Japanese Language Kit. Converted the .c file to a .htm, adding html tags to make it appear to be a valid html file. No luck.
2. Attempted to add Kanji font from Japanese version of Win-95. The .ttf files are rejected as improperly formatted, and the .fon files are not recognized.
3. Attempted to change the DOS code page to 932. I apparently do not have a country.sys with that code page available. (Not even the one off the Japanese SDK CD.)

Other suggestions I have received include:

Format hard drive

Install Japanese Win-95

Rename various files (win.com etc.)

Install English language Win-95

Then have the option of jwin.com or win.com Even if this worked, and I dodged the obvious version-skew problems of having two incompatible versions of Windows on my system at the same time, I would still have to be able to read Japanese to run the jwin version.

If anybody knows of a simple way of viewing shift-JIS encoded text, please let me know.

Thanks.

Re: Viewing Kanji on US Win-95
Written by Mike Fulton on Thu Mar 27 19:59:10 1997

To David Crane:

I'm afraid that the only machines here at Sony where we've viewed Kanji are those with the Japanese version of Windows installed.

You may consider trying the HTML thing again using Netscape instead of Internet Explorer, but aside from that I don't have any ideas.

Mike @ SCEA

Boot from OSD fails on PAL

Written by Rob Vawter on Wed Feb 19 17:58:09 1997

chaplin:

We're doing a PAL version of Virtual Pool, and when the program boots up from the memory card/CD player screen, it hangs during the memory card initialization. Is there a difference in how the memory card expects to be initialized for PAL boxes? Mind you, this is a blue box, not a production unit but I doubt there would be a difference between the production and blue boxes in this matter. Unfortunately, it doesn't even get to doing anything with memory card blocks. It hangs during the initialization of the memory card events! Basically, when the user boots up from the memory card/CD Player screen, then closes the CD tray, the program hangs during my memory card

initialization. It's pretty standard code, taken from the sample programs and works perfectly for the NTSC versions.

vawter: This problem has been reported a few times before. In the past, it has always been a problem with initialization. Calling InitCard(1) (which is in the sample code between opening the events and enabling them) before controller initialization is the most likely cause of the problem. Let me know otherwise.

Memory cards & PAL Playstations...

Written by jens on Thu Feb 06 22:11:47 1997

Hi, I (seem to) have a problem with my memorycard code when I have an unformatted card in the machine at power-on. Sony Europe tells me that when my title tries to save to such a card, it gives the "save error" message instead of the "unformatted memory card" message.

So far I could not reproduce that phenomenon on 3 blue PAL debugging machines with either brand-new or software unformatted cards.

Does anybody out there know of this problem? SCEE says that they get it 100% reproduced and I'm a little confused here.

Any help appreciated,
Jens

Re: Memory cards & PAL Playstations...
Written by Rob Vawter on Tue Feb 18 23:04:32 1997

To jens:

I'm awaiting info from SCEE on the exact nature of the problem. I'll let you (and all other interested parties) know hopefully tomorrow (if you haven't already solved this)

PAL Development - Memory Cards

Written by Rob Vawter on Sat Feb 01 01:34:56 1997

Forwarded from the BBS:

Romesh: We're doing a PAL version of Virtual Pool, and when the program boots up from the memory card/CD player screen, it hangs during the memory card initialization. Is there a difference in how the memory card expects to be initialized for PAL boxes? Mind you, this is a blue box, not a production unit but I doubt there would be a difference between the production and blue boxes in this matter. Any help would be appreciated.

Rob: There is a difference in the format of memory card file headers. If you display more than one icon for your U.S. saved games (i.e. type 0x12 or 0x13), the UK machines won't read the file properly. Either remove the animation from the save file or set aside multiple file blocks (2 icons needs 2 file blocks, 3 needs 3). If you try to use one file block for a type 0x13, the following two file

blocks will be overwritten. It sounds like you have a saved file with multiple icons, but only one file block.

Forum: CD Gen

CD Cutters

Written by Bob Koon on Fri May 02 06:20:25 1997

I seem to remember a while back that there was talk or information about another CD cutter that can burn PlayStation discs other than the 900W unit. Is this accurate? And if it is, is it available from Sony?

Has anyone else used a 'regular' cutter to make PlayStation discs? How involved was the process?

What I'm after here is a way to cut discs at a rate other than 1x speed. 8-/

Thank you for the information.

- Bob Koon

Re: CD Cutters
Written by Brian Leake on Wed Jun 04 04:45:02 1997

To Bob Koon:

You could try the 921S, which we just got from SCEA. This is the other 'official' burner. It is a lot cheaper than the 900W, and works under Win95. I get away with 2x too, which I'm not sure it's supposed to do.

- Brian.

Re: Re: CD Cutters
Written by frank on Wed Jun 04 15:35:09 1997

To Brian Leake:

I'm not sure what you mean when you "get away" with burning at double speed. Both burners are capable of double speed operation. The 921 is the more accurate of the two in that mode, in fact masters may be burned at double speed on the 921.

Frank
SCEA

Re: CD Cutters
Written by Mike Fulton on Fri May 02 18:11:51 1997

To Bob Koon:

The CDGEN program that is used for mastering PlayStation CDROM discs only works with SONY CD-R machines. Version 1.3 will work with either the 900 or the newer 920/921 model. (These are internal/external versions of the same drive.)

The 921 is not the current model, but is still available from several mail-order houses as well as from SCEA directly.

Note that the 900 cannot be used with CDGEN 1.3 & Windows 95 because it does not conform to the ASPI specification closely enough to work with the Win95-compatible SCSI driver.

You cannot use a CD mastering program other than CDGEN to make a bootable PlayStation disc, because nothing else knows how to write out the special information that makes the PlayStation recognize it as bootable, such as the license area file and the region code.

The 900 can burn discs at double-speed. The restriction is that cannot make discs for submission at double-speed because the error rate is a bit higher. However, you can use the other burners at double-speed. You'll have to check the SCEA Technical Guidelines for the exact information

Mike @ SCEA

P.S. For the swiftest response, please always remember to direct your messages to the appropriate message area.

Re: Re: CD Cutters
Written by Mike O'Connell on Tue May 06 04:03:05 1997

To Mike Fulton:

Must the 900/920/921 be used? I remember reading that the earlier betas of 1.3 added support for the "Spessa" CD-ROM drives.

BTW, if I'm not losing my mind, do you know if that support extends to the current model?

(960/9611)

Re: Re: Re: CD Cutters
Written by Mike Fulton on Tue May 06 19:49:25 1997

To Mike O'Connell:

I believe the 920 and 921 did go by "Spressa". CDGEN 1.3 added support for these drives.

I think the 924 and later drives may work with CDGEN 1.3, but they are not approved for mastering purposes. I haven't seen the 960 or 9611 drives, or heard anything about them.

I've sent off a message to SCEI asking what is the exact current status of this stuff. I'll let you know what they tell me.

Mike

Mike

Re: CD Cutters
Written by Mike O'Connell on Tue May 06 21:51:08 1997

To Mike Fulton:

Ah. I think Sony has a couple different model numbers for the same unit, if 920/921/924 are Spressa models.

I assumed those models were simply updated versions of the 900E (given the price difference vs. a retail Spressa).

Thanks for looking into it, I'll check back regularly.

Re: Re: CD Cutters
Written by Mike Fulton on Tue May 06 22:53:04 1997

To Mike O'Connell:

The 920 & 921 are basically the same drive either for internal mounting or already in a case.

The CDW-900E is a completely different animal, and it's physically much different. It's basically a first-generation CD-R burner, and is probably around 3 times larger than the 921 and other more current drives.

Mike

Re: Re: Re: CD Cutters
Written by Mike O'Connell on Wed May 07 20:19:48 1997

To Mike Fulton:

Ah. Okay, then the 920 is the same, and the 921 is a 9211.

Yup, the 900E is a huge beast - we have several. But now one of our teams needs Win95 support... Please let me know whether CDGen 1.3 works with the 960/9611 (maybe that's a 924?). Thanks!

Re: Re: Re: Re: CD Cutters
Written by Lee Jones on Wed May 28 19:34:41 1997

To Mike O'Connell:

Hi, my name is Lee Jones. I'm kinda in the same boat... I need win95 support, and I have a CDW-900E. So, if you get yours working in Win95, please let me know as well. If have had some success using real mode drivers.

Lee Jones

Using CDW-900E with Jazz Drive?
Written by Lee Jones on Wed May 28 19:19:11 1997

Hi y'all. I'm a little new to this type of stuff, but I'm trying to set up my Sony CDW-900E on a computer, along with a Jazz Drive as well.

Has anyone done this? I cannot get them both to operate at the same time.

I would appreciate any help.

Lee Jones

Re: Using CDW-900E with Jazz Drive?
Written by Stephane Leblanc on Fri May 30 18:13:50 1997

To Lee Jones:

I did use a Jazz drive once. I've setup the Jazz drivers under DOS and used Win 3.11 for running CdGen and it has works flawlessly.

Re: Using CDW-900E with Jazz Drive?
Written by Mike Fulton on Thu May 29 19:27:58 1997

To Lee Jones:

There's no inherent reason why you could not use a Jaz drive at the same time as a CDW-900E, provided you have not selected the same SCSI ID number for both devices.

Keep in mind that the CDW-900E is not compatible with any known SCSI ASPI driver that works under Windows 95, so that burner cannot be used under Windows 95. (There are rumors of an older

Adaptec DOS APSI driver that will work under Windows 95 if you disable the Windows 95 ASPI DLL, but we cannot confirm this. If anybody does get it working, please let us know.)

Mike @ SCEA

CD burning tools

Written by Matt Fairweather on Tue May 27 21:15:29 1997

What do I need to burn a CD. One of my co-workers told me that there is a program called CD Gen that only works under Windows 3.1. Please help. Thanks, Matt

Re: CD burning tools
Written by frank on Wed May 28 19:11:24 1997

To Matt Fairweather:

You need three things: hardware, software, and media.

Your hardware will either be the Sony CDW-900E or CDU-921S write-once drive. The 900 is an older unit, not ASPI compliant, and -WILL NOT- work under Win95. It's also bulkier, and I believe more expensive. The only advantage it has is that it can be daisy-chained to do multiple burns at one time. The 921 is a later model, follows the ASPI standard, and works fine under Win95.

The software will be CDGEN 1.3. This is the only software capable of burning Playstation discs. It's a Windows app, either 3.1 or 95.

The media are write-once (commonly called "gold") discs. They are widely available, under many brand names, lengths, etc. The important thing to note is that master discs for submission MUST be burned on a particular Sony media, the CD-R 71PS.

More info on these and other development tools can be found in the "Products" area of this website.

Hope this helps,

Frank Sheperson
SCEA

Bad Burns

Written by Chris Reese on Thu May 15 16:43:39 1997

I have been trying to burn discs for submission and out of 4 burns I have only had one that is a good burn. The other 3 bad burns have movies that run choppy and the program fails the soak test about 2 hours into it. We have all the necessary read retries, but that doesn't seem to help. One of the discs wouldn't even boot it was such a bad burn. My problem is that the burns are not consistent. I can burn 4 discs, one right after the other, and only one of them came out fine. Any suggestions?

We are running CDGEN from windows 3.1 on a 486 66. The drive it is burning from is an IDE and has been optimized before the burn. All the discs were verified and have been run on two different blue stations. The burns have been made at single speed.

Re: Bad Burns
Written by Mike Fulton on Thu May 15 19:57:33 1997

To Chris Reese:

It's possible that the discs are bad in the first place, before you burn them. That sort of thing is rare, but it does happen from time to time.

I don't know if we are aware of any particular batch of discs that have problems, but you can contact Soo Chung and find out from her. She's also the one to contact about getting more discs.

Mike

CCS2CTI with CDGEN 1.3
Written by frank on Fri Mar 21 17:41:25 1997

CCS2CTI was made to work with CDGEN 1.1. If you try to use it with a .ccs file from CDGEN 1.3, you get a "null pointer assignment" message. There happens to be a simple work around for this problem.

Compare the beginnings of .ccs files from the different versions of CDGEN:

```
CD-ROM Generator for Windows 1.10  
Sun Mar 16 13:02:28 1997  
Standard  
0  
XA2  
2048
```

```
CD-ROM Generator for Windows 1.30  
Wed Feb 05 19:30:09 1997
```

```
Standard  
0  
XA2  
2048
```

The difference (aside from the time stamp) is the empty line in the 1.3 version. If you delete this, CCS2CTI will work.

Frank

Unable To Burn Test CDs

Written by Andrew H. Marrinson on Wed Mar 19 20:28:20 1997

We have been having this problem for quite some time and haven't been able to track it down. When we burn test CDs and try to run them on a Playstation they freeze at different points early in the game.

The CDs run just fine on an H2510, so I started verifying the items in the CD FAQ for this situation. I know I have good license data (the game loads and starts.) I am not executing pollhost. We have started running in 2MB even when debugging, and I just verified it runs with SetMem(2) in that environment. We are not using PCOpen and friends, nor are we using open("sim:..."). We do have a large amount of garbage data at the end of the CD.

That leaves accessing the system area in the first 64K of RAM. Is there any easy way to see if this is happening?

What other things could be a problem? We have tried both 74 and 71 minute media. We are using an old Playstation with the disk swapping trick. Would a genuine debug station make a difference? Are there any other gotcha's I've missed?

Help!

Andrew H. Marrinson
NuFX, Inc.

Re: Game freezes up...
Written by Mike Fulton on Thu Mar 20 03:25:37 1997

To Andrew H. Marrinson:

Absolutely do not use a standard PlayStation for testing gold discs. We can't help you with anything related to that sort of situation.

It sounds to me like your code is either running out of memory or it's using a bad pointer and hanging when it tries to access past the end of RAM. Past that I really can't speculate until we know what happens when you are trying this on a real test PlayStation.

Mike

Re: Re: Game freezes up...
Written by Andrew H. Marrinson on Fri Mar 21 13:37:36 1997

To Mike Fulton:

It seems you are quite right about testing with an ordinary playstation. At the same time we sent you the message, we sent a disk that failed to our customer for testing on a proper test station. It worked fine.

Needless to say, we will be placing an order for test stations soon!

Thanks,

- andy

DecDCTinCallback
Written by Dave Marshall on Fri Mar 14 20:00:36 1997

Any idea if SCE is planning on fixing this broken callback?

thanks,
dave

Re: DecDCTinCallback
Written by frank on Mon Mar 17 20:04:20 1997

To Dave Marshall:

I'm curious what makes you say it's broken. This function has been around a long time, and I'm unable to find a single complaint about it anywhere in the BBS archives. The sample program using it works fine.

Are you having some problems using it, or did someone tell you it was broken?

Frank

Re: Re: DecDCTinCallback
Written by Dave Marshall on Mon Mar 17 22:02:21 1997

To frank:

Which sample program uses DecDCTinCallback()?

BBS note in the CD conference dated 1/6/97 from Mike Fulton states that this callback and the Sync routine are broken. I saw the same problem that that original email was asking about.

Currently, I don't seem to get a DecDCTinCallback() unless DecDCTOut() is called. So if you put your call to DecDCTOut() inside your DecDCTinCallback routine, nothing ever happens.

dave

Re: Re: Re: DecDCTinCallback
Written by Mike Fulton on Thu Mar 20 03:29:34 1997

To Dave Marshall:

The message you referred to on the BBS should be a lot older than this past January... that was like last summer or so.

Didn't you also see a message from me to the effect of it doesn't matter because there's not really any overwhelming reason to call DecDCTinCallback? You can't spool in data a little bit at a time like you can read it out, so what's the point?

What are you trying to accomplish that requires this callback?

Mike

Re: Re: Re: Re: DecDCTinCallback
Written by Dave Marshall on Thu Mar 20 20:44:55 1997

To Mike Fulton:

You might need this callback if you were developing a driver to play movies on a textured poly during gameplay. If the DecDCTinCallback() worked, then as soon as I called DecDCTin() I could return control to the game loop, since everything else after that would be run from a callback. Currently, I have to wait for DecDCTin() to finish and then call DecDCTOut() before leaving.

thanks,
dave

Re: DecDCTinCallback
Written by Mike Fulton on Fri Apr 04 02:49:05 1997

To Dave Marshall:

I've been doing some more experimenting and looking through the library source code.

I have discovered that the DecDCTinCallback does work... just not the way you might first expect. I'm getting the callback routine at the end of the decoding process.

When I had posted that original message on the BBS many months back, the thing that had sparked it off was that I was getting lock-ups or timeouts from DecDCTinSync. Now I realize it was because the data transfer hadn't started yet because I hadn't asked for any output yet.

I'm going to ask SCEI to clarify a few things, but it looks like it doesn't actually start sucking in data until you ask it to decode something... so it doesn't finish reading data until the last bit of the decoded data is requested.

For the situation you described, the callback isn't really necessary anyway. Do your DecDCTin() call and then your DecDCTOut call right afterwards. You don't have to wait for anything.

Then you can use a callback for the output if you want. This isn't going to wait any longer than necessary.

Mike

Re: Re: DecDCTinCallback
Written by Dave Marshall on Wed Apr 09 19:48:04 1997

To Mike Fulton:

thanks for the info.

Calc Movie VLC buffer size

Written by Dave Marshall on Tue Mar 11 18:42:12 1997

Is there a way to calculate the VLC buffer size needed to decode a movie using DecDCTvlc()?

Maybe based on the movie size and bit depth?

All the Sony movie demos have this value #defined with a comment that says "not the correct value".

By using DecDCTBufSize(), I can tell that a 64x64x16bit movie needs a smaller VLC buffer than a 324x240x24bit movie. I need a way to create my buffers accordingly.

thanks,
dave

Re: Calc Movie VLC buffer size

Written by Mike Fulton on Tue Mar 11 23:46:55 1997

To Dave Marshall:

The function DecDCTBufSize() is used to determine how much buffer space is required to decompress an MDEC bitstream from the completely-compressed stage to the intermediate stage.

In other words, it tells you how much memory it will take to decompress the Huffman (VLC) encoded data and get back the run-length encoded quantized DCT data that is what you actually send to the MDEC.

Once you've sent the data to the MDEC, the buffer size for obtaining data back from the MDEC is entirely up to you. It isn't necessary to base this on the overall movie size. I suggest you make this one macroblock (i.e. a 16 x 16 pixel block) in length.

When you retrieve your data from the MDEC, you do it a macroblock at a time, and then send that macroblock to VRAM before getting the next one. This has the advantage of requiring a minimal amount of RAM for the decompressed data.

Also, when you do it this way, the movie bit depth is not really that important, since all movies are treated as 24-bit at encoding time. The only thing you have to worry about is when you allocate the buffer for getting back a macroblock of data, since obviously 24-bit data will require 1.5 times as much space as 16-bit data.

Mike

Re: Re: Calc Movie VLC buffer size

Written by Dave Marshall on Wed Mar 12 02:03:29 1997

To Mike Fulton:

If I call DecDCTvlc() with a buffer as small as one macroblock, then it crashes because the buffer is not big enough for what DecDCTvlc() is giving me. So I'm not clear on what you mean or there must be some new library call that I missed.

I guess the buffers my original message referred to was the VLC to DCT buffers.

thanks,

dave

Re: Re: Re: Calc Movie VLC buffer size
Written by Mike Fulton on Thu Mar 13 03:19:13 1997

To Dave Marshall:

I think you're doing it wrong.

The minimum buffer size required by DecDCTvlc is specified by the header on the BS data you pass to this function. The function DecDCTBufSize simply parses the header to return this value.

You cannot decode the VLC data piecemeal. You have to decode the entire bitstream in one step.

The basic steps are:

```
bufsize = DecDCTBufSize( my_bs_data );  
buffer = malloc( bufsize * 4 );  
DecDCTvlc( my_bs_data, buffer );
```

Keep in mind that some of these parameters, such as the buffer size returned by DecDCTBufSize, are the number of LONG words, not the number of bytes.

Please show me your decode routine if this doesn't clear things up.

Mike Fulton

Re: Re: Re: Calc Movie VLC buffer size
Written by Dave Marshall on Thu Mar 13 16:51:07 1997

To Mike Fulton:

```
>>bufsize = DecDCTBufSize( my_bs_data );  
>>buffer = malloc( bufsize * 4 );  
>>DecDCTvlc( my_bs_data, buffer );
```

Thanks, that does clear things up. I've never seen a Sony movie demo do memory allocation in this manner. All the SCE demos use hardcoded arrays. Did the above code sample come from a developer demo?

Re: Re: Re: Re: Calc Movie VLC buffer size
Written by Mike Fulton on Thu Mar 13 19:16:57 1997

To Dave Marshall:

No, the sample code came off the top of my head, and generally it would be a BAD idea to do things this way for Movie playback because you'd be allocating and freeing memory so often that things could get pretty fragmented. I would recommend dynamic allocation for this only when doing still images.

For movie playback, what you need to do is find your worst-case frame and set your buffer size accordingly before you start the movie.

The samples you've looked at undoubtedly use a buffer size that they've determined to be large enough for virtually all movies. And if you think about it, it's not too hard to come up with the right size.

If you have 300k/sec coming off the CD to play a 30fps movie, that's 10k per frame, including audio. The Huffman stage of the encoding typically gives you about 50% compression, so you have to figure that each frame will require in the neighborhood of 20k for the DecDCTVvlc buffer. You might add a little breathing room to that figure and bump things up to, say, 25k or 30k. Or you might actually run through the movie to find out the worst-case buffer requirements.

I wrote a little tool sometime back that goes through a movie file and prints out a few pieces of information about the movie, including the buffer size requirements for each frame. I'll see if I can clean it up a little bit and release it. I don't know what the timeframe on that will be, however.

Mike

Creating subheaders for sectors
Written by Ron Nakada on Wed Mar 12 16:16:15 1997

I think I may have found the way to accomplish what I want. As far as I can tell, I can simply create all of my files using Mode2 Form2 sectors, and create my own subheaders, depending on the data required. The CD-Rom generator will write the error detection code, etc. based on what the subheader is for each sector. I can then write my normal data (2048 bytes/sector) with an 8 byte subheader, and NULL for the remaining 280 bytes, and the EDC and ECC will be automatically generated when I create the disc.

I haven't tried it yet, but I think this will work.

Ron

Re: Creating subheaders for sectors
Written by frank on Tue Feb 25 23:54:34 1997

To Ron Nakada:

You're right, it should work if CDGen behaves as promised. To be on the safe side, you may want to manually build a small test file before you spend too much time writing your tool. Please let us know how things turn out.

Frank

Re: Re: Creating subheaders for sectors
Written by Ron Nakada on Fri Feb 28 18:56:02 1997

Well, it seems to work! Since we already had a tool to build our data file, it was fairly easy to implement the subheaders and NULL padding. I burned a disc and it reads everything fine. Thanks!

Ron

Merging data files w/fmv
Written by Ron Nakada on Thu Jan 23 18:34:55 1997

I'm writing the routines for file handling and fmv. Is there a way to merge Mode 1 files and Mode2 files into one file? I'd like to have all of my data in one file, but this doesn't seem possible. I'm getting the feeling that I'll need separate files(one for game data, and one for fmv). Is this true?

Ron Nakada
Big Grub

Re: Merging data files w/fmv
Written by Mike Fulton on Sat Jan 25 03:13:02 1997

To Ron Nakada:

Don't get hung up on "Mode 1" and "Mode 2" files. You're worrying about that too early in the whole process. That really only becomes important when mastering your CD.

Basically, if a file contains no XA-audio, then it would use standard 2048-byte data sectors. If the file does contain XA audio (such as an STR file with video+sound) then it contains 2336 byte sectors. When you tell CDGEN mode 1 or mode 2, you're basically specifying how much data is already in each sector in the original source file, and how much CDGEN has to generate in order to create the 2352-byte sectors that actually end up on the CD.

The MovPack utility can be done for combining everything together, but you also have to make allowances for combining the data when you compress the movie in the first place.

Basically what you need to do is encode your video with the idea that so much room will be required for other things. At double-speed you have a throughput of 150 sectors per second. If you know you're going to have a certain audio format and so much data to load, you can figure out how many sectors per second to allocate to video, then you can set up your frame rate and compression accordingly.

Once you've got your video data, your sound, and your game data, then you can use MovPack to combine everything together. Have you played around at all with MovPack?

If this isn't enough to point you in the right direction, then let me know and I'll try to take a few days to put together a basic sample and maybe a tutorial.

Mike

Re: Re: Merging data files w/fmv
Written by Ron Nakada on Mon Jan 27 19:53:36 1997

To Mike Fulton:

Since we haven't gotten to the point of burning any CD's, and have only been using the emulator, perhaps I should have posted this elsewhere. Anyway....

Okay - actually, I'm not trying to interleave any data into STR files. Basically, for our game data (2048-bytes/sector) files, we've concatenated the files into one file, with a header sector containing location information on all of the files. This works fine.

What we'd like to do, is have our fmv (2336-bytes/sector) files in this one large file also, to reduce seek times. In order to create this file on the emulator (or to burn it), all of my sectors need to have 2336 bytes.

I've tried using MovPack (ver 1.3e) to convert my normal files. It has a minor problem of corrupting my data (i.e. the 4th character in each sector of a text file changes; a static pop occurs near the beginning of vab file.). Otherwise, it seems to work. If a newer version of MovPack doesn't corrupt the data, then I can expand all of my normal data files, concatenate them with the fmv, and I should be happy....

Or, is there another way of accomplishing this?

Ron

Re: Re: Re: Merging data files w/fmv
Written by Mike Fulton on Wed Feb 12 01:52:02 1997

To Ron Nakada:

Merging a lot of small data files together in order to reduce loading times is a good idea, provided these files are typically loaded at the same time. You avoid wasting a portion of a sector at the end of each file, and you avoid initiating a new read operation for each tiny file.

However, STR files are typically large files that you stream off the disc one at a time. I do not understand how merging multiple STR files together is going to reduce seek times. If you're playing one movie right after the other, then there's no reason to seek. If you're not, you still have to seek to the start of each movie, even if that's within a "file".

Keep in mind that a "file" on the PlayStation is more or less just a convenient means of being able to locate the beginning of a chunk of data. You know that a particular file starts at a certain location on the disc and you start reading from that position. But you can seek to and start reading from any position on the disc, not just the positions corresponding to the beginnings of files.

You're using an old version of MovPack. The current version is 1.5e and is contained in the MOVPACK.ZIP file here on the web site.

Mike

Re: Re: Re: Re: Merging data files w/fmv
Written by Ron Nakada on Tue Feb 18 17:38:30 1997

To Mike Fulton:

Most of our movies will occur at various times between gameplay(i.e. finish a level and play a movie). If the game level's data and movie are not contiguous, the seek times will be longer. Not only would the lens need to be repositioned, but the disc speed would need to increase or decrease. This takes time we want to avoid.

The use of MovPack is bothering me now. It doesn't create the error correction code (as far as I can tell) that is needed for normal game data. This could be a problem. Unless I can create a way of having the cyclic redundancy checks and error correction, I'll have to generate a map of all our data and movies, and use this information for the file handling routines. This would result in a large number of files (and directories) on the gamedisc.

Is there any other way of accomplishing what I'm trying to do? Or am I simply attempting the impossible....

Ron

Re: Re: Re: Re: Merging data files w/fmv
Written by Mike Fulton on Tue Feb 18 22:13:23 1997

To Ron Nakada:

You are correct in thinking that MovPack won't create error correction code for user data, since it doesn't create error correction code for other types of data like video and XA audio. These other data types use error DETECTION but not correction.

You would get an error code if an error occurred on a sector with your game data, but you won't get the automatic retry and correction that would happen with a regular data file. So if you interleave game data into your movies, it is up to you to perform any data-verification and error correction that may be required. If you get an error, for example, you may want to come back after the movie is done playing and re-read just the error sector(s).

Mike @ SCEA

Forum: Development Environment

DTL-H3050

Written by Carl Stika on Thu May 29 01:18:45 1997

I read a response to a question about connecting the blue play station to a pc. The reply was that the DTL-H3050 cable was only used to get printf's from the code on the blue to be sent over to be recieved by a terminal program on the pc. I have tried that but I am unable to get the printf's to come over the cable. I did link the libcomb.lib library in with my code. What exactly do you have to do to get the printf's to come over?

Another question I have is can you get a two way communication going between the blue playstation and the pc. I link in the libcomb.lib library and do all the set up to get the link cable going (i.e. doing Addcomb and opening the streams to read and write.) and I recieve garbage on the pc. I can get two blue stations to communicate fine but can't get any useful info to come over to the pc. Tried all kinds of baud rates. But doing only asynchronous reading and writing, should I got to synchronous reading and writing?

Any help in either of these two problems would be helpful. Thanks in advance.

Re: DTL-H3050
Written by Mike Fulton on Thu May 29 19:31:53 1997

To Carl Stika:

You do not have to link in LIBCOMB to get printf calls to work over the serial cable. In fact, if you've made calls to LIBCOMB, you may have prevented the printf calls from being able to access the cable.

The printf calls are automatically routed to the serial connetion on a test PlayStation. All you have to do is configure your PC serial connection properly and then you should be able to view the output in most PC terminal emulation programs, including the simple "terminal" program that comes with Windows.

Mike @ SCEA

Overlays

Written by Stephane Leblanc on Thu May 22 13:54:17 1997

I would like to know how to use overlays with assembly code. What are the switches for asmpsx ? How do I setup my makefile and my linkfile ? I've red the overlays tutorial from the fall conference but it's only discussing the method for C files.

Re: Overlays
Written by Stephane Leblanc on Mon May 26 19:06:50 1997

Ok, everything seems fine except for one thing. My GTE commands are not transform in my .bin files meaning that an RTPT remain a "dw \$bf". I'm doing the dmplx on my .obj before psylink is executed. What am I doing wrong ?

Re: Re: Overlays
Written by Stephane Leblanc on Tue May 27 02:06:00 1997

To Stephane Leblanc:

Replying to myself here. Seems to me that DMPSX do is job ONLY in the .text section. On the other hand, I can't place my overlay's code in that section. So am I still doing something wrong or do I have to write the COP2 intructions myself.

Re: Re: Re: Overlays
Written by frank on Tue May 27 21:02:30 1997

To Stephane Leblanc:

DMPSX seems to work fine on overlays for me, although I've been using it with c files rather than assembly files. Perhaps you could post or E-mail me the pertinent lines from your makefile.

Frank Sheperson
SCEA

Re: Re: Re: Overlays
Written by Mike Fulton on Tue May 27 20:20:53 1997

To Stephane Leblanc:

I think it would help if you detailed what you're doing to create your overlays.

Also, what version of the libraries and DMPSX are you using? It sounds like it might be an old version.

Mike

Re: Re: Re: Re: Overlays
Written by Stephane Leblanc on Wed May 28 13:32:42 1997

To Mike Fulton:

I'm using DMPSX 3.01 with libraries 3.7 which I think are the latest versions. Comparing my .OBJ before and after DMPSX reveals that the files are identical. Somehow there's something that DMPSX doesn't like in my OBJ.

For now I've wrote my own GTE macro to be able to continue to work but would like very much to go back to the old ones so here how I've setup my overlays:

```
;Overlay1
gpoly group
section spoly,gpoly
;-----
; Overlay2
gpolycf group
section spolycf,gpolycf
;-----
; Linkfile
text group org($80012200)
bss group bss
gpoly group org($80010200),file("z:\usr\jd\overlays\polya.bin") gpolycf group
over(gpoly),file("z:\usr\jd\overlays\polya_cf.bin")

section .rdata,text
section .text,text
section .spoly,text
section .spolycf,text
section .data,text
section .sdata,text
section .sbss,bss
section .bss,bss
```

Re: Re: Re: Re: Re: Overlays
Written by frank on Wed May 28 17:15:47 1997

To Stephane Leblanc:

I think the problem may be in the way you're naming your sections. For example, in your first overlay you put:

```
gpoly group
section spoly,gpoly
```

This sets up a section type that the assembler, as well as other tools(DMPSX) may not know what to do

with. Your overlay modules should have sections named text, data, and so on just as your main module

does. My best guess is that DMPSX parses only the text section, so by giving your section a unique name, DMPSX can't find the changes you'd like it to make. Try changing your section directive to section text,gpoly

and see if that helps. You will also need to get rid of your section directives in the linker command file that refer to the spoly and spolycf sections. These sections are part of your overlays, and as such should not be included in groups of your main program(perhaps this was part of your attempted fix).

Please let us know if this helps.

Frank Sheperson
SCEA

Re: Re: Re: Re: Re: Re: Overlays
Written by Stephane Leblanc on Wed May 28 20:01:06 1997

To frank:

Here's what I get when doing this:

Using section text,gpoly

D:\Sony\jd\src>psymake

PsyMake version 1.12 copyright SN Systems Software Ltd 1993 D:\SONY\BIN\psylink.EXE /n1024 /c /m

/v @linkfile.lnk,.. \exe\main.cpe,.. \exe\main.sym,.. \exe\main.map

PSYLINK version 2.42

POLYA.OBJ : Error : group text redefined as section POLYA.OBJ : Error : Attempt to switch section to

group 'text' Errors during pass 1 - pass 2 aborted Linking completed. 2 error(s) in 1.92 seconds Bad return

code - make aborted

```
;------  
;------  
;------  
;------
```

Using section .text,gpoly

D:\Sony\jd\src>psymake

PsyMake version 1.12 copyright SN Systems Software Ltd 1993 D:\SONY\BIN\asmprx.EXE /l /zd
polya.asm,polya.obj

Psy-Q PSX version 2.25

Assembly completed.

0 error(s) from 3949 lines in 0.6 seconds D:\SONY\BIN\dmpsx.EXE polya.obj
D:\SONY\BIN\psylink.EXE

/n1024 /c /m /v @linkfile.lnk,.. \exe\main.cpe,.. \exe\main.sym,.. \exe\main.map

PSYLINK version 2.42

POLYA.OBJ : Error : Section '.text' placed in two different groups Errors during pass 1 - pass 2
aborted

Linking completed. 1 error(s) in 1.98 seconds Bad return code - make aborted

So on one hand DMPSX can't do his stuff and on the other the linker can't. Any other suggestions ?

Re: Re: Re: Re: Re: Re: Re: Re: Overlays
Written by frank on Wed May 28 21:52:28 1997

To Stephane Leblanc:

Sorry, my mistake. Try

section gpoly.text,gpoly

Re: Re: Re: Re: Re: Re: Re: Re: Overlays
Written by Stephane Leblanc on Fri May 30 18:28:55 1997

To frank:

It works. Thanks.

Re: Overlays
Written by Christer Ericson on Fri May 23 00:12:22 1997

Stephane,

As far as I understand it there are absolutely no differences in overlaying an assembly file compared to that of overlaying a C file. Or, at least, I have overlays working fine including one small assembly file in the overlaid section without having done anything special (I just did what the conference proceedings said to do).

The only "tricky" thing I can think of is that you have to remember to specify the proper group and section directives in your assembly file for the stuff that should be overlaid.

Christer,
Neversoft

Re: Overlays
Written by frank on Thu May 22 22:28:53 1997

To Stephane Leblanc:

You should be able to use ccpsx (with all the necessary switches to build overlays), whether you're using C source or assembly source. I'm not aware of any equivalent switches for asmplx, but I'll see what I can dig up.

Frank

Dev. Studio 5.0 and psymake

Written by Stephen Magladry on Tue May 06 21:52:42 1997

I am trying to set up a tool in Microsoft Developer Studio 5.0 with psymake so I can compile my code from MDS. Currently, When I try to run psymake from MDS, I get an "Tool returned code: -1" message.

First of all is it possible to set up psymake to run within MDS 5.0? If so what could I be doing wrong?

I currently the have following settings for my psymake tool:

Command: E:\dev\ps\bin\psymake.exe

Arguments: (None)

Initial Directory: (None)

Use output window (checked)

I have tried psymake without a directory, relying on the path; The results were the same. I set up the initial directory to pot to the file's path; The results were the same.

If I use psymake from a DOS shell in my development directory, I do not have a problem. Anyone have any ideas?

Please email me any suggestions.

Thanks

Stephen Magladry
smagladry@bigrain.com

Re: Dev. Studio 5.0 and psymake
Written by Mike Fulton on Tue May 06 22:55:12 1997

To Stephen Magladry:

It's not completely clear if PSYMAKE was not being executed, or if it was not returning a valid result. I'll guess the former, but correct me if I'm wrong.

My guess would be that either MS Developer Studio cannot find the PSYLINK program itself, or else PSYLINK isn't finding the makefile when it executes. But I can't tell which is the case from the information you've provided.

Fortunately, I've got a small program which is useful for debugging such problems. DUMPPARM does nothing more than dump a list of the parameters it was passed on the commandline, followed by a list of all the environment variables which it inherited. Then it exits with a code of 0 (meaning no error). Here's some sample output:

Arguments passed on command line are:

```
argv[0] = 'D:\DUMPPARM\DUMPPARM.EXE'
```

Environment variables are:

```
env[0] = 'TMP=C:\WINDOWS\TEMP'
```

```
env[1] = 'winbootdir=C:\WINDOWS'
```

```
env[2] = 'COMSPEC=C:\COMMAND.COM'
```

```
env[3] = 'LMOUSE=C:\WINDOWS\MOUSE'
```

```
env[4] = 'PROMPT=$PSG'
```

```
env[5] = 'TEMP=C:\WINDOWS\TEMP'
```

```
env[6] =
```

```
'PATH=C:\WINDOWS;C:\WINDOWS\COMMAND;C:\HW\CPQ;C:\WINDOWS\NLS;C:\PSX\BIN;C:\NOTES;Z:.;Y:.'
```

```
env[7] = 'COMPILER_PATH=C:/PSX/BIN'
```

```
env[8] = 'PSYQ_PATH=C:/PSX/BIN'
```

```
env[9] = 'C_INCLUDE_PATH=C:/PSX/INCLUDE'
```

```
env[10] = 'CPLUS_INCLUDE_PATH=C:/PSX/INCLUDE'
```

```
env[11] = 'LIBRARY_PATH=C:/PSX/LIB'
```

```
env[12] = 'TMPDIR=C:/WINDOWS/TEMP'
```

```
env[13] = 'CLASSPATH=C:\WINDOWS\Java\Lib;'
```

```
env[14] = 'CMDLINE=WIN'
```

```
env[15] = 'windir=C:\WINDOWS'
```

```
env[16] = 'BLASTER=A220 I7 D1'
```

This information is written to the console device as well as to a text file called "PARAMS.TXT" which will be created in the current directory.

Typically, all you need to do is look at the output of DUMPPARM and you see what's wrong right away. Or if DUMPPARM isn't being executed, then you know the problem is within your MS Developer Studio setup.

The attached ZIP file includes the executable for DUMPPARM. The source and MSVC++1.5 makefile are also included in case you want to make changes. (Please share them with me if you do that.)

What you do is replace PSYLINK with DUMPPARM. First, place DUMPPARM.EXE in the same directory as PSYLINK. Next, you can either just change the name of the program being called from within MS Developer Studio, or you can actually move PSYLINK to a temp filename and then rename DUMPPARM.EXE to PSYLINK.EXE.

Once you've done that, then you just try again from within MSDS. The output should give you a clue as to what's wrong.

Mike

Name: Dumpparm.zip

Content-type: application/x-zip-compressed

Size: 12362 bytes

Re: Re: Dev. Studio 5.0 and psymake
Written by Stephen Magladry on Thu May 08 17:42:24 1997

To Mike Fulton:

Okay, I ran you DUMPPARM.EXE and got the following results:

Arguments passed on command line are:

argv[0] = 'E:\dev\ps\bin\Dumpparm.exe'

Environment variables are:

env[0] = 'TMP=C:\WINDOWS\TEMP'

env[1] = 'TEMP=C:\WINDOWS\TEMP'

env[2] = 'PROMPT=\$p\$g'

env[3] = 'winbootdir=C:\WINDOWS'

env[4] = 'COMSPEC=C:\COMMAND.COM'

env[5] = 'MSINPUT=C:\MSINPUT'

env[6] = 'PATH=d:\Program

Files\DevStudio\SharedIDE\BIN\;E:\DEV\PS\BIN;C:\WINDOWS;C:\WINDOWS\COMMAND;E:\DEV\PS\BAT;'

env[7] = 'PSX_PATH=e:\dev\ps\bin' env[8] = 'LIBRARY_PATH=e:\dev\ps\lib' env[9] =

'C_PLUS_INCLUDE_PATH=e:\dev\ps\include' env[10] = 'C_INCLUDE_PATH=e:\dev\ps\include'

env[11] = 'PSYQ_PATH=e:\dev\ps\bin' env[12] = 'COMPILER_PATH=e:\dev\ps\bin' env[13] =

'GO32=DPMISTACK 1000000 ' env[14] = 'G032TMP=C:\WINDOWS\TEMP' env[15] =

'TMPDIR=C:\WINDOWS\TEMP' env[16] = 'CMDLINE=C:\WINDOWS\SYSTEM\redirect.mod

ERIVHDNIABM' env[17] = 'windir=C:\WINDOWS' env[18] = 'BLASTER=A220 I5 D1 H5 P330 T6'

env[19]

= '_MSDEV_BLD_ENV_=1' _____

First of all, dumpparm ran correctly, so it doesn't seem to be a problem with MSDS. Secondly, all of the environment variables seem to point to the right place.

Any other suggestions?

Re: Re: Re: Dev. Studio 5.0 and psymake
Written by Mike Fulton on Thu May 08 19:20:15 1997

To Stephen Magladry:

First of all, please confirm that you called DUMPPARM exactly as if it were PSYMAKE?

I notice that nothing was passed on the commandline. Is that correct, or is something missing?

By the way, if you're using PSYMAKE anyway, what's the point of doing it from within MSDS? Are you just using the MSDS editor, or what?

Finally, can you send me whatever configuration you've got setup with MSDS so I can try it on my system?

Mike

Blue Playstation

Written by Jeff Beckham on Tue May 06 17:54:19 1997

Recently we bought a DTL-H3050 cable for our blue playstation, but noone here knows anything about it. I recently signed on with this team, so I know nothing about this blue playstation, and noone here seems to know where any documentation is on the blue playstation. They have asked me to setup the blue playstation to work with one of the computers here, but I have no idea of how to get them to communicate. Is this newly released cable for transferring data to the debugging station during runtime, or can it actually be used to download code to the psx machine and run it? Any helpful hints pointing me in the right direction would be appreciated.

Clueless,
Jeff

Re: Blue Playstation
Written by Mike Fulton on Tue May 06 22:44:52 1997

To Jeff Beckham:

The blue test PlayStation is essentially identical to a retail unit with two small exceptions. First, it does not check for the copy protection when booting a disc, in order to allow developers to use CD-R discs they've burned themselves. Second, it does not check the license area code, so it will play discs from Japan or Europe as well as North America.

The DTL-H3050 connects to the "Serial I/O" connection on the back of the PlayStation. The other end connects to your PC serial port.

Any "printf" calls within your code will be sent out the serial port on the blue PlayStation. On your PC you would run a basic terminal package of some kind in order to view the messages. This is the primary purpose for this cable for commercial developers.

This cable is also used to communicate with a Net Yaroze Black PlayStation. But because it is a LOT slower than using the standard dev boards, we don't really recommend it as a good option for commercial development.

Mike Fulton @ SCEA

Re: Re: Blue Playstation
Written by Brian Leake on Sat May 10 21:19:38 1997

Doesn't the Blue PlayStation have 8Meg, compared to the 2Meg of the retail version? Or am I getting confused and thinking of just the development boards?

Re: Re: Re: Blue Playstation
Written by frank on Mon May 12 17:00:01 1997

To Brian Leake:

Only the dev boards have 8 meg.

Frank

More CD emulator woes
Written by Brian Leake on Sat May 03 22:19:16 1997

I read a few of the problems people had been having with the CD emulator, but didn't really see any solution to a similar problem I'm having.

Basically I'm booting from the emulator and then attempting to locate a file using CdSearchFile..., unfortunately my emulator is complaining with "Disk format error in cd_read(PVD)" repeatedly, causing CdSearchFile to fail. I'm not doing anything particularly special, I could have sworn my emulation worked six months ago!

Help anyone?

Re: Even more CD emulator woes
Written by Brian Leake on Sun May 04 03:15:40 1997

Okay, I now run the risk of advising and talking to myself, but I managed to get things a little more 'alive'

by changing the

'Disc CDROM CD.EMU' line from my .cti file into 'Disc XA_PSX CD.EMU' and changed all my file declarations to be XAFiles.

e.g. -

File PSX.EXE

XAFileAttributes Form1 data

Source PSX.EXE

EndFile

This fixed the problems I was having with CdSearchFile() but it's now pushed me on to discover a new problem. I want to put a couple of CD-DA tracks into the emulation image.

My CD-DA declarations are of the style -

Track AUDIO

Pause 150

Source c:\cdda\mytrack.da

EndTrack

And my code has -

```
tracks=CdGetToc(loc); // Get CD-DA TOC param[0] = CdlModeDA; /* CD-DA on */
```

```
CdControlB(CdlSetmode, param, 0); CdControlB(CdlSetloc, &loc[0], 0); /* seek to start of track */
```

```
CdControlB(CdlDemute,0,0); // Unmute CD audio CdControlB(CdlPlay), 0, &result[0]); /* play track */
```

The TOC seems to read okay, and with CdSetDebug(1), things look healthy.

If I run my code directly following a 'run selcd', with an audio CD in the drive, then it does actually work (although there is some popping at the start). If I run following a 'run selemu' and boot from the emulation image using 'run cdexec' then I get no audio, even though there are CD-DA track in the emulation image.

Any hints and tips from those that have done CD-DA on the emulator?

- Brian.
ps. Does it matter that I'm not setting a DMA with cdbios?

Re: Re: Even more CD emulator woes
Written by Dave on Tue May 06 16:57:21 1997

To Brian Leake:

I'll look into this and post some sample code later today... DaveC

Re: Re: Re: Even more CD emulator woes
Written by Dave on Tue May 06 19:00:01 1997

yo brian,

ok, here is my simple sample of playing da either from cd or emulator. Check the cti file, I suspect that you

may not be creating your da files correctly, sorry I couldn't send you any data, its way too big... hope this

helps... Dave

Name: Cddademo.zip

Content-type: application/x-zip-compressed
Size: 8941 bytes

using CodeWarrior and CCPSX interchangeably?
Written by Dave Scheele on Wed Apr 30 15:51:54 1997

We're getting a dev system for our second PSX programmer, and he's used CodeWarrior in the past on his Mac. If we were to get the CW PSX setup, would his source code be compileable by my standard Sony/GNU compiler? If it just needed a few #defines, that would be ok. Is anyone out there using both compilers? Any thoughts?

Dave

Re: using CodeWarrior and CCPSX interchangeably?
Written by Mike Fulton on Wed Apr 30 19:32:03 1997

To Dave Scheele:

It's certainly possible to write code that would work without change on either compiler. The main thing is

avoiding compiler-specific things like "#pragma" or some of the GNU-specific keywords like "__attribute__"

It would be a good idea to go back and forth with the source code on a regular basis to catch problems one at a time as they pop up, rather than discovering a lot of problems all at once. Also, that way you'll get an idea of what things to be careful about.

Mike

ccpsx compiler switch
Written by John Bryant on Tue Apr 22 01:55:58 1997

I know that there is a compiler switch to help stop the 'stack overflow' problem when compiling optimized code on large files. I am not sure what the switch was.

Anyone have any ideas what the switch is?

Thanks

John Bryant

Re: ccpsx compiler switch

Written by frank on Tue Apr 22 16:31:27 1997

To John Bryant:

Just add a line similar to the one below in your autoexec.bat:

```
set GO32=dpmistack 1000000
```

The default value is 256k, set the value as high as you need (memory permitting, of course).

Frank Sheperson
SCEA

No Type Information With Overlays?

Written by Andrew H. Marrinson on Fri Apr 11 19:12:18 1997

We just started using overlays due to the usual running out of space problems. I have just discovered that the data symbols that are in our overlays do not have any associated type information in the debugger.

All the variables look to the debugger like they are of type unsigned long. This is very inconvenient when the type is actually a large struct.

We are using the /v linker switch, and the line numbers for the code are just fine. Any idea what we are doing wrong?

Thanks,

Andrew H. Marrinson
NuFX, Inc.

Re: No Type Information With Overlays?
Written by Ian Schmidt on Mon Apr 14 22:39:01 1997

To Andrew H. Marrinson:

We are getting this problem also...I've discussed it a bit with Andy Beveridge of SN Systems via email. The best workaround right now is to get the latest v4.108 or later PSX debugger, which allows you to typecast expressions in the watch window. Additionally, the type will always be proper if the PSX's program counter is in the source file where the structure is instantiated. I'm trying to come up with a 'trivial' program which has the problem to send to Andy so he can debug things better.

Good luck,
Ian Schmidt

Re: No Type Information With Overlays?

Written by frank on Fri Apr 11 21:56:24 1997

To Andrew H. Marrinson:

Could you possibly e-mail me a copy of your linker command file and your makefile?

Thanks

Frank

Re: No Type Information With Overlays?
Written by Mike Fulton on Fri Apr 11 21:23:52 1997

To Andrew H. Marrinson:

I'll ask Frank and/or Brian to check this out, since they've dealt with overlays the most.

Have you reported this to SN Systems? (EMAIL: support@snsys.com)

Mike

CD library (3.5) and EMU problem

Written by Dedan Anderson on Fri Apr 11 18:58:03 1997

The only way I could get my CD functions to read off of the EMU was to replace them with the old (3.2 i think) libraries. Has anyone else discovered this or have any information on why this solution works?

Re: CD library (3.5) and EMU problem
Written by Stephane Leblanc on Tue Apr 15 13:16:57 1997

To Dedan Anderson:

Make sure that your CTI file specified to build a CDROMXA_PSX disc image. I had that problem when upgrading from 3.3 to 3.5 and this has solve my problems.

Re: CD library (3.5) and EMU problem
Written by Mike Fulton on Fri Apr 11 21:26:18 1997

To Dedan Anderson:

I think Tom Rolfs may have hit on the problem in his message.

However, it would help to confirm this if you can provide more information about your problems.

Mike

Re: CD library (3.5) and EMU problem

Written by Tom Rolfs on Fri Apr 11 19:18:59 1997

To Dedan Anderson:

Are you using the low level CDControl interface? If so then this info might be helpful.

In V3.5 and higher of the CD library all CD high level functions (such as CdRead, CdSearchFile etc...) enable sector header information using CdlModeSize1. If you mix high level commands in with you low level stuff (for example CdSearchFile) then you need to issue a CdlSetMode command to disable header bytes after each high level call. NOTE: Opening the CD drive door also causes header bytes to come back again.

Alternatively you can write your driver with CdlModeSize1 active and either ignore the header bytes or check them to verify the sector is the one you expect.

Refer to the section entitled "Special CD-ROM Notes" in 3.6 and 3.7 of the overview docs for more info.

can link cable hose interface bet. h2000 and cd emu card?

Written by Jim Buck on Fri Apr 11 03:32:09 1997

Nagging hardware problem...

I have been just running the dtl-h2000 cards perfectly fine for a year now (a1360 i11) on a Digital Celebris GL5120st P133. A week ago I started doing development for the link cable using a second similarly configured unit (a1360 i10). Now I made sure to only connect the link cable when both machines were turned off. Ok, everything was fine for almost a week. A few days ago, I installed v1.2 of the cd emu system with a Micropolis external hd (a310 i11 d7, scsi id 2). Ok, now all of THAT worked fine for 1.5 days.

Now, everything works fine until I run selemu.cpe. If I do a resets anytime after I run that, my h2000 cards die. They won't reset and can't run anything due to 'cannot connect to target'. (If I run another cpe that uses the cd immediately after selemu, it fails when it does the CdInit().) For the hell of it, I swapped h2000 cards with someone else in the office, and all works well. So, did something hose the interface bet. the cd emu card and my h2000 cards? (Btw, the same exact thing happened with the second computer as well.)

Also, if this helps at all... before the selemu problem came up, I was starting to get a weird thing happening with respect to the link cable. If the second machine was off, and I was running something on the first machine, the controllers were behaving like they were unplugged. But, if while keeping the program running, I turned on the second computer, the controllers were suddenly recognized as being plugged in (and I could shut off the second computer with the controllers still maintaining their 'plugged-in-ness'). Now, the program I was running had no link code whatsoever.

That's about all. I appreciate any help. I have been going nuts for two days now swapping all sorts of hardware and changing interrupts and memory addresses... ugh.

Re: can link cable hose interface bet. h2000 and cd emu card?
Written by Jerry Jessop on Tue Apr 15 00:28:25 1997

To Jim Buck:

It's not a good idea to run a DTL2000 or 2500 linked to another development system that is off. The primary reason is that the control signals (DTS,RTS,CTS,DSR) are unbuffered unlike a consumer Playstation. Consumer units use some common emitter transistors as buffers on these lines, this also inverts these signals as well. So remember that these signals are unbuffered and inverted vs a consumer unit. I would expect any damage would be on the unit that is off.

Another common failure is to "accidentally" hook up the link cable to a controller port (very easy on the DTL2000). This usually blows one of the integrated circuit protectors on the PCB (controller PCB)F700 or F701. While looking like a transistor these act as fuses and just check them with a meter for continuity. The part is an ICP N38 and cost less than \$1. The video PCB uses ICP N20's. F701 accounts for almost %50 of all DTL2000 returns. There are two circuit protectors on the controller PCB, and 4 on the video PCB.....check them all is my general rule of thumb.

As for your specific problem. It sounds like a conflict in a base address or IRQ setting or that maybe you have one of the gate arrays failed (U42,43 or 730). Verify that your 2000 is set for interrupts off. As others are coming forth with similar problems, I will see what I can do to reproduce your failures on some lab units.

Jerry

Re: Re: can link cable hose interface bet. h2000 and cd emu card?
Written by Jim Buck on Mon Apr 14 23:11:53 1997

To Jerry Jessop:

Thanks for the information regarding both the link cable and my other problem. I will only use the link cable when I have both systems on.

As for the other stuff... it seems that disabling interrupts was the solution. I had placed the jumper across two different IRQ settings, but I didn't realize that that effectively disables the interrupts. Now the only anomaly I have is that all my printf's are WAY slower than they used to be, but I believe this is because I need an IRQ setting for it to work at normal speed, but I can't have an IRQ setting to use the CD emu. Tradeoff?

Re: can link cable hose interface bet. h2000 and cd emu card?
Written by Stephane Leblanc on Fri Apr 11 13:38:38 1997

To Jim Buck:

The same thing happened to me for about a week now. My CD emulator never gaved me any problem before and suddenly it started to do the exact same thing as you mention. If I run selemu.cpe, my DTL-H2500 card die. I can't reset anymore, can't run any program and even if I reinstall the driver, I still can't reset it. To make it work again,I must reboot the computer. I've tried my CD emulator with another DTL-H2500 board and the same thing happened. What's really odd in all this is that I can build an image on my HD using BuildCD but I can't see any of the content of my partitions with CdDisk.

I haven't tried it yet on our DTL-H2000 board, but I'm really thinking about ordering a new CD emulator board unless the cause of the problem is found.

Re: Re: can link cable hose interface bet. h2000 and cd emu card?
Written by Jim Buck on Fri Apr 11 20:05:48 1997

To Stephane Leblanc:

Interesting you having the same thing happen to you with the h2500 dev board. As an update on my problem, I sort of found out what was going on. I looked closer at the h2000 cards a guy here is using with the cd emu stuff. He has a jumper placed across the lower pins of both irq 10 and 11. I did the same with mine and voila! It works. I had mine on just irq 11. But what's weird is that is DID work for a day or two with just a jumper on irq 11. Now, with the second computer I did the same thing, and it still wouldn't work. I used someone else's Micropolis hd with the system, and now it works. I guess the other Micropolis hd I had on there was crap. (But I was able to buildcd and cddisk all I wanted to it, which is strange.) I assumed that it didn't even need a hd attached to the cd emu card in order for selemu to be run though, and that it didn't matter whether the hd connected was good or not (much less connected or not). I just thought I would get errors later on during either CdInit() or CdRead()'s.

So, my new questions is... is it supposed to be that a jumper needs to be placed across two irq's (on the bottom pins, or does THAT even matter) for the cd emu to work?

I still have the problem that with the link cable connected to two dev systems that the second computer needs to be turned on in order for the controller stuff to work on the first system (even if the first computer isn't even running link code).

Where's new CCS2CTI?

Written by John Hamilton on Mon Apr 07 21:30:41 1997

Where is the new version of CCS2CTI? The one that has been needed ever since CDGEN 1.3 started putting the country character in the CCS file. Also, I discovered yet another problem with CCS2CTI. If you don't have all the files for one directory one after another in the ccs file, CCS2CTI will abort with an error message about the file being in the wrong location. So, don't even try to add a file later in CDGEN. Unless, of course, you hand move it in the layout screen to be with the other files for that directory.

Maybe this "feature" could be added to list of problems with CCS2CTI.

Re: Where's new CCS2CTI?
Written by Mike Fulton on Wed Apr 09 21:07:18 1997

To John Hamilton:

I've just EMAIL'ed you a new version of CCS2CTI that we just got from SCEI this morning.

It does not take care of ALL the problems, in particular the new one you've just mentioned. Apparently this is at least partially due to the fact the the CTI format cannot deal with this. To work around this would require much more sophisticated parsing and sorting of the CCS data before creating the CTI file.

I'm not sure if this item is considered a bug, or just part of the way it's designed to work. I've asked what the plans are regarding any changes.

But regardless of that particular item, this version should otherwise be an improvement and it should work with CCS files from CDGEN 1.3.

It will be posted on the web site within a few days.

Mike

Re: Where's new CCS2CTI?
Written by Mike Fulton on Tue Apr 08 18:21:19 1997

To John Hamilton:

As you know, we've been trying to deal with this situation for awhile now and we still haven't gotten anything new from SCEI.

I've fired off yet another message to SCEI about CCS2CTI. However this time we have the advantage that Bill Rehbock and some of our other guys are at SCEI this week, so I've also send Bill an EMAIL asking him to check into it.

And I also included info about the new problem you mentioned.

So hopefully we'll have some news about it later this week.

Mike

sample: OVMENU w/ CDROM 1.8
Written by Dave Miller on Fri Mar 28 17:07:04 1997

I recently upgraded to the Sony libs v3.6 and Psy-Q development tools on CD-ROM 1.8. I have been unable to get the sample \PSX\SAMPLE\MODULE\OVMENU to work with these libraries/tools. This sample works correctly with the tools and v3.5 libraries on CD-ROM 1.7.

The OVMENU program always fails when searching for the file MUTUAL.VB on the CD emulator. The CdSearchFile() statement in file_read() is the exact location of the failure. Searches for other files, such as BALLS.BIN, fail after the MUTUAL.VB failure.

What is going on here? Do I need a new cdbios.com to use the libraries/tools on CD-ROM 1.8? I have have cdbios version 1.00.

Do I need a different snpatch.cpe? I've tried both snpatch.cpe and snpatchw.cpe found on CD-ROM 1.8.

The documentation on CD-ROM 1.8 makes no mention of such changes.

I have the DTL-H2000 boardset and Psy-Q CD emulator board with a 4GB Micropolis hard drive.

Dave Miller
PyroTechnix, Inc.

Re: sample: OVERMENU w/ CDROM 1.8
Written by frank on Thu Apr 03 17:49:37 1997

To Dave Miller:

The problem lies with the sample's file reading code.

When you call CdSearchFile(), a list of files in the target directory is cached. Consecutive calls to CdSearchFile() that access the same directory won't parse the directory structure, but will just grab the info it needs from the cached list.

When MUTUAL.VH is read, CdRead() is reading in two sectors. The file length is only 3104 bytes (as is the destination array), and the extra stuff picked up by the read is tacked on as zeroes. One guess as to where these zeroes end up—on top of the cached file list. That's why CdSearchfile() fails when it goes after MUTUAL.VB.

You say that CdSearchFile() fails for the .BIN files as well. Are they in the same directory as MUTUAL.VH, or in a separate directory as the readme file for the sample suggests? I have had no problems using the structure from the readme file.

I had always assumed that the file list was stored somewhere in lower memory where it wouldn't get trashed, but obviously that isn't the case. It's strange that this sort of thing hasn't happened more often.

There are a few easy work arounds. You can reorder the sound files so that the vab header is read in last. You can change the value of VH_SIZE to 4096. Or you can simply remove the -DSOUND switch from the makefile.

It definitely has nothing to do with which CDBIOS or SNPATCH you're using. Sorry for any inconvenience this has caused.

Frank

Re: Re: sample: OVERMENU w/ CDROM 1.8
Written by Dave Miller on Fri Apr 11 20:26:11 1997

To frank:

I finally got around to reloading Sony CDROM 1.8 onto my machine and revisiting the OVERMENU example. I increased the size of sound_seq[] and sound_vh[] in MENU.C to prevent memory overwrites from trashing the cached CD directory. This fix worked like a charm.

As you speculated, I moved all the required files into a single directory on the CD emulator. This is why the .BIN files could not be found previously.

Dave Miller

Re: sample: OVERMENU w/ CDROM 1.8
Written by Mike Fulton on Mon Mar 31 20:22:43 1997

To Dave Miller:

The \PSX\SAMPLE\MODULE\OVERMENU demo isn't specifically aimed at the emulator, but it should work correctly if you have everything setup correctly.

CDBIOS is now at version 1.04. See the FILES section of the web site to download EMU.ZIP if you don't have this version. (However, this version has been out for awhile and you should have it already.)

Yes, you should be using SNPATCH.CPE. The version on CDROM 1.8 should be OK.

Have you successfully used the emulator with anything else with the newer libraries?

Can you provide your source files for CDBUILD so we can take a look?

What is the sequence used for CDBIOS, DEXBIOS, etc.?

Mike

Stack fault in DOS box

Written by David Crane on Mon Mar 17 22:36:57 1997

When running CCPSX from within a DOS box under Windows 95 the program invariably crashes due to a Stack Fault. This problem goes away after rebooting to Command Prompt only. I have tried adjusting the STACK amount in CONFIG.SYS to no avail. Anyone have any ideas? Thanks.

Re: Stack fault in DOS box
Written by Mark Fong on Wed Mar 19 02:40:13 1997

To David Crane:

In your autoexec.bat, try adding the line, "set GO32=DPMISStack 500000" and then reboot. This is a fairly common problem.

Good Luck,

Mark Fong
Stormfront Studios

Re: Re: Stack fault in DOS box
Written by David Crane on Wed Mar 19 06:14:13 1997

To Mark Fong:

Thanks for the response... works great.

Re: Re: Re: Stack fault in DOS box
Written by Mike Fulton on Thu Mar 20 03:33:14 1997

To David Crane:

By the way, you should also be aware that while setting the DPMM stack size will fix the problem, the reason the problem occurs in the first place is that you've got some source code that is unusually complex, such as a routine that is many, many pages long, or perhaps there are some very large switch statement blocks.

You may wish to take a look at the source code in question and see if it can be simplified. Not only will this make life easier for the compiler, but the code will probably run better because of register allocation and other optimization improvements.

Mike

printf problems revisited

Written by Tom Plunket on Wed Mar 12 00:13:37 1997

I just started (last night) trying to get my emulator up and running. It's the H2000, I guess, the two ISA card one, and I've got the CD Emulator card, too. Actually, the cards seemed to work fine. I was doing some MIME experiments a few weeks ago, and got that boxer on the screen, and some of my own experimentation worked just fine, too. Now, in trying to run our game, I don't see printf messages. The card seems to run ok, the game gets running and all, but I can't figure out why the messages aren't being output like they are on the other devstations that we have here in the office. I've got my PSX cards set to a1340, i11, no DMA, and the Emulator card is at a390, no dma or interrupt (any info into what those are for?), and I have a batch file which is this: (paths deleted for simplicity)

```
dexbios /a1340 /i11  
cdbios /a390  
mess l
```

Then the batch file that runs this whole thing runs snpatch, selemu, then my .cpe file. Can anyone point

to anything that I might be having troubles with? oh yeah, another question: Does anyone know what the dongle is for? I can run the debugger without it, and messages don't start appearing when it's plugged in, I'm just wondering if I need it at all. Thanks, Tom Plunket

Re: printf problems revisited
Written by Tom Plunket on Sat Mar 15 23:00:47 1997

more experimentation yields something interesting...

I swapped devsystem cards, to no avail. However, when I rebooted my machine into MS-DOS compatibility mode, it works just fine. Any ideas on what Windows might be doing? I presume that it could be feeding this data into another driver, but there is nothing that lies even close to the card's address, nor is there anything sharing its IRQ...

help.

- Tom

Re: printf problems revisited
Written by Mike Fulton on Thu Mar 13 03:28:08 1997

To Tom Plunket:

Try compiling this program:

```
main()
{
printf( "Hello World from PlayStation!\n" );
}
```

And then run it after doing this:

```
dexbios /a1340 /i11
messl
resetps
run c:\psx\bin\snpatch
run hello
```

This assumes that "hello.cpe" is the name of the program and that SNPATCH.CPE is located in the C:\PSX\BIN directory. If not, change these sections.

After you've entered these commands, enter the command:

testmess

on the command line. This is the program that actually sits and listens for the printf calls and displays them on your PC.

Once you've verified that this works, then you can try getting it going within the debugger.

Mike Fulton

Re: Re: printf problems revisited
Written by Tom Plunket on Fri Mar 14 20:29:35 1997

To Mike Fulton: Here it is: test.bat (so I wouldn't miss anything?):

```
ccpsx main.c
dexbios /a1340 /i11
mess1
resetps
run /w4 %psxbin%snpatch
run /w4 a.out
testmess
```

resultant output:

```
C:\Temp>ccpsx main.c
C:\Temp>dexbios /a1340 /i11
DEXBIOS version 1.23
Installed OK - card address 1340, Interrupt 11
```

```
C:\Temp>mess1
Psy-Q Message Server TSR ver 1.02 Installed
```

```
C:\Temp>resetps
Sony Playstation DTL-H2000 RESET utility 1.03
Copyright © 1994 S N Systems Ltd.
Done.
C:\Temp>run /w4 D:\PSX\bin\snpatch
RUN.EXE Executable/Binary file downloader version 2.33
Copyright © 1992 to 1995 S N Systems Ltd.
Target 0 is R3000 - SONY_PSX4.00
Loading file D:\PSX\BIN\SNPATCH.CPE
Starting target 0
OK
```

C:\Temp>run /w4 a.out

RUN.EXE Executable/Binary file downloader version 2.33

Copyright © 1992 to 1995 S N Systems Ltd.

Target 0 is R3000 - SONY_PSX4.03

Loading file A.OUT

Starting target 0

OK

C:\Temp>testmess

Psy-Q Message buffer test program - press a key to exit

At this point, I waited about a minute, although I would expect that the message should just pop right up,

right? I could (will) try different devstations in this machine... I know that the card works (it worked in my

old machine), and my current computer is the same (as far as I can see: same sound and video cards, and

nothing else besides the devstation in there) as another setup here in the office. The IRQs do not conflict

with anything, either. -Tom

Re: Re: Re: printf problems revisited
Written by Mike Fulton on Thu Mar 20 03:43:42 1997

To Tom Plunket:

I didn't anticipate your changing the commandlines I asked you to try, or I would have said explicitly not to create a batch file as well as not to add "/w4" to the "run" command.

I use the MSDOS "pause" command in my batch files, following the "resetsps" command, which requires me to hit a key. That way I know it's always going to have enough time to complete.

I'm glad to see things are working OK for you now.

Mike

Re: Re: Re: printf problems revisited
Written by Tom Plunket on Wed Mar 19 20:27:51 1997

Christer and Dave:

Ok, thanks for your help; it's now solved. Somewhere in the documentation I read that the /w switch in run inserted a wait, I took it to mean that it was actually pausing for this time... Plus, I didn't want

to bother figuring out where the wait program was to put it in my search path... Anyways, putting in the wait like I should have in the first place solved this problem...

Thanks,
Tom

Re: printf problems revisited
Written by Christer Ericson on Mon Mar 17 00:29:00 1997

Tom,

The last time we had problems with printf's, it was caused by not having proper delays between resets, downloading snpatch, and running the program. So, your "/w4" could be the culprit.

Try issuing the commands manually (not in a batch-file) and wait at least 10 seconds or so inbetween issuing each command (then if it works, you can trim the delays).

Good luck!

Christer,
Neversoft

Re: Re: printf problems revisited
Written by Tom Plunket on Mon Mar 17 20:28:20 1997

To Christer Ericson:

I'll give it a try doing it manually, however, the only problem anyone here has ever had with improper delays is with "cannot connect to target," and the program completely failing to execute... However, I'll give that a try in any case...

Thanks,
Tom

Re: Re: Re: printf problems revisited
Written by Tom Plunket on Fri Mar 14 21:33:00 1997

To Mike Fulton:

Another point to note is that I got the same results on three other devstation-equipped PCs here, and I know that two of them _do_ show printf(s) during debugging. The third hasn't had much use, so I can't verify it.

Plus, it seems that it could very well not be the card, since I swapped in a different set of cards, and get the same results.

tom

Re: Re: Re: Re: printf problems revisited
Written by Dave Marshall on Tue Mar 18 16:49:08 1997

To Tom Plunket:

The problem you are having is an old problem. Below is my batch file that works all the time. The trick is to put delays between the run calls. A delay is different than a /w10. The /w10 will tell the run command to wait 10 before timing out. It does not cause the run command to stall for 10 if the boards are ready to accept the run command.

For some magic reason, putting in the delays will get you your printf's when running on the emulator.

dave

```
rem boot.bat just does a reset and loads snpatch call boot.bat delay 10 run /w10 c:\psyq\selemu delay 10
```

```
run /w10 c:\psyq\cdexec delay 10 testmess
```

updatecd

Written by Fred Shic on Mon Mar 10 23:12:05 1997

I'm having trouble getting UPDATECD to work. I've heard that some of the problems are hardware related. BUILDCD and CDDISK seem to work just fine.

I'm running a P200 digital with a Micropolis drive as an emulator. I'm using BUILDCD v.2.31, CDBOOT v.1.8, and UPDATECD v.1.18. The problem I'm having is that UPDATECD never works at all. I get this error:

A SCSI error was encountered during a read off the hard drive - Error Number 0x0300. This program is aborting, we hate you, etc., etc.

Any information would be greatly appreciated.

Re: updatecd
Written by Mike Fulton on Tue Mar 11 23:49:49 1997

To Fred Shic:

Can you please give us a screen dump showing the command as entered on the command line and the exact response back from the program?

Mike

Re: Re: updatecd
Written by Fred Shic on Wed Mar 12 06:42:12 1997

To Mike Fulton:

This is my command line entry:

```
updatecd -l usanoxa.map > output
```

the contents of output are attached.

(incidentally, line 183 marks the end of the usanoxa.map file)

Re: Re: Re: updatecd
Written by Mike Fulton on Thu Mar 13 03:31:05 1997

To Fred Shic:

Have you tried the ideas from the other messages? If so, what was the result?

Mike

Re: Re: Re: updatecd
Written by Tom Plunket on Wed Mar 12 19:08:06 1997

To Fred Shic:

I was getting a similar error message when I first installed my CD Emulator and tried to get it to work. CDDisk would give me that error 0x0300, and I couldn't do anything... What I did for experimentation was remove the DMA jumper, and that solved the problem. I have since removed the IRQ jumper, too.

Incidentally, does anyone know what effect the DMA or IRQ settings for the emulator card have? I _do_ know what DMA and IRQ is for, I'm just wondering how it relates to this card.

- Tom

Re: Re: Re: Re: updatecd
Written by frank on Thu Mar 13 01:26:04 1997

To Tom Plunket:

The IRQ is needed if you want to use the CDMON utility. Perhaps the DMA is as well, although it's not mentioned as such in the docs.

Frank

Re: Re: Re: Re: Re: updatecd
Written by Tom Plunket on Thu Mar 13 01:38:10 1997

To frank:

Since the DMA doesn't get a switch in DexBIOS, either, I presumed that perhaps it was for DMA transfer between the CD Emulator card and the H2000, maybe if it's "too far away" inside the case for the cable to reach?

I dunno. I've just disabled DMA on both cards, and disabled the IRQ on the Emulator card...

disassembly command
Written by Mark Harmon on Mon Mar 10 03:03:53 1997

Hello again.

I am trying to find a bug in my program that causes the whole thing to freeze. I have ran the program in the DOS version of the debugger and have located the area of the code where the program always stops.

In the disassembler I can see that the program is stopping on a break command.

It reads:

```
nop  
break $1c00
```

I've included the previous line in case it matters.

My C source code has no break in it.

I compiled the source code to assembly source to see if the break command would be there and it wasn't.

My question is why would there be a break command in the disassembly code?

Please excuse my ignorance in this matter as I have little experience with assembly language.

Re: disassembly command
Written by Stephane Leblanc on Mon Mar 10 14:29:39 1997

To Mark Harmon:

You are simply doing a division by zero which result in causing the break \$1c00.

Re: Re: disassembly command
Written by Mark Harmon on Mon Mar 10 21:28:03 1997

To Stephane Leblanc:

The break is there before I run the program. Is the break put there by the compiler in case of a divide by zero? I'm just trying to understand how this works.

Thanks.

btw. Where could I find information regarding the commands I see in the disassembler? Is there a book you could recommend that I could use as a reference?

Re: Re: Re: disassembly command
Written by Stephane Leblanc on Tue Mar 11 00:54:22 1997

To Mark Harmon:

Yes, you are right, the break instruction is put there by the compiler in case of a division by zero. If you want information about the R3000, try "the MIPS programmer's handbook" written by Erin Farquhar and Philip Bunce.

Re: disassembly command
Written by Mike Fulton on Mon Mar 10 05:16:59 1997

To Mark Harmon:

From what you say, it sounds to me like perhaps the code is becoming corrupted AFTER it's loaded and starts executing.

My suggestion is to examine this portion of the program after the program has been loaded, but before you start running it.

If it's getting corrupted, you may have a bad pointer that is referencing the program code area of memory instead of the data area.

It might also help if you posted the disassembly along with the assembly source output by the compiler.

Just include the 10 lines or so on each side of the problem, not the whole thing.

Mike

Re: Re: disassembly command
Written by Mark Harmon on Mon Mar 10 21:13:30 1997

To Mike Fulton:

Here is a portion surrounding the break command at 8006031C.

```
----- 800602E0 24A40050 addiu a0,a1,$50 800602E4
8FC50024 lw
a1,$24(fp) 800602E8 8FC60018 lw a2,$18(fp) 800602EC 00000000 nop 800602F0 00A62823 subu
a1,a1,a2 800602F4 8FC60000 lw a2,0(fp) 800602F8 8FC70004 lw a3,$4(fp) 800602FC 8CC60010
lw a2,$10(a2)
80060300 8CE70010 lw a3,$10(a3) 80060304 00000000 nop 80060308 00C73023 subu a2,a2,a3
8006030C
00063A03 sra a3,a2,$8 80060310 00A7001A div a1,a3 80060314 14E00002 bne a3,zero,$80060320
80060318
00000000 nop 8006031C 0007000D break $1c00 80060320 2401FFFF addiu at,zero,-$1 80060324
14E10004
bne a3,at,$80060338 80060328 3C018000 lui at,$8000 8006032C 14A10002 bne a1,at,$80060338
80060330 00000000 nop 80060334 0006000D break $1800 80060338 00002812 mflo a1 8006033C
00000000 nop 80060340 AC850004 sw a1,$4(a0) 80060344 8FC50014 lw a1,$14(fp) 80060348
00000000 nop 8006034C 24A40001 addiu a0,a1,$1 -----
- This is before running the program.
```

Thank you for replying. This is driving me nutso.

Re: Re: Re: disassembly command
Written by Mike Fulton on Mon Mar 10 22:01:50 1997

To Mark Harmon:

From the disassembly you included, if this is before running the program, then it looks as though the program contains this break instruction on purpose. To verify that at this point, I want to go a step farther back and look at the EXE file to see what it contains.

If you want to send me the EXE file, I can take a look at it for you. Otherwise, what you need to do is take the address of the disassembly below, subtract the load address of the program, and that will give you an offset into the file.

We also have to account for the EXE file header, and I'm not sure off the top of my head how big that is. So the quickest method is probably for you to jump to the offset, then search from there for the hex string 24a400508fc50024 which is the first two instructions of the disassembly (shown below):

```
800602E0 24A40050 addiu a0,a1,$50
800602E4 8FC50024 lw a1,$24(fp)
800602E8 8FC60018 lw a2,$18(fp)
800602EC 00000000 nop
800602F0 00A62823 subu a1,a1,a2
800602F4 8FC60000 lw a2,0(fp)
800602F8 8FC70004 lw a3,$4(fp)
800602FC 8CC60010 lw a2,$10(a2)
```

```

80060300 8CE70010 lw a3,$10(a3)
80060304 00000000 nop
80060308 00C73023 subu a2,a2,a3
8006030C 00063A03 sra a3,a2,$8
80060310 00A7001A div a1,a3
80060314 14E00002 bne a3,zero,$80060320
80060318 00000000 nop
8006031C 0007000D break $1c00
80060320 2401FFFF addiu at,zero,-$1
80060324 14E10004 bne a3,at,$80060338
80060328 3C018000 lui at,$8000
8006032C 14A10002 bne a1,at,$80060338
80060330 00000000 nop
80060334 0006000D break $1800
80060338 00002812 mflo a1
8006033C 00000000 nop
80060340 AC850004 sw a1,$4(a0)
80060344 8FC50014 lw a1,$14(fp)
80060348 00000000 nop
8006034C 24A40001 addiu a0,a1,$1

```

Once you find the hex string, jump ahead \$3C bytes and look for the “break \$1c00” pattern. If you find it, then you know that it’s being compiled into your code.

And speaking of code, I should have asked earlier, but can you show me the C source for this function?

In fact, you may want to do that before digging into the EXE file.

Mike Fulton

Re: Re: Re: Re: disassembly command
 Written by Mark Harmon on Mon Mar 10 22:23:00 1997

To Mike Fulton:

This is the section of code that contains the break. ----- else {

//----- coming -----/

```

EXIST[line+1][SIDE].length = teki_r - zako_l;
EXIST[line+1][SIDE].count = (teki_r-zako_l)/((zako->vx.val-teki->vx.val)>>8);
    }

```

For now I have changed it to this:

else { //----- coming -----/

```

    {

int tmp;
EXIST[line+1][SIDE].length = teki_r - zako_l;

```

```
if(tmp=((zako->vx.val-teki->vx.val)>>8))
EXIST[line+1][SIDE].count = teki_r-zako_l/tmp;
else {
EXIST[line+1][SIDE].count = teki_r-zako_l;
    }
    }
    }
```

----- This may be kludgy but it keeps the divide by zero from happening (may be reason for break according to Stephan).

Re: Re: Re: Re: disassembly command
Written by Mike Fulton on Wed Mar 12 00:06:41 1997

To Mark Harmon:

I think that the idea that this is a divide by zero trap is probably correct.

However, the question that comes up next is, is it valid for these values to be zero? I mean, is this an indication of some other problem in your application?

Mike

Overlaying Libraries

Written by Nigel Spencer on Thu Mar 06 18:58:30 1997

I'd like to include a library module in an overlay, but need a utility to modify the group/section to allow me to include it in my overlay file. I remember using such a tool on Saturn Psyq libraries some time back, but couldn't see it listed on this site.

Re: Overlaying Libraries
Written by frank on Thu Mar 06 19:59:56 1997

To Nigel Spencer:

The utility is called prefsect. Unfortunately I cannot attach it due some problems we're having with our firewall. You can get it directly from SN's ftp site, or if you have BBS access I attached it to a message in the Development Environment conference along with a sample demonstrating its use back in October.

I'll go ahead and E-mail it to you as well, and then we'll have all the bases covered.

Let me know if you have any problems.

Frank

Re: Overlaying Libraries
Written by Ian Schmidt on Thu Mar 06 19:49:30 1997

To Nigel Spencer:

SN Systems has a tool you use to do this called "PREFSECT". Ask them for info.

Re: Overlaying Libraries
Written by Mike Fulton on Thu Mar 06 19:37:00 1997

To Nigel Spencer:

I believe the information you need was included in the presentation on overlays from last year's developer seminars.

As far as I recall, this was something that you did with PSYLIB, not some other special tool. However, for the details you'll want to look at the presentation. It's available for downloading in the file area.

Mike

Compiler switch for automatic alignment
Written by Mark Harmon on Wed Mar 05 03:48:27 1997

Is there a compiler switch that will automatically align addresses to the correct multiples? eg. multiple of 4 for word-length data or multiple of 2 for half word-length data.

Thank You.

Re: Compiler switch for automatic alignment
Written by Mike Fulton on Thu Mar 06 19:41:20 1997

To Mark Harmon:

Can you give more detail and perhaps some sample source code about what you're trying to do?

As far as I know, any "int" or "long" declarations in your C code should be 16-bit aligned automatically.

Are you having some specific alignment problem? If so, please let us know what is going on.

As far as I know, and I'll admit I'm not the world's biggest R3000 expert, there is no advantage to having data 32-bit aligned.

Mike

Re: Re: Compiler switch for automatic alignment
Written by Mark Harmon on Thu Mar 06 22:04:22 1997

To Mike Fulton:

According to the PSX Hardware Guide Page 13, R3000 requires strict address alignment.

The problem is that I'm porting a humungous game from a dedicated arcade platform to the PSX. There are a ton of include files and I don't want to have to comb through everyone to find out if structures are following the alignment rules. The problem is that if a structure has (for eg.) 3 char variables and then 1 int variable, the int isn't going to be aligned correctly and can cause problems. I'm currently experiencing a lockup that has been giving me a lot trouble. I thought I had it's origin pinned down and then after a complete recompilation of the code the error moved somewhere else. I'm pretty sure it's a memory related problem. I'm pretty close to tracking it down again but I would like to try compiling with some type of automatic alignment to see if this clears up the problem.

Re: Re: Re: Compiler switch for automatic alignment
Written by Mike Fulton on Wed Apr 02 19:32:26 1997

To Mark Harmon:

I was taking a bigger look at the GNU CC manual, and I think the portion that I directed you to may have been talking about options for recompiling the GNU CC compiler itself, rather than options available for compiling your own programs.

I don't know of any other commands or options which would force automatic structure alignment.

Mike

Re: Re: Re: Compiler switch for automatic alignment
Written by Mike Fulton on Mon Mar 10 05:28:18 1997

To Mark Harmon:

Please see the GNU CC documentation, around page 240 or so is some information about various types of alignment.

Mike

DTL-H2500 works with H25DRV.EXE but not H25bios.com

Written by Chia-Ming Wang on Wed Feb 26 17:38:24 1997

We've encountered the following common problem:

1. Users were able to run DECICONS okay and see the BALLS demo running. 2. Users removed the "h25drv.exe" from their config.sys and rebooted.
3. Users ran "pflash" to update the FLASH ROM of the DTL-H2500.
4. Users ran "h25bios.com" in an MS-DOS box.
5. The dos-box locked up and would not receive any more keyboard input.

Key problem is step 3 -- you have to be in the directory "\ps\psyq\bin" (or wherever the pflash.bat is), and you must explicitly type out "pflash.bat", because another file called "pflash" exists in that directory. ("Pflash.bat" floods the "pflash" file.)

This is already in the DTL-H2500 installation manual, but we've updated it to emphasize this common user error. It'll be posted on the WebSite shortly.

Psy-Q tools 8 character file name limit

Written by John Alvarado on Thu Feb 20 18:48:09 1997

Why do the psy-q tools go out of their way to limit file and directory base names to 8 characters? Are they going to fix this any time soon? I know how to get around it, but it sure is annoying! Is this DOS curse going to haunt us forever? I fear that 100 years from now programmers of the future will be programming Holodeck adventures and still have an 8 character file name limit.

- John

Re: Psy-Q tools 8 character file name limit
Written by Mike Fulton on Thu Feb 20 23:33:58 1997

To John Alvarado:

As I'm sure you know, the PSY-Q tools are MS-DOS programs, which have an 8.3 filename length limitation.

Windows 95 and Windows NT support the idea of a "console application" which is a commandline-based application which is mostly similar to a DOS application, except they are expected to do no graphics output, they use native 32-bit code, not 16-bit, and they are "Windows" applications which can use the appropriate functions for long filename access.

We don't have any plans at the moment to ask SN Systems to create Windows 95/Windows NT console application versions of the tools. I'd like to see it myself, but I don't think it's going to happen in the short term future. Of course, you can always feel free to ask SN Systems about this yourself. (email to "support@snsys.com")

Alternately, the Codewarrior package from Metrowerks, while still ironing out a few other problems, supports long filenames just fine.

Don't forget that Windows 95 still maintains 8.3 filenames for files with long filenames. You can still use longer filenames with other tools like your program source editor under Windows 95, as long as you specify the 8.3 filename in your MAKEFILE. (i.e. something like "gamemaincode.c" would be "gamema~1.c")

Mike

Re: Re: Psy-Q tools 8 character file name limit

Written by Tom Plunket on Tue Mar 11 23:58:47 1997

Mike wrote:

>

We don't have any plans at the moment to ask SN Systems to create Windows 95/Windows NT console

application versions of the tools. <

Do "we" think that it would be more involved than just re-compiling the code on a Win32 machine?

Tom

Re: Re: Re: Psy-Q tools 8 character file name limit

Written by Mike Fulton on Wed Mar 12 00:11:38 1997

To Tom Plunket:

Yes, it would be more involved than simply recompiling the program.

As I understand it from my own Windows programming experience, the support for long filenames would require the use of Windows-specific functions and the program would have to register a window class with windows just like a graphics-based Windows application.

It isn't just a matter of linking with a different library. Anywhere filenames are used, the code would have to be changed.

I've heard that SN Systems do have Win32-based executables on their BBS, but they are just recompiled versions and do not support any of the Win32-specific features like long filenames.

Mike

Re: Re: Re: Re: Psy-Q tools 8 character file name limit

Written by Tom Plunket on Thu Mar 13 20:33:47 1997

To Mike Fulton:

- Yes, it would be more involved than simply recompiling the program.

As I understand it from my own Windows programming experience, the support for long filenames would

require the use of Windows-specific functions and the program would have to register a window class

with windows just like a graphics-based Windows application. <

Quite to the contrary, Mike, long filenames are no more Windows-specific than are any of the file functions in DOS. Many of the tools that we have in house are Win95 command-line apps, and I didn't do

anything fancy to enable long filename support. To open a file I just declare
fstream file ("filename", flags);

It's that easy in DOS, UNIX, and Win95. The standard C fopen works the same; the long filename support is inherent in the operating system's file system, so any valid filename would work...

The problem being, of course, that DOS programs don't recognize valid Win95 filenames. I don't know how fancy those guys needed to get, but all of the programs that I've written would compile under DOS if I took out the MessageBox() calls and replaced them with printf(s).

Also, for command line apps in Win95, there's no need to register any windows at all.

- Tom

Re: Re: Re: Re: Re: Psy-Q tools 8 character file name limit
Written by Mike Fulton on Fri Mar 14 02:16:35 1997

To Tom Plunket:

Well, now that I've actually tried a simple test, I guess I was wrong about the long filename support being included in the standard library. I was remembering some article I'd read about using Windows 95 long filenames from a DOS program, and I've obviously forgotten a few things.

Aside from the long filename stuff, however, I'm still not convinced that it would just be a simple recompile to go from the GNU compiler to a Win32 compiler like VC++.

Mike

Re: Re: Re: Re: Re: Psy-Q tools 8 character file name limit
Written by Tom Plunket on Fri Mar 14 19:07:12 1997

To Mike Fulton:

>> Aside from the long filename stuff, however, I'm still not convinced that it would just be a simple recompile to go from the GNU compiler to a Win32 compiler like VC++. <<

Good point.

However, if it is "straight C", then I would imagine that it should be easy enough...

I've sent support@snsys.com email, but never heard anything back...

- Tom

Re: Psy-Q tools 8 character file name limit
Written by Mike Fulton on Thu Mar 20 03:46:26 1997

To Tom Plunket:

Whenever you send mail to "support@snsys.com" you may wish to also CC the message to "andy@snsys.com" and "duncan@snsys.com".

Mike

Re: Re: Psy-Q tools 8 character file name limit
Written by Tom Plunket on Thu Mar 20 19:30:06 1997

To Mike Fulton:

Actually, support@snsys.com worked just fine, I got a response (from another person, Ian), and I've seen that they indeed _have_ ported most of their tools to Win32...

- Tom

h2500 test works, but...

Written by Magnet on Wed Feb 19 21:18:39 1997

Here's my situation: I can use decicons to run balls.exe, but that's the only way my h2500 card will co-operate with me.

I upgraded to the h2500 board. I installed the drivers et al according to the software installation guide.

When I ran freset, I didn't get the numbers that the booklet mentioned. What it returned to me was PCI

```
ver 2.00 Special Cycle 0 config 1 1 bus bus 0 dev=14 func = 0 irq = 12 io = 0009f000 unit = 0 addr 0009f000
```

```
irq = 12 (v = 0074,8259= a0) dipsw 0x00
```

hmmm, looks different but doesn't look fatal.

(following installation guide rev 001.02, p 17)I took h25drv.exe out of config.sys and rebooted. I opened a dos windows on my w95 desktop. I ran flash.bat from c:\psx\bin. I ran h25bios. I locked. I rebooted.

I can run decicons and run balls.exe, but that's about the only way I can anything up to my display.

Just as a test I put h25drv back in my config and stepped through my config output. Seems that it was returning the line "PS25\$0 irq=0" which I also considered strange. I then took h25drv out of config.sys again.

What file have I forgotten to copy/update? What flags have I set incorrectly? What detail have I failed to provide?

Any help would be greatly appreciated.

- jason mceachen

Re: h2500 test works, but...

Written by Chia-Ming Wang on Wed Feb 26 18:49:30 1997

To Magnet:

Are you sure you ran "flash.bat" from "\psx\bin" and then tried to run "h25bios.com"? If so, that's not quite the combo you want to run. Here's a summary of what file works with what driver:

flashb8.bat works with h25drv.exe

pflash.bat works with h25bios.com

----- Are you sure that your IRQ 12 is solely used by the PCI card? If not, you will have to shift other peripherals to other IRQ's using a configuration utility usually accessible by hitting the F10 key after the memory test during a cold boot. (Your system may require a different keyboard character—some computer have you hit ESC, or DEL, or whatever).

----- Do you have an AMI bios? You can tell if you turn off your computer and turn it back on again. The type of BIOS should be printed on the screen. If you have an AMI BIOS, then the memory is being allocated to the wrong location. If so, then we need to change where memory is being allocated for the PCI card. After the computer finishes booting in Win 95, right click on the "My Computer" icon. Right click on it; select "properties". Select the "Device manager" tab. Select the "computer" icon. Double click on "Other devices" then select the "PCI card" icon (which has a question mark and an exclamation point on top). Then click on the button "remove". Shutdown Windows 95. Reboot. When Win95 displays the message "New Hardware found", select "Do not install driver". After Windows 95 finishes rebooting, shut down again. Reboot. Again Win 95 will display a dialog, and choose to ignore whatever actions it asks of you. Go back into the "Device manager" and select the "PCI card" icon. Double click it. Select the sources tab. This may show the memory range of E7000-E7FFF. Clear the "automatic settings" check box. Change the memory to a non-conflicting memory address such as C8000-C8FFF. After modification, Win 95 asks you "Do you want to restart now?" Select "Yes".

Now try the "pflash.bat" and the "h25bios.com".

(This information appears in the DTL-H2500 installation guide Revision 1.007, which will be posted shortly on the Web).

----- Please let us know if you have any problems.

Re: h2500 test works, but...

Written by Stephane Leblanc on Thu Feb 20 14:25:42 1997

To Magnet:

Make sure that your board is compatible with the DTL-H2500. We had that problem. We thought we had a faulty board but before sending it back, we tried it in another computer and it worked the first time.

printf message

Written by simon on Fri Feb 14 20:42:07 1997

I have my new lib3.6 installed, but for some reason in my debugging window there is no print out of message any more. I had my printf working before. I checked to make sure the mess1 TSR installed and testmess seems functional. So why the hell the printf messages I used to have disappear now ? Please advise!

Thanks!

Re: printf message
Written by simon on Tue Feb 18 20:55:08 1997

It turns out it's the Metrowork's CodeWarrior DR2 PSX driver cause the problem. Once I remove the driver, the dbugpsx works just fine now.

Simon

Re: Re: printf message
Written by Mike Fulton on Tue Feb 18 22:15:48 1997

To simon:

Glad to hear you figured out what was happening. Did you also pass this information along to Metrowerks?

Mike @ SCEA

Re: Re: Re: printf message
Written by simon on Mon Feb 24 17:48:23 1997

To Mike Fulton:

I havn't pass this to Metrowerks yet. But I guess I can get by this by reinstall their driver whenever I want to use their IDE compiler. Do you guys know whether there is a 32bit version PsyQ tools, like "run selcd.cpe" from Psygnosis yet? Because that is the major reason I have to switch back to Sony dbugpsx as I have to emulate constantly from cd or harddrive and rebuild my disc image. If I use Metrowerk IDE, I have to drop to dos 16bit and then remove dexbios to switch back to 32bit environment.

Simon@Readysoft

Re: Re: Re: Re: printf message
Written by frank on Wed Feb 26 01:28:18 1997

To simon:

SN has 32 bit versions of the compiler, assembler, and linker available. I think they may have a 32 bit librarian as well, but no corresponding version of the "run" utility exists that we know of.

If you have specific features you would like to see from SN, I'm sure they would like to hear from you.

Just don't make it sound like "hey, could you do XXX so that it's easier for me to use Metrowerks?"

Frank

CD Emulator

Written by Jeff Beckham on Fri Jan 31 21:40:51 1997

I have read that CD emulator is a big headache, and I am finding that out too. Does anyone know if the emulator in addition to not responding and the like, might also provide garbled data? I am trying to load a TIM from the emulation HD, and the same code that displays it fine when using the pqbload utility to get it on the psx board, displays garbled data when using the cd emulator. I cut and pasted the cd code from CD4.PDF [1.6.39] How can TIM files be read in from CD-ROM? I am using the gencti.exe utility to generate the cti file, with the wild.wcd file containing *.tim. Anyone see something I don't?

Jeff

Re: CD Emulator
Written by Mike Fulton on Tue Feb 18 22:21:13 1997

To Jeff Beckham:

Reading TIM files from the CD is not different from reading any other data. However, you aren't using the same method with PQBLOAD and the emulator. One method is "magically" placing the file in memory for you, the other requires that you read it from the CD. Can you post the code you're using for loading the TIM from the CD?

Mike @ SCEA

Re: CD Emulator
Written by Les Hedger on Mon Feb 10 23:24:02 1997

To Jeff Beckham:

I too was seeing what looked like garbled data in my files until I switched the .cti to create XA formatted

sectors. Model sectors just won't work. (See the thread from Jan 9 '97 titled "CD Emulation from RUN.EXE") I'm not using cdgen so I'm not sure what it uses. Good Luck

Les Hedger
PF.Magic

Register local variable
Written by Martin Ross on Fri Jan 31 01:46:44 1997

Is someone know how many "register" local variable GNU can allocate ex:

```
void Didi(void)
{
    register long var1;
    register long var2;
}
```

Re: Register local variable
Written by Phil Weeks on Thu Feb 20 17:02:31 1997

To Martin Ross:

My understanding is that the 'register' connotation only works with debug info on and NO optimizations.

As soon as you turn optimizations on, the compiler decides which local variables go into registers. Maybe someone else can say how many registers the compiler will use, but I would think that two would definitely make it.

Re: Re: Register local variable

Written by Steve King on Fri Feb 21 13:58:10 1997

It's stated in the GNUCC manual that the 'register' keyword is only honored when all optimizations are turned off. (And even then, you're only giving the compiler a hint, not an absolute commandment.) If you turn on any optimizations (the '-O' flag) the compiler decides for itself which variables go in which registers. In my experience the optimizer is pretty good, generating code considerably faster than compiling without it.

If you *need* a variable in a certain register (say, to pass to an assembly-language function) you should be able to load it explicitly with an inline asm() statement.

Re: Register local variable
Written by Mike Fulton on Tue Feb 18 22:27:22 1997

To Martin Ross:

There are a lot of factors which control register allocation for local variables. This information can be found in the GNU documentation (GNUCC.PDF) on your Programmer's CD.

Keep in mind that the "register" keyword is a suggestion to the compiler, not an absolute directive.

Mike @ SCEA

More H2500 woes—printf?

Written by Andrew H. Marrinson on Tue Jan 28 14:03:16 1997

Well, one more problem has cropped up with the H2500. We no longer see our printf's in the debugger message window. We don't get them with testmess either.

I just verified that the same binary does display printf's with testmess on the H2000 board. So it appears to be something about the H2500 bios or hardware.

We really need this to work. Can you help us?

Given the number of problems and questions we have had with/about the H2500 perhaps a FAQ posted somewhere on this web site is in order.

Thanks,

Andrew H. Marrinson
NuFX, Inc.

Re: More H2500 woes—printf?
Written by Jason Curtice on Fri Mar 07 01:34:33 1997

To Andrew H. Marrinson:

After many hours of trying to find out why some of our programs worked with the 2500 boards and some didn't we finally found the problem : libc2. If we link with libc our printf's work. Unfortunately, libc is very slow compared to libc2. Hopefully Sony will fix this ASAP!

Jason Curtice
Westwood Studios

Re: More H2500 woes—printf?
Written by Chia-Ming Wang on Thu Jan 30 01:17:12 1997

To Andrew H. Marrinson:

1. Are you running "mess1.com" in your DOS-box? And is mess1.com already invoked on the DTL-H2000 machine (which would explain the discrepancy).
2. Make sure that the H25drv.exe program is not running. If it's sitting in your config.sys, remove it or comment it out.

To run your program on the DTL-H2500:

1. Run the "h25bios.com" in a DOS-box (under Windows 95).

2. Run "mess1.com" to set up the message handler TSR.

3. Run "resetps" to reset the board.

4. Run your program.
5. Run "testmess.com"
6. Hopefully, you'll see the printf's popping out.

(Steps 1 and 2 are described in the document "2500.doc", in the section "Installing the h25bios.com driver", however the above steps are not explicitly spelled out.)

Does this help? Please let me know...

Chia-Ming

Re: Re: More H2500 woes—printf?
Written by Andrew H. Marrinson on Mon Feb 17 13:24:09 1997

To Chia-Ming Wang:

I replied directly to Chia-Ming that this was exactly what I was doing and got no response, so I'll try again here.

I do exactly the steps you describe and don't get printf's. Do you have any suggestions?

Has anybody gotten printf's to work with a 2500 card?

This all worked fine until I switched to the 2500 and replaced dexbios with h25bios.

- andy

Re: More H2500 woes—printf?
Written by Jeremy Gordon on Sat Feb 22 17:56:16 1997

To Andrew H. Marrinson:

Hi Andy, we've got a couple of DTL-H2500's over here and the printf's work fine. The only wierdness seems to be that we have to printf("\r\n") now instead of printf("\n").

Did you use pflash.bat to install the PSY-Q DTL-H2500 side driver in place of the Sony deci driver? Also I'm pretty sure that the printf's require the DTL-H2500 to talk to the PC via hardware IRQ, when you run h25bios does it report back a valid non-conflicting IRQ?

The Sony Windows 95 driver (the .vxd and two .dll's) definitely screw up the operation of the SN tools in a DOS box, so make sure that you go in and disable them in the control panel while using the SN tools.

Good luck!

- jeremy
jgordon@magicarts.com

Re: Re: Re: More H2500 woes—printf?
Written by Mike Fulton on Tue Feb 18 22:29:44 1997

To Andrew H. Marrinson:

Did you change anything else at the same time? Somebody else noted a problem with printf calls not working when they had installed the driver for Metrowerks' CodeWarrior.

Mike @ SCEA

Re: Re: Re: More H2500 woes—printf?
Written by Andrew H. Marrinson on Wed Feb 19 20:20:30 1997

To Mike Fulton:

I did see the message about Metrowerks. I don't think I changed anything other than going from a 2000 to a 2500 board and upgrading to all the software on the 1.8 CD-ROM release.

When I retested with the H2000 board on another programmer's system, I copied the identical binary and ran on a system with the H2000 board, so it seems unlikely it is something about the software on the 1.8 CD unless it is H2500 specific.

It sounds like you have had success getting printf's out of an H2500 board?

All I can say is both me and another programmer here definitely noticed the printf's which had been working go away when we upgraded to H2500+1.8 CD.

- andy

Re: Re: Re: Re: More H2500 woes—printf?
Written by Mike Fulton on Fri Feb 21 00:01:43 1997

To Andrew H. Marrinson:

I have a 2500 board, but I'm still waiting for the CD mechanism so I haven't installed it yet. I am still using the 2000 board set.

However, we do have other 2500's installed around the office and they have no problems with printf() calls as far as I am aware.

Can you please specify the EXACT sequence of driver loading and so forth? Also, are you viewing the printf() calls with the TESTMESS utility, the debugger, or what?

Mike

Re: Re: Re: Re: Re: More H2500 woes—printf?
Written by Andrew H. Marrinson on Wed Feb 26 15:04:12 1997

To Mike Fulton:

Well, I've made significant headway on this. At first I thought I had totally solved the problem by following Jeremy's advice about making sure I'd run pflash.bat.

It turns out that printf's work until the first time I run resetps, and then never again. If I run h25bios, and mess1 and then run my program, I get printf output with either testmess or the debugger. If I run h25bios, mess1, and then resetps, I don't get printf.

As a workaround, when I need to resetps I find that running h25bios twice (once to remove it, once to re-install it and reset) works.

So, basically this sequence works:

```
h25bios
mess1
run aprogram
testmess
```

but this sequence doesn't

```
h25bios
mess1
resetps
run aprogram
testmess
```

(It doesn't seem to matter if resets 1 is substituted for resets in the above.)

At least I have a workaround now, so I care much less about a complete solution.

Thanks,

- andy

More H2500 woes—printf?

Written by Mike Fulton on Tue Mar 04 02:32:04 1997

To Andrew H. Marrinson:

Glad to hear you've got something working. I'll pass this info along to SN Systems to see if they have any ideas.

Mike

H2080 Control Boxes

Written by Andrew H. Marrinson on Fri Jan 24 20:52:26 1997

Hello,

When we got our K2500 packages we did not receive the H2080 boxes that allow us to connect joypads, etc. When we asked, we were told that these were backordered.

Any idea when these devices will be shipped? It is kind of silly to ship the boards without them isn't it?

I fairness, I should say that this information comes not direct from Sony, but from Electronic Arts (Canada) for whom this project is being done.

So far we have been able to make do with controllers for the old H2000 boards we own directly, but I am about to start work on the Multi-tap support and for that we need these boxes.

Is there any way to get a conversion pin-out so we can attach a DB-9 to the multi-tap we bought and work that way?

Thanks,

Andrew H. Marrinson
NuFX, Inc.

Re: H2080 Control Boxes
Written by Mike Fulton on Wed Feb 12 01:56:13 1997

To Andrew H. Marrinson:

For information about availability for things like this, you need to contact your company's account executive.

Mike

H2500 board w/Win95 Debugger

Written by Andrew H. Marrinson on Tue Jan 21 19:54:00 1997

It appears that the new Windows 95 debugger only supports the old 2000 board. Is this true? Is there any plan to have support for the new PCI board in the new debugger?

Re: H2500 board w/Win95 Debugger
Written by Chia-Ming Wang on Thu Jan 23 22:43:47 1997

To Andrew H. Marrinson:

The Windows debugger is only valid for DTL-H2000 boards. PsyQ is currently working on a version that will work with the DTL-H2500 boards, but they haven't set a definite date when it will be available. Stay tuned...

Re: H2500 board w/Win95 Debugger
Written by Andrew H. Marrinson on Wed Jan 22 17:28:43 1997

Replying to myself here. We found a readme that confirmed that the debugger does not work with the H2500 board. That is very unfortunate. Are there any plans to support the PCI card with the W95 debugger?

Re: Re: H2500 board w/Win95 Debugger
Written by Mike Fulton on Tue Feb 18 22:32:17 1997

To Andrew H. Marrinson:

Yes, of course we want the PCI card to enjoy full support. Keep an eye out for news about support with the Win95 version of the debugger. We do post announcements about important updates.

Mike @ SCEA

H2500 Software

Written by Andrew H. Marrinson on Tue Jan 21 17:29:30 1997

Hello,

We just got our K2500 kits with the new PCI boards and the latest CD-ROM version 1.8. With the boards we got floppy disks with driver updates. We got one with a printed label and one with a hand written label that said run me first.

I'm wondering which software to use. It appears that the CD has the latest stuff, and that the floppy disks were for using the 2500 board with older (<1.8) CDs, but I am not sure.

Can someone at Sony clarify this for me?

Thanks,

- andy

Re: H2500 Software
Written by Chia-Ming Wang on Thu Jan 23 22:40:53 1997

To Andrew H. Marrinson:

You should use the stuff on the CD. I'll notify shipping to discontinue handing out the floppies, since the floppies were update disks for the CD-ROM 1.7.

Sorry for the inconvenience.

CD Emulator Woes

Written by Steve King on Sun Jan 19 17:24:14 1997

I'm having some major trouble with the CD Emulator. About half the time it works fine, and the other half I get DiskErrors and Retries. When it works, it works great. When it fails, it fails completely. It's an all-or-nothing situation.

Here's the details: I'm using an internal Micropolis 4221AV drive addressed as SCSI ID 5. There are no other devices on the SCSI bus, just the drive and the emulator card. Both of these are properly terminated. The emulator card is strapped for I/O port 310, DMA 6, IRQ 12. My PSX board set is strapped for I/O port 1340, IRQ 15. These settings do not conflict with any other device in my system. (I've disabled my PS2 mouse port and my IDE ports, so IRQs 12 and 15 are available.) I'm using Win95 for my OS.

I'm using CDBOOT30.BIN as the emulator's boot program, and I'm using the latest version of the emulator tools, downloaded from here just last week (emu.zip). Prior to downloading the latest, I was using the version found on the Developer's CD 1.7. (Same results.) I'm using version 3.5 of the PSX libraries, and the compiler and linker from the Developer's CD 1.7.

CDDISK never fails to detect and manipulate the drive. However, programs running on the PSX boards fail about half the time. I've tried using both the DOS-based and Win95-based (from SN's web site) loading programs with identical results.

DOS:

```
resetps 1
delay 60
run c:\psx\cdemu\exec\selemu.cpe
delay 10
run c:\psyq\snpatch.cpe
delay 10
run <my_program>.cpe
```

Win95:

```
(start fileserver)
pqrun reset 0
delay 60
pqrun run c:\psx\cdemu\exec\selemu.cpe
delay 10
pqrun run c:\psyq\snpatch.cpe
delay 10
pqrun run <my_program>.cpe
```

IF the CD emulator functions at the start, it'll function fine for the remainder of the program. If it fails initially, I cross my fingers and reload. About half the time it works—sometimes it works the first time, sometimes I have to load two or three times.

Any ideas? This is really starting to tick me off. I've tried running CDMON (with CDBOOT60.BIN), but I don't know how to interpret the output.

Re: CD Emulator Woes
Written by kozman on Tue Apr 15 00:50:02 1997

To Steve King:

Try taking off the termination to see if you find different results. Also, make sure all the SIPS(terminators) on the emulator are plugged in and the dot on them matches up with PIN 1.

Michael K.

Re: CD Emulator Woes
Written by Dedan Anderson on Fri Apr 11 18:47:27 1997

To Steve King:

I was having problems every time I ran CDDISK, the emulation will not work, the drive remains inactive. However, if I go in CDDISK and reload whatever bootbin I am using my emulation magically works! I rarely use CDDISK anymore because of this. Emulations work fine before CDDISK and dont work at all afterwards! Well, hope that helps.

Re: CD Emulator Woes
Written by Simon Freeman on Thu Feb 06 01:27:29 1997

To Steve King:

You might want to try using an older CDBOOT (like 15 or 25) We had similar problems to those you describe with Quantum Fireball 1.28GB drives. After switching to CDBOOT15 the problems have disappeared.

Let us know if this works for you,

Best,

#6

Re: Re: CD Emulator Woes
Written by Steve King on Thu Feb 13 13:50:08 1997

To Simon Freeman:

Actually, I went the other way and got the latest and greatest CDBOOT34 off of SN's ftp site, ftp.snsys.com. This seems to have cured my CD emulator problems. I've had it installed for a week now and have yet to have a single failure. (Knock on a REALLY BIG piece of wood!)

Re: Re: CD Emulator Woes
Written by Simon Freeman on Thu Feb 13 00:32:30 1997

To Simon Freeman:

As a follow up, it turns out that the development system board was to blame, rather than the emulator board or hard drive.

The non-working board had a PD2 chip on it (as mentioned in the manual) however the working board had a PD3 chip.

Whether this is the only problem, or whether it is something else I don't know.

Anyway it's almost usable now,

#6

Re: CD Emulator Woes
Written by Chris Reese on Wed Jan 22 01:15:21 1997

To Steve King:

It sounds like you are experiencing the same problem I have had to deal with for the past year now. I have found that the success of the emulator depends on the board that you have installed. I have three emulator board and all three behave differently. I have one that works 95-99% of the time. I have another that fails about 3 out 5 times and yet another that fails only 1 out of 5 times. All three boards have been connected to the exact same development boards and run the exact same hard drive. After talking with Sony about it and getting the response that the Emulator boards are inherantly flaky I gave up and went to the expense of burning CD's on a regular basis. I only use the emulator when we have data that is rapidly changing. I can definitely sympathize with you because there is nothing more frustrating than going into a debugging session and it takes a good 10 fails before it catches!

Re: Re: CD Emulator Woes
Written by Mike Fulton on Tue Feb 18 23:05:57 1997

To Chris Reese:

Did you try the new CDBOOT34 that Steve King mentioned? I don't think we have this new version on the web site yet, but you can get it from SN System's site in the meantime. See Steve King's message for more info.

Mike

R3000 Disassembler

Written by William B. Norris IV on Fri Jan 17 13:29:23 1997

Is there an R3000 disassembler available?

I'm looking for a command line utility (not in the debugger) which can disassemble R3000 code. An added bonus would be a disassembler which could extract the code from an object file, not just a binary.

William B. Norris IV
Recombinant Limited

Re: R3000 Disassembler
Written by Simon Everett on Sun Jan 19 01:29:19 1997

To William B. Norris IV:

SCEE released a disassembler which disassembles object files. It is called omi_psx.exe, developed by Visual Sciences Ltd in the UK. Contact SCEE Developer Support or email brianm@vissci.demon.co.uk

Re: R3000 Disassembler
Written by Mike Fulton on Fri Jan 17 22:54:50 1997

To William B. Norris IV:

I'm afraid we don't have a disassembler other than the one built into the debugger.

Mike

SNLink format object file

Written by William B. Norris IV on Fri Jan 17 13:34:05 1997

Is there any documentation on the SNLink format object file?

Our game's .exe file is currently 866,304 bytes in size and I am working on getting the code and assets to fit into the 2MB of PlayStation main memory.

I'd like to write two programs based on the object file format:

- 1) A program which will analyze all object files to generate a cross-reference listing, showing which functions call which other functions. This would help me determine "dead code".
- 2) The second, more ambitious program is a linker which can link on a function-by-function basis. The Psygnosis linker will bring in the whole object file even if only a single function is called within that object file. By looking at our map file, I have determined that *at least* 100K of the executable is not called and can therefore be safely eliminated from the program. As I'm unwilling to place each function in a separate file and #ifdef'ing code would be a maintenance nightmare, I'm seeking an automatic solution by coding this tool.

William B. Norris IV
Recombinant Limited

Re: SNLink format object file
Written by Jeremy Gordon on Fri Feb 28 17:28:28 1997

To William B. Norris IV:

Hi William, check out dumpers.zip which contains the same thing as the files dumpsym.zip, dumpobj.zip, and dumpcpe.zip. All these files are located on the SN Systems 'BBS'. Thier object file format is like a simplified version of COFF, and it's pretty easy to figure out from thier code. I recently wrote my own specialized version of thier linker, so if you run into problems figuring out thier patch format (which isn't really 'documented' other than source code comments) feel free to send me some email. The guys at SN are also pretty cool about answering questions, duncan@snsys.com is the first line Good luck!

- jeremy
jgordon@magicarts.com

Re: SNLink format object file
Written by Bill Long on Fri Feb 28 00:53:15 1997

To William B. Norris IV:

William

Have you had any success in finding the compiler's object format? I am also looking to rewrite the linker in order to improve cache performance.

Bill Long

Re: SNLink format object file
Written by Mike Fulton on Fri Jan 17 21:58:33 1997

To William B. Norris IV:

You'd have to contact SN Systems to get information on their file formats.

As far as your second program idea goes, a linker that links function by function instead of by object module is sort of the holy grail of development tools. I've seen dozens of linkers for different systems and different compilers, and I've never seen one that was capable of linking in anything smaller than a complete object module.

The problem is that no object file format has the information necessary to reliably suck out individual functions or individual variables. The best you can hope for is a reasonable guess, and in most cases it doesn't take too much to make it guess wrong.

My suggestion is that you look into other methods of eliminating unused code. You shouldn't have to go to the extreme of breaking each and every function into a separate source module, but taking a few steps in this direction may help significantly.

As far as your first program idea goes, there are existing C/C++ source code cross reference utilities available, both commercial shrink-wrap packages and downloadable shareware/freeware/PD (some even with source code) that can do what you want.

Mike @ SCEA

RAM Load timings discrepancies

Written by Christer Ericson on Fri Jan 17 02:11:20 1997

RAM Load timings discrepancies

Working on optimizing some low-level code, I've been trying to figure out some more exact cycle timings for reading/writing main RAM.

According to information posted earlier by Sony representatives (Mr Okamoto, and, hmm, someone else) to the SN Systems BBS, memory reads on the Playstation are 5 cycles: 4 cycles to read from memory to the R-buffer (a single four-byte register) and 1 cycle to read from the R-buffer to the CPU.

This seems rather simplistic however, and in fact, the Fall '96 Playstation Developer Seminar Notebook says the following (page 10):

“First read from a particular page in RAM takes 5 ticks. 4 ticks for RAM to R3000 R-buffer, 1 tick R-buffer to CPU. Second and subsequent reads take just 2 ticks. 1 tick for RAM to R3000 R-buffer, 1 tick R-buffer to CPU. True for both data reads and I-cache burst reads.”

Unfortunately, this doesn't make sense either. The last sentence seems to imply that data reads (i.e. data loaded from RAM by a load-instruction) are affected by page faults, whereas this doesn't seem to be the case at all.

Running code_a and code_b below, with a0 pointing at one page and a1 pointing at another page, code_a should cause 6/6 load page faults, whereas code_b should cause only 2/6 load page faults, and consequently code_b should execute faster than code_a. This is not the case, they execute in the same amount of time.

Conclusion 1: data reads are not (visibly) affected by page faults.

```

code_a li    t0,20000      code_b li    t0,20000
@loop lw     t1,0(a0)     @loop lw     t1,0(a0)
      lw     t2,0(a1)     lw     t2,4(a0)
      lw     t3,4(a0)     lw     t3,8(a0)
      lw     t4,4(a1)     lw     t4,0(a1)
      lw     t5,8(a0)     lw     t5,4(a1)
      lw     t6,8(a1)     lw     t6,8(a1)
      subu   t0,t0,1      subu   t0,t0,1
      bnez   t0,@loop     bnez   t0,@loop
      nop
      jr    ra            jr    ra
      nop                nop

```

Furthermore, replacing the 6 loads with a small delay that takes 36 cycles results in a piece of code that again executes in the same amount of time.

Conclusion 2: a load (seemingly) takes 36/6=6 cycles, not 5 and/or 2 cycles.

However, this is not true either! Digging deeper, loads are not consistently 6 (nor 5 and 2) cycles. It seems the above results are caused by pipelining effects, so the next thing to look at is how exactly loads are pipelined.

Running code_c and code_d below results in code_c being faster than code_d. This result implies that the two loads are not pipelined, but that the loads are fully (or close to fully) executed before the next instruction is scheduled, otherwise the added nop in code_d would have been hidden by the 5-something cycle load delay slot.

Conclusion 3a: loads are (seemingly) executed in full before the next instruction is scheduled.

```

code_c li    t0,20000      code_d li    t0,20000
@loop lw     t1,0(a0)     @loop lw     t1,0(a0)
      lw     t2,4(a0)     lw     t2,4(a0)
      nop
      addu   t1,t1,t2      nop
      subu   t0,t0,1      addu   t1,t1,t2
      bnez   t0,@loop     subu   t0,t0,1
      nop                bnez   t0,@loop

```

```

jr   ra           nop
nop                jr   ra
                    nop

```

Strangely enough, it is possible to arrive at exactly the opposite conclusion! Look at code_e and code_f below. They execute in the same amount of time, yet code_f has 5 extra nops in the loop, which, if loads very fully executed before the next scheduled instruction, would cause the execution time to be much higher than for that of code_e.

Conclusion 3b: loads are pipelined.

```

code_e li    t0,20000    code_f li    t0,20000
@loop lw     t1,0(a0)    @loop lw     t1,0(a0)
      lw     t2,4(a0)    lw     t2,4(a0)
      lw     t3,8(a0)    lw     t3,8(a0)
      addu   t1,t1,t2     nop
      addu   t1,t1,t3     nop
      subu   t0,t0,1     nop
      bnez   t0,@loop    nop
      nop                nop
      jr     ra          addu   t1,t1,t2
      nop                addu   t1,t1,t3
                        subu   t0,t0,1
                        bnez   t0,@loop

      nop
      jr     ra
      nop

```

As conclusion 3a and 3b cannot co-exist, we must conclude:

Conclusion 4: Loads are pipelined, but the delay slot of a load instruction varies between 1 and 6 cycles.

Could someone from Sony (or anyone else who have any idea) please explain the discrepancies between the Sony documentation and the actual behaviour of the DTL-2000?

Are loads affected by page faults or not?

If so, why is it possible to construct examples which seemingly contradict the fact, i.e. under what circumstances are they affected by page faults?

How are loads pipelined? How can the results of my example programs code_c through code_f be explained in light of this?

Christer Ericson
 Programmer, Neversoft Entertainment

Re: RAM Load timings discrepancies
 Written by Jerry Jessop on Wed Jan 22 22:36:34 1997

To Christer Ericson:

The reply from Simon Everett was right on the money, but maybe I can add some details. The text below is directly from the PSX CPU manual compiled by LSI Logic and Sony.

(Please keep in mind this was translated from Japanese)

“ The CPU reads data through a 1 word READ_BUFFER, and read accesses have higher priority than write accesses, since the CPU stalls when it can't read data or load instructions. Data to be written is stored in 4 words WRITE_BUFFER without penalty if it's not full. BIU (Bus Interface Unit of CPU) writes

data in WRITE_BUFFER when main bus is idle. If the CPU attempts to write data when WRITE_BUFFER

is full, it stalls till the end of the one word write transaction. When the CPU attempts to read data which is

at the same address stored in WRITE_BUFFER, write access has higher transaction. In general, it is difficult to anticipate when an actual write transaction will occur. It is recommended to execute a dummy

read command to the same address of the previous write command, then the write command has to be done immediately”

A somewhat confusing translation, but I think you get the picture in regards to how unpredictable the actual write transaction is. Below is an example of forcing the write transaction from the LSI CPU manual.

Example: Unmask V-BLANK interruption and go ahead.

```
li s0,0x1f801074 #Interrupt Mask register Address
lw t0,(s0) #read Mask register
li t1,1 #bit1=Vblank
or t1,t1,t0 #Vblank Unmask
sw t1,(s0) #write Mask Register
li t0,(s0) #Mask Register dummy read
```

For the record the published (LSI CPU Manual) access timing are as follows.

Main Memory

Single READ	4 Cycles
Single WRITE	4 Cycles
Burst READ (4 words)	4 Cycles for 1, 1 for rest
Burst WRITE (in same page)	4 Cycles for 1, 2 for rest
DMA	4 Cycles
Internal Registers	2 Cycles
GPU I/O access	2 Cycles
GPU DMA	1 Cycle

Jerry Jessop
Hardware Engineer
Sony Computer Entertainment of America

Re: RAM Load timings discrepancies
Written by Christer Ericson on Thu Jan 23 00:26:18 1997

Jerry, thanks for the reply.

That's a good answer, unfortunately to a different question. I thought I had made it pretty clear that I was asking about reading (and not writing) to main RAM, but since two people have misread it to be about writes, I seems I somehow wasn't clear enough, for which I apologize.

Anyway, my question(s) still stand as asked.

Let me just again point out that the code I provided did not have any writes in it whatsoever which leaves only the loading of words from memory to account for the timing discrepancies.

I urge anyone interested to run and time the pieces of code I provided to verify the results.

Christer Ericson,
Neversoft Entertainment

Re: RAM Load timings discrepancies
Written by Christer Ericson on Tue Jan 21 22:07:28 1997

Simon, I'm afraid you must have misread what I wrote. You are talking about writing (storing) to memory whereas I am talking about reading (loading) from memory.

I do not dispute that page faults affect (main) RAM writes, and that the flush of the W-buffer makes cycle counting somewhat difficult when it comes to write instructions.

What I am saying, and what the example code I gave proves, is that the information given to us concerning reading from main RAM doesn't make sense. Reading from main RAM is not a simple 5 cycles/read as we have been told.

As you can easily verify, none of the code examples I gave contain any store instructions and thus the timing discrepancies that I show cannot simply be explained by referring to the sporadic behaviour of writes and the W-buffer.

My code implies that there is potentially a lot to be gained by scheduling loads from RAM in a particular way: judging from code samples e and f, up to 20%!

Of course, the answer could simply be that the DTL-2000 simply does not behave as a real Playstation, what do I know? But if that's not the case, then I'm sure every developer out there would be interested in knowing if and how they can get an extra 20% out of their inner loops by, say, pairing loads 3 and 3. I know I am. Aren't you?!

That's why I'd like Sony to tell us what's up with the timings of load instructions. So, guys, what's up with the timings of load instructions, eh?! :)

Christer Ericson,
Neversoft Entertainment

Re: Re: RAM Load timings discrepancies
Written by Simon Everett on Mon Jan 27 19:59:51 1997

To Christer Ericson:

Do you understand how the r3000a instruction pipeline works? Your code examples are demonstrating three characteristics of the CPU: RAM latency, the r3000a pipeline, and load instruction delay slots.

A & B : Loads are not affected by page faults.

C & D : The r3000a has a 5 stage pipeline. The result of a load is available before the instruction is complete (code c), but if the result is not used at this point, the pipeline is stalled until the load instruction is completed (code d).

E : t2 is available in the pipeline (the instruction is not complete), t1 is not available (the instruction is not complete, availability has been missed). Thus the pipeline is stalled until both t1 & t2 loads are complete.

F : t1 and t2 loads are both complete, CPU is not stalled as t1 & t2 are not required before instructions are complete (NOPs in your code).

Conclusion.

Loads can be pipelined (as all instructions are), but you must be careful exactly when you use the result of a load. If you do not use the result of a load when it is available in the instruction pipeline, you have to wait until the instruction is completed.

Be careful when talking about RAM latency timings, and CPU pipeline stages. Be aware that an instruction may be stalled in the pipeline (effectively one stage for an instruction is stalled) if exclusive resources are required, giving an appearance of variable cycle timings for instructions.

If instructions are scheduled correctly, the number of cycles per instruction can potentially be reduced by the depth of the pipeline. So it is possible to obtain results from series of loads at 1 per cycle (if enough loads are performed).

Simon Everett
FHT Development Corporation

Re: Re: RAM Load timings discrepancies
Written by Christer Ericson on Tue Jan 28 20:51:55 1997

Simon, the only reference on the R3000a pipeline that I know of is the Kane/Heinrich "MIPS RISC Architecture" book, which barely scratches the surface on how the pipeline works. It certainly does not mention for what duration of the pipeline the result of a load is available and when it is not.

Anyway, you say: "If you do not use the result of a load when it is available in the instruction pipeline, you have to wait until the instruction is completed [before you can use the result]." and comment on the behaviour of my examples in light of this.

However, for code C:

```

        lw    t1,0(a0)
        lw    t2,4(a0)
nop
addu   t1,t1,t2

```

you say the result of t1 and t2 are available when doing the add (and not one cycle later as in code D) and that there is no stall in this case. Then for code E:

```

        lw    t1,0(a0)
        lw    t2,4(a0)
        lw    t3,8(a0)
        addu  t1,t1,t2
        addu  t1,t1,t3

```

you go on to say that t1 is not available at the time of the add, and that there is a stall—even though the instruction is in the same position in the pipeline in both code examples (immediately following the load delay slot instruction).

With the pipeline looking like this for the two examples:

code C:

```

lw    t1,0(a0): |IF|RD|ALU|MEM|WB|
lw    t2,4(a0): |IF|RD|ALU|MEM|WB|
nop                                     |IF|RD|ALU|MEM|WB|
addu  t1,t1,t2: |IF|RD|ALU|MEM|WB|

```

code E:

```

lw    t1,0(a0): |IF|RD|ALU|MEM|WB|
lw    t2,4(a0): |IF|RD|ALU|MEM|WB|
lw    t3,8(a0): |IF|RD|ALU|MEM|WB|
addu  t1,t1,t2: |IF|RD|ALU|MEM|WB|
addu  t1,t1,t3: |IF|RD|ALU|MEM|WB|

```

would you care to explain when exactly the result of a load is available in the pipeline, and why the “addu t1,t1,t2” would stall in code E and not code C?

In fact, with code G and H (as given below) executing in the same amount of time, it seems likely that both code C and E would be stalling on the add, no? Which then, of course, again begs the question: why is code C faster than code D?

code G:	code H:
lw t1,0(a0)	lw t1,0(a0)
lw t2,4(a0)	lw t2,4(a0)
lw t3,8(a0)	nop
addu t1,t1,t2	addu t4,t4,t2
addu t1,t1,t3	lw t3,8(a0)
nop	nop
nop	addu t4,t4,t3

Christer Ericson,
Neversoft Entertainment

Re: Re: Re: RAM Load timings discrepancies
Written by Simon Everett on Tue Jan 28 22:58:52 1997

To Christer Ericson:

In your diagrams you've assumed each stage in the pipeline takes the same amount of time to complete for all instructions, and you haven't allowed for pipeline stalls.

Load instructions do not take the same amount of time to complete. From the data given by Jerry Jessop, a single load from RAM is not as efficient as a burst load. A single load will take longer to be executed than the second load in a series of pipelined loads (burst load), as the data is delivered faster in burst mode. The case is even more dramatic with a MIPS CPU which has a proper D-Cache, operating on a bus which has a burst size of 128 or 256 bytes.

The difference between your code c & e examples is the NOP in the load delay slot. Your timings for the pipeline stages of a NOP assume that the timings are the same for the pipeline stages of a load. Clearly a single load (not burst) will take longer than a NOP, as data must be fetched from RAM. Subsequent loads will also take longer than a NOP as the data must still be delivered from the R-buffer.

Kane/Heinrich states 'the time required for register-to-register operations is much less than the time needed to operate on memory operands'.

Simon Everett
FHT Development Corp.

Re: Re: Re: Re: RAM Load timings discrepancies
Written by Mike Biddlecombe on Thu Jan 30 20:46:23 1997

To Simon Everett (and others):

I have been following this thread with some interest. My problem is that I need run through a large number of 16bit values transforming them with a routine that runs exclusively on registers (and scratchpad)

Sort of like:

```
for (0...100) {  
  read 8 values from main memory  
  process 8 values  
}
```

Which can be written in asm as:

test1:

```
la a0,table_o_data  
addi v0,r0,100
```

```
bloop: lh t0, 0(a0) lh t1, 2(a0) lh t2, 4(a0) lh t3, 6(a0) lh t4, 8(a0) lh t5,10(a0) lh t6,12(a0) lh  
t7,14(a0)
```

```

;;; simulate processing
nop
;;; end of processing

```

```

addi a0,a0,16
bne v0,r0,bloop
addi v0,v0,-1

```

```

jr ra
nop

```

In test1 the bottleneck is the read from main ram as each read takes a 4 cycle stall (because the PSX has no data cache burst read). This got me thinking, why not try and use the I-cache line fill? I can replace every 16bit value in the table with a 32bit ADDI instruction.

Sort of like:

load 8 immediate values

process 8 values

load 8 immediate values

process 8 values

:

:

< repeated 100 times >

Which can be written in asm:

```
test2
```

```
la v0,ProcessRoutine
```

```
addi t1,r0, 8 ;; constants would vary addi t2,r0,12 ;; to reflect table data addi t3,r0,16 addi t4,r0, 4
addi t5,r0,
```

```
4 addi t6,r0, 8 addi t7,r0,12 jalr v1,v0 addi t0,r0, 4
```

```
addi t1,r0, 8
```

```
addi t2,r0,12
```

```
addi t3,r0,16
addi t4,r0, 4
addi t5,r0, 4
addi t6,r0, 8
addi t7,r0,12
jalr v1,v0
addi t0,r0, 4
< repeated 100 times >
```

```
jr ra
```

```
ProcessRoutine:
```

```
::: simulate processing
```

```
nop
```

```
jr v1
```

```
nop
```

```
::: end of processing
```

```
*****
```

In test2 I am taking I-cache hits as I move through the 'data', but the process routine stays in the cache. Rather than taking the 4-cycle stall every time I read in a data value (and Mr Okamoto is quite correct, the penalty is paid regardless of whether you read bytes, halves or words) I get a boost from the I-cache line fill. In test2 only something like 1 in 4 'reads' incur a stall (the JALR must be accounted for too).

I timed both of these methods and found that test2 took only %60 of the time need to complete test1.

So if you have the space to double the size of a large 16bit table by converting it to immediate loads, why

not give this a try

Mike

Speedy processing of 16-bit table (was: re: RAM Load timings discrepancies)

Written by Christer Ericson on Fri Jan 31 05:21:21 1997

Mike, that's definitely an interesting technique! The drawback, as you noted, is the increase in size, which is okay for a small to moderate size table.

There are a few things you could do to your basic code though to make it run faster. Assume you start with code looking something like this, which is pretty close to what you posted:

```
code1
li    t0,10000
@loop
        lh    t1,0(a0)
        lh    t2,2(a0)
        lh    t3,4(a0)
        lh    t4,6(a0)
        lh    t5,8(a0)
        lh    t6,10(a0)
        lh    t7,12(a0)
        lh    t8,14(a0)

rept 8
nop          ; does some processing
endr

rept 8
nop          ; does more processing (eg. bumping a0)
endr

        subu  t0,t0,1
        bnez  t0,@loop
nop
jr    ra
nop
```

As you point out, the reading from RAM kills this routine. But instead of reading halfwords, lets read full words and then reconstruct them into halfwords. I'm here assuming that you actually need them to be sign-extended. If not, so much the better of course.

Okay, let's read 4 words and split them into 8 halfwords, making sure they're sign-extended. The trick is (and this is purely empirical on my part, which is why I started this thread in the first place) to make sure you do the load as far ahead as possible from where you are using it. Actually, make that 6 instructions or more. This results in this code:

```
code2
li    t0,10000
@loop
        lw    t1,0(a0)
        lw    t3,4(a0)
        lw    t5,8(a0)

        move  t2,t1
        sll  t1,t1,16
```

```

sra  t1,t1,16
sra  t2,t2,16

move  t4,t3
sll  t3,t3,16

lw   t7,12(a0)

sra  t3,t3,16
sra  t4,t4,16

move  t6,t5
sll  t5,t5,16
sra  t5,t5,16
sra  t6,t6,16

move  t8,t7
sll  t7,t7,16
sra  t7,t7,16
sra  t8,t8,16

rept 8
nop      ; does some processing
endr

rept 8
nop      ; does more processing (eg. bumping a0)
endr

subu  t0,t0,1
bnez  t0,@loop

nop
jr   ra
nop

```

That's not too bad, but it's still not as fast as we can go. The problem is the three loads at the start of the code. By unrolling two of them, moving them out of the loop at the start, and then inserting them (still separated by 6 or more instructions before they are used) at the end of the code we'll go even faster. This results in the final code looking like this:

```

code3
li   t0,10000

lw   t1,0(a0) ; moved out of the loop
lw   t3,4(a0) ; same here

@loop

lw   t5,8(a0)

move  t2,t1
sll  t1,t1,16

sra  t1,t1,16

```

```

sra    t2,t2,16

move   t4,t3
sll    t3,t3,16

lw     t7,12(a0)

sra    t3,t3,16
sra    t4,t4,16

move   t6,t5
sll    t5,t5,16
sra    t5,t5,16
sra    t6,t6,16

lw     t1,0(a0) ; folded back into the loop (at "end")

move   t8,t7
sll    t7,t7,16
sra    t7,t7,16
sra    t8,t8,16

rept 8
nop    ; does some processing
endr

lw     t3,4(a0) ; also folded back into the loop (at "end")
rept 8
nop    ; does more processing (eg. bumping a0)
endr

subu   t0,t0,1
bnez   t0,@loop

nop

jr     ra

nop

```

This will not get down to the 60% you report for your technique, but my very rough timings indicate that code3 executes at roughly 70% of the time for code1, which has to be considered pretty good. (Routines timed exactly as they're given here.)

BTW, glad to hear someone finds the "load" thread interesting! :-)

Christer Ericson,
Neversoft Entertainment

Re: Speedy processing of 16-bit table (was: re: RAM Load
timings discrepancies)
Written by Mike Biddlecombe on Sun Feb 02 10:18:49 1997

To Christer Ericson: Timing results for consecutive loads

I am new to the PSX (and MIPS asm) so I hope you will indulge me as I try to figure out just what is going on WRT to optimizing memory reads.

See if the following code frags make sense to you:

```
delay macro
```

```
rept \1  
nop  
endr  
endm
```

```
code1 ; Running time: 369 units  
; Fastest time of all  
; No extra opcodes added/rescheduled
```

```
li t0,500
```

```
@loop
```

```
lw t1,0(a0)  
lw t2,4(a0)  
lw t3,8(a0)
```

```
nop  
add t4,t3,t3
```

```
subu t0,t0,1  
bnez t0,@loop  
nop
```

```
jr ra
```

```
nop
```

```
code2 ; Running time: 441 units  
; One extra instruction follows each load
```

```
li t0,500
```

```
@loop
```

```
lw t1,0(a0)
```

```
delay 1
```

```
lw t2,4(a0)
```

```
delay 1
```

```
lw t3,8(a0)
```

```
delay 1
```

```
nop  
add t4,t3,t3
```

```
subu t0,t0,1  
bnez t0,@loop
```

nop

jr ra

nop

code3 ; Running time: 441

units

 ; Five extra instructions between loads

 ; Exactly the same running time as code2

li t0,500

@loop

lw t1,0(a0)

delay 5

lw t2,4(a0)

delay 5

lw t3,8(a0)

delay 5

nop

add t4,t3,t3

subu t0,t0,1

bnez t0,@loop

nop

jr ra

nop

code4 ; Running time: 593 units

 ; Same number of instructions as code3

 ; ...but extra code is not interleaved

 ; Longest running time of all

li t0,500

@loop

lw t1,0(a0)

lw t2,4(a0)

lw t3,8(a0)

nop

add t4,t3,t3

delay 5

delay 5

delay 5

subu t0,t0,1

bnez t0,@loop

nop

jr ra

nop

Comments

As far as I can see, consecutive loads (ie, each new load immediately follows the previous load) will execute faster than loads which are interleaved with other (non-conflicting) op-codes . However each of the these additional loads cannot take advantage of the previous delay slot and the processor stalls.

This would indicate that it would make sense to insert other op-codes, thus trying to get the load down to a magic 1/cycle per instruction.

But (and a big but it is) when we insert any instruction at all between two consecutive loads (or between the load of a target reg and an instruction that uses that reg as a source) we immediately incur a 5-cycle penalty when we go to start the next load (or use the target). Exactly why this occurs I cannot completely explain (anyone?... anyone?... Bueller?).

My thinking is that the circuitry that allows for the published statement, “the second and subsequent reads take just 2 ticks”, gets reset when the interleaved instruction is executed. In this case, the subsequent read is now a ‘first’ read, and will take 5 cycles to complete rather than 2.

But we still get a win if we can stuff at least 5 instructions between the loads (or 6 between the load op and the op that uses the target register), because then we can take advantage of the load delay slots for instructions that we would have had to execute sooner or later. This is what is what makes code3 faster than code4.

I guess my whole point (and why I am posting to this conference) is:

Does the current version of GCC on the DevCD take these delays into account? Without the source it is hard to tell.

If not, is anyone interested in getting this optimization working?

I have got a cross-compiler based on GCC 2.7.2.1 (built from DJGPP) working as a replacement to CC1PSX.EXE.

Is it worth the effort to modify the MIPS.C module to allow for funky load reordering. Or has this already been done, and I just missed it?

.biddle

Re: Speedy processing of 16-bit table (was: re: RAM Load timings discrepancies)
Written by Christer Ericson on Mon Feb 03 02:18:06 1997

Mike, your analysis pretty much agrees with mine. The one thing I’m sceptical about though is this whole “burst read, 4 cycles for first, 1 cycle for rest”. I suspect it is baloney, that there is no burst read, or if there is, consecutive reads are being stalled at the one word R-buffer (or at least being stalled, for whatever reason) so it amounts to nothing anyway. I don’t know how else to interpret that 3 consecutive loads (with no access to the registers read, and enough nops to allow them to be fully processed before looping again) roughly amounts to the same time as that for 15 nops, and 4 consecutive loads (still no accessing the registers and enough nops) roughly amounts to 20 nops.

Thus it seems we have that loads are always 5 cycles, but as they are pipelined we can insert 5 cycles worth of instructions between loads to avoid stalling on consecutive loads, or 6 cycles worth of

instructions between a load and first use of the register loaded (5 cycles + 1 cycle load delay slot) to avoid stalling on accessing the register.

However, there are still questionmarks remaining. For one, if the above is true, then why is it that a single load (followed by nops) roughly equates to 3 nops?! I've been trying to draw pipeline diagrams, inserting stalls of different lengths in all possible places, but I never come up with something that accurately covers the (seemingly) contradictory instruction timings that I've measured; this is what I had hoped someone here would be able to help me with. Anyway, for all purposes, and until something better comes along, I'm sticking to the previous paragraph as a model for how loads work.

To answer your other question: no, AFAIK, GCC/G++ does not take this into account, or if it does, it does a terribly bad job out of it because I've never seen it do it. Actually, overall it produces rather lousy MIPS code, at least compared to the standard of its 680x0 code; it rather consistently fails to do even the simplest things. (I don't know, but I get the impression that in my case, it could be objects that really confuses the optimizer.)

Getting the compiler to schedule loads properly would definitely be worth it, but I have no idea how much effort it would take. It would also be good to have it try minimizing write page faults (I'm not sure if it does this, but I doubt it), but again I have no idea how to do this as I've never played with the GCC source code. Anyone?

Christer Ericson,
Neversoft Entertainment

Re: RAM Load timings discrepancies
Written by Christer Ericson on Wed Jan 29 03:20:00 1997

Simon,

as it's not clear to me for how long, where and when the pipeline stalls are, I obviously couldn't put them in the diagrams, instead the diagrams were intended as a starting point for discussing where the stalls are actually taking place.

Perhaps you can provide the full pipeline diagrams for code C and code E? Even though you imply you understand what is going on, I seriously don't think you can. Because, even though what you say is theoretically sound, it simply doesn't agree with empirical data; this discrepancy being the source of my confusion.

It seems I've tried every possible interpretation, but each time I can come up with a simple piece of code that wreaks havoc on the theoretical model. Have you actually tried timing some code yourself?

Looking for Metrowerks comments
Written by Stephen Magladry on Thu Jan 16 00:26:58 1997

I am looking for comments on the Metrowerks Development Environment. If you can share your thoughts on the following, it would be much appreciated.

How does generated code compare to GNU code? How does the MW editor compare to other compiler editors? How does the MW debugger compare to the Sony debugger? Are compiles noticeable faster? We are very early in our PlayStation Development. Is DR2 stable enough? Do you know how soon release 1 is expected? Other general comments not covered from the above list?

thanks,

Please reply to:

Stephen magladry
smagladry@bigrain.com

Re: Looking for Metrowerks comments

Written by simon on Fri Feb 14 20:45:51 1997

To Stephen Magladry:

I got the CW DR2 running. But I noticed the compiling is a lot slower than I use ccpsx. Another thing is I don't know whether they support cd emulation from hard drive or CD. You also have to rename the final executable to .exe file when you burn a CD. As the the IDE system itself, it's very handy and decent.

simon
readysoft Inc.

CD Emulation from RUN.EXE

Written by Les Hedger on Thu Jan 09 20:04:14 1997

I am having trouble getting my CD emulator to work from a test program that uses cdrSearchFile() to find a file. The program I'm using is tuto7.cpe from the developer's CD tutorials. I have succeeded in running the emulator example program RCUBE by creating an image and then running CDEXEC.CPE. This works fine. However, when I try to run the tutorial program using "RUN TUTO7" I get a message that the disk has a format error.

Here are the only changes I have made to the sources: I have changed TUTO7 to set CD debug level to 3, in order to get the error and status messages, and I have changed the name of the file to "\\RCUBE.EXE;1" in order for it to work with the working RCUBE partition.

Here's the output from the execution:

```
ResetGraph:jtb=8002fde4,env=8002fe2c
TYPE : 6 free button or flying-V form
PS-X Control PAD Driver Ver 3.0
CD_init:addr=800302ac
CdINop...
CdIReset...
CdIDemute...
CdISetloc...
CdISetmode...
CdIReadN...
CdIPause...
CD_newmedia: Disc format error in cd_read(PVD) \RCUBE.EXE;1: not found CdISetloc...
CdIReadN...
CdIPause... CD_newmedia: Disc format error in cd_read(PVD) \RCUBE.EXE;1: not found
```

... and it just retries repeatedly.

Am I going about this in the right way? Is it possible to access the CD emulator from a program started with RUN, or do I have to create (and run from) a bootable image in order to be able to access the CD emulator?

Re: CD Emulation from RUN.EXE
Written by Mike Fulton on Mon Apr 14 20:50:48 1997

To Les Hedger:

The endless retry is happening because the sample program apparently expects you to hold down the "Select" button on your controller in order to make it stop looping.

However, it only checks this once on each pass of the loop, and it uses the old (no longer recommended) PadRead function to read the controller. This means you'd have to hold down the button until it looped around.

There are apparently some problems using CdSearchFile with the emulator. Unfortunately, there have been so many contradictory reports that I haven't quite gotten a handle on the situation yet.

I've attached an object file named "iso9660.obj" to this message. Please download it and link your sample program with it. This should give an additional bit of debug information that may help us figure out the problem.

Please let me know what debug information is printed when you use this object module.

Mike

Name: Iso9660.obj
Content-type: application/octet-stream
Size: 8742 bytes

Re: CD Emulation from RUN.EXE
Written by Dedan Anderson on Fri Apr 11 18:54:13 1997

To Les Hedger:

I had problems with my CdSearchFile also. I also hardcoded sector values in to force a Cdread of anything, and discovered that for some reason it was returning an error. I then in fit of confusion changed my current libraries (3.5) to an older library one of my co-workers had and viola CdSearchFile works!!! You should be able to use RUN or selemu.cpe, cdexec.cpe to get your to read off a CD emu. Well, I hope that works.

Re: CD Emulation from RUN.EXE
Written by frank on Fri Jan 10 01:25:34 1997

To Les Hedger: Some more info would be useful in helping to figure out your problem. You put RUN TUTO7 in quotes, is that what you're entering on the command line? What file did you rename as RCUBE.EXE? I'm not sure why you renamed it in the first place, since you need to change the .cti file to reflect the file/directory structure expected by tuto7 anyhow before building a new image.

You can access the emulator without using a bootable image, but you must run SELEMU.CPE first. For example, if you create a simple image with just one file, \DATA\MOV.STR;1, you can type "run tuto7" and it will work.

I'm assuming the "tuto7" you refer to is the one in the \psx\sample\cd\tuto directory. There are 4 or 5 tuto7's in the samples, but I don't think any others use any CD accessing.

If this doesn't help, a more detailed step by step description of what you're doing could point us the right direction.

Frank

Re: Re: CD Emulation from RUN.EXE
Written by Les Hedger on Fri Jan 10 03:58:39 1997

To frank:

Thank you for your prompt reply.

You are correct in your assumption that my starting source file is TUTO7.C from the \psx\sample\cd\tuto folder.

I have not changed CD.CTI for rcube, in order to be sure that the CD image is correct. The `_only_` change that I made was to TUTO7.C, where I changed the debug level to 3, and I changed the string which contains the name of the file to open. You can see from the output that tuto7 is trying to open "\RCUBE.EXE;1" when it fails. I figured that it didn't matter which file I tried to open because the program never gets past the CdSearchFile() call.

When booting DOS, I run these commands:

```
DEXBIOS /a 1340 /i 12  
MESS1  
CDBIOS /a 390 /i 10 /d 7  
CDMON 4
```

and then I run:

RESETPS 1

```
RUN /w4 PATCH.CPE  
RUN /w4 SELEMU.CPE  
RUN /w4 TUT07.CPE (or RUN CDEXEC.CPE to run RCUBE) TESTMESS
```

...and I see the output that I included in my first message. Since I can run CDDISK.EXE without problems, I think that the BIOS stuff is OK, as well as the SELEMU program, or else the RCUBE program would not run. I can see from the output that tuto7 is indeed talking with the CD, because

the output shows that it does a seek without complaining, and the CDMON program indicates disk activity on my screen.

Please let me know if you have any other questions or any ideas as to what the problem might be.

Perhaps you could create some sort of test suite which includes a proper CD.CTI file, a simple data file, and a simple program to use it? It would be most helpful if the same test program could be compiled to be executed from a bootable image, as well as using the RUN command. I don't see a lot of info on making a simple bootable image, but I may just need to dig a little deeper. Since RCUBE.EXE and PSX.EXE are pre-compiled, they are not a complete example.

I'd appreciate any help you could give. Thanks. Les Hedger doc@mail.pfmagic.com

Re: Re: Re: CD Emulation from RUN.EXE
Written by frank on Fri Jan 10 22:45:34 1997

To Les Hedger:

I obviously misunderstood your original post, and your clarification helped a bunch.

In order to do what you'd like, you DO need to edit the cti file. These are the changes/additions:

1. For the type field of the "disc" command, change CDROM to CDROMXA.
2. Following the "LeadIn" command, change MODE1 to XA.
3. Ditto for the "Track" and "LeadOut" commands.
4. For each of the files in the image, between the "File" command and the "Source" command, insert the following—"XAFileAttributes Form1 Data".

The reason for this is that the Playstation doesn't like mode1 sectors. I'm not sure why the PsyQ supplied cti file makes a working image, perhaps cdexec.cpe somehow corrects the situation.

Frank Sheperson

C++ Compiler bug slowing things down

Written by Mick West on Tue Jan 07 20:57:47 1997

I have problem with the GNU C++ compiler generating spurious code with a fixed point class that uses operator overloading. I've made the following fragment of code that can be compiled under C++ to illustrate the problem:

```
////////////////////////////////////  
  
// the class, just like my fixed point class, one member // variable, an int.  
  
    class hmm {  
    public:  
  
int value;  
  
        };
```

```

// only one operator defined, for simplicity
inline hmm operator + (register hmm a, register hmm b)
    {

hmm temp;
temp.value=a.value+b.value;
return temp;
    }

// globals variable, just to simplify the example
hmm mick_result;

// silly function, just illustrate the problem
void mick( hmm a, hmm b)
    {

mick_result = a+b;
    }

////////////////////////////////////

```

this is the code it produces for mick(hmm a, hmm b)

```

mick__FG3hmmT0
80278E54>27BDDFFF8 addiu sp,sp,-$8

80278E58 00852021 addu a0,a0,a1
80278E5C AFA40000 sw a0,0(sp) // totally not needed
80278E60 AF8405AC sw a0,$5ac(gp)
80278E64 27BD0008 addiu sp,sp,$8
80278E68 03E00008 jr ra
80278E6C 00000000 nop

when it should just be
mick__FG3hmmT0
80278E58 00852021 addu a0,a0,a1
80278E60 AF8405AC sw a0,$5ac(gp)
80278E68 03E00008 jr ra
80278E6C 00000000 nop

```

I think the problem is not with the “temp” variable, but rather with the “return”, which seems to want somewhere to put the value, so it uses the local heap, but as it never reads the value again, it’s all wasted memory accesses, which are very very slow.

Now, as we are using a fixed point class with operator overloading to replace all the uses of “float” in the MDK code (which we are converting to the PSX from the PC version, being done by Shiny), we are getting a LOT of spurious code which slows things down a lot, and also bloats to code size.

Do you guys have any idea of a way round this that still allows me to use operator overloading (as this greatly reduces the amount of code I need to convert)?

Thanks for any help here,

Mick.

Re: C++ Compiler bug slowing things down
Written by Mike Fulton on Wed Jan 08 18:56:44 1997

To Mick West:

First of all, please understand that C++ is not really “officially” supported. We’ll try to help as best we can, but keep that in mind.

Now that I’ve given the basic disclaimer, let’s see if we can figure out what’s going on.

A lot of the code you cut out when you did your “this is what it should be” listing was the function prologue and epilogue for the “mick” function, not code related to the operator overloading. I’d recommend something I usually do when I’m trying to figure out what code is being generated by the compiler for something: include a call to a dummy function immediately before and after the code I’m trying to figure out. In other words, change your test function to:

```
void mick( hmm a, hmm b)
{
    dummy();
    mick_result = a+b;
    dummy();
}
```

Of course, you should have the dummy() function somewhere, although it doesn’t need to do anything.

This allows you to better isolate the code that is being generated.

Secondly, you said you don’t think the problem is with the “temp” variable, but rather with the “return”, but you don’t indicate what results you got when the operator function was defined as:

```
inline hmm operator + (register hmm a, register hmm b)
{
    return( a.value + b.value );
}
```

Also, I’d be interested to know what level of compiler optimization you have enabled on your command line options.

Mike Fulton @ SCEA

Re: Re: C++ Compiler bug slowing things down
Written by Mick West on Thu Jan 09 23:29:01 1997

Question

To Mike Fulton:

Ok, as you suggested, I put a dummy function around the + to delineate the function, though this did not actually affect anything.

I also changed the + operator so it just return the first value, for simplicity, this dies away with the temp variable

This is the code:

```
// only one operator defined, for simplicity
inline hmm operator + (register hmm a,register hmm b)
{ return a; // just returns a, so should all be optimized away
}
```

```
int g1;
hmm g2;
void dummy ()
{ g1++;}
```

```
// silly function, just illustrate the problem
void mick( hmm a, hmm b)
{
    dummy();
    g2 = a+b;
    dummy();
}
```

This is what it generates with -O0 (no optimization)

```
move fp,sp ; just lots of crap
sw a0,$18(fp)
sw a1,$1c(fp)
jal dummy__Fv
nop
lui v0,$802c
addiu v0,v0,-$2d00
move a0,v0
lw a1,$18(fp)
lw a2,$1c(fp)
jal __pl__FG3hmmT0
nop
jal dummy__Fv
nop
move sp,fp
lw ra,$14(sp)
lw fp,$10(sp)
addiu sp,sp,$18
jr ra
```

and with -O1 and -O2

```
addiu sp,sp,-$20
sw s0,$18(sp)
sw ra,$1c(sp)
jal dummy__Fv
move s0,a0
sw s0,$10(sp) <<<<< spurious sw
lui at,$802c
sw s0,-$2ff8(at)
jal dummy__Fv
nop
lw ra,$1c(sp)
lw s0,$18(sp)
addiu sp,sp,$20
jr ra
```

finally with -O3

```
lui v1,$802c
lw v1,-$2ffc(v1) <<< note it has placed the dummy function inline
addiu sp,sp,-$8
sw a0,0(sp) <<<< spurious sw
lui at,$802c
sw a0,-$2ff8(at)
addiu v0,v1,$1
addiu v1,v1,$2
lui at,$802c
sw v0,-$2ffc(at)
lui at,$802c
sw v1,-$2ffc(at)
addiu sp,sp,$8
jr ra
```

as an illustration of the greater problem, look at this

```
void mick( hmm a, hmm b)
{
hmm c,d,e,f;
c=a+b;
d=a+b+c;
e=a+b+c+d;
f=a+b+c+d+e;
g2 = a+b+c+d+e+f;
}
```

which compiles to:

```
addiu sp,sp,-$38
sw a0,$10(sp) <<< spurious write
sw a0,0(sp) <<< spurious write
sw a0,$18(sp) <<< spurious write
sw a0,$14(sp) <<< spurious write
```

```

sw a0,$4(sp) <<< spurious write
sw a0,$20(sp) <<< spurious write
sw a0,$1c(sp) <<< spurious write
sw a0,$18(sp) <<< spurious write
sw a0,$8(sp) <<< spurious write
sw a0,$28(sp) <<< spurious write
sw a0,$24(sp) <<< spurious write
sw a0,$1c(sp) <<< spurious write
sw a0,$20(sp) <<< spurious write
sw a0,$c(sp) <<< spurious write
sw a0,$30(sp) <<< spurious write
sw a0,$2c(sp) <<< spurious write
sw a0,$1c(sp) <<< spurious write
sw a0,$24(sp) <<< spurious write
sw a0,$28(sp) <<< spurious write
lui at,$802c
sw a0,-$3018(at)
addiu sp,sp,$38
jr ra

```

Which, I'm sure you will agree, is a little silly.

The writes to the stack will appear regardless of what the "+" operator does, so if was defined as

```

inline hmm operator + (register hmm a, register hmm b)
{
temp t;
t.value = a.value+b.value;
return t;
}

```

then it would still hav the same number of spurious writes to the stack.

We estimate that this is bloating our code size between 20% and 30%, and also impacting on the speed as writes to memory are very slow.

Does anyone have any ideas on either:

- 1) How to fix the compiler
- or
- 2) How to patch my class to avoid it happening.

All info gratefully recieved

Thanks,

Mick West (Mick@Neversoft.com)
Neversoft Entertainment.

Re: C++ Bugs/Language Features Bloat
Written by Jeremy Gordon on Sat Feb 22 18:09:11 1997

To Mick West:

Hi Mick,

We're using C++ for our project here and I would recommend the book 'Inside the C++ Object Model' by Stanley B. Lippman (one of the Bell labs guys that worked with Bjarne on the original CFront compiler and C++ language spec). His book talks about how C++ language features are typically implemented without getting too machine specific. The sections on the 'Named return value optimization' might be interesting reading.

It's definitely been an invaluable reference in understanding what C++ language features will do to our code. This probably doesn't really address your question at hand but I couldn't imagine attempting a large C++ project on an embedded system without it

Good luck!

- jeremy
jgordon@magicarts.com

Re: Re: Re: C++ Compiler bug slowing things down
Written by Mike Fulton on Thu Jan 23 02:04:39 1997

To Mick West:

I've looked over your sample code and frankly I don't think this is so much a bug in the compiler as a case where you are expecting certain optimizations that the compiler just doesn't try to do.

For one thing, your comment makes the assumption that the inline "+" function is going to be completely optimized away because it just returns back the value of the first argument. This is obviously not the case. There is a store instruction generated.

So the result where you get 20 store instructions in a row is simply due to having done that many additions in a row. The code doesn't make sense, but that's the way you have the operator defined, so it's not surprising.

You never did include the definition of the "hmm" class. That, of course, does affect the situation.

If you really wish to pursue this, send email to "support@snsys.com" and give them all the details. However, I'm pretty sure they'll tell you that it's a basic GNU R3000 issue, and not something for which they're responsible.

Aside from all of the above, I find it hard to believe that this could increase your code size by 30%, as you indicated. However, don't forget that you can always hand-code things in assembly language if you want.

Mike

aspsx listing file

Written by Michael Koziniak on Tue Jan 07 00:35:43 1997

From: Edgar Tu
Note: Moved from FirstClass

Does anyone know how to make aspsx generate a listing that contains both the source and assembly code together? Other gnu assembler seems to be able to do this with the -ahld option.

Edgar

Metrowerks Libraries

Written by Michael Koziniak on Tue Jan 07 00:29:42 1997

From: Bob Koon
Note: Moved from FirstClass

Uhm...I thought the Sony includes bundled with the CodeWarrior libraries are supposed to work with CodeWarrior. :/

I get the same problems using those new includes as I do with 3.5's includes; things like redefining CVECTOR, it not knowing what an SVECTOR is (even though it's in the same file), etc.

Is it something I'm doing/not doing in CodeWarrior??

- Bob Koon
Vortex Media Arts

Forum: Peripherals

StopCARD rules

Written by Dave Marshall on Tue Jun 03 19:35:44 1997

Are there any rules as to how/when StopCARD() should be called?

I have inited my controllers and memcards using the Sony recommended way: Initpad

Startpad InitCard(1) StartCard _bu_init() ChangeClearPad(0) open/enable memcard events

Later on, if I just call StopCARD(), the code hangs and never comes back from the StopCARD() call. If I put something like DrawSync(0), VSync(0), or printf right before the CardSTOP() call everything works fine.

Any thoughts?

dave

Memory card error messages?

Written by Jason Bell on Thu May 29 05:01:29 1997

My producer would like to know whether there are any Sony standards about the error message that should be displayed if there is an error while writing to a memory card. For instance, should we mention that the user might want to try a different card?

Thanks,

-Jason Bell

Re: Memory card error messages?
Written by Mike Fulton on Thu May 29 19:37:06 1997

To Jason Bell:

The Sony standards regarding memory cards are contained in section 12 of the document "Technical Requirements Checklist for PlayStation Software".

Beyond that, you can word your messages the way you want. But keep in mind that we reserve the right to suggest changes when the title is submitted. If you want to provide us with what you've come up prior to final submission, to get our feedback, that's OK too.

Mike @ SCEA

Memory card and SPU streaming?

Written by Ian Schmidt on Wed May 28 19:14:24 1997

We're using SPU streaming to play background music for our game. The problem is, this seems to interfere with large (12 blocks. I know that's a lot) memory card reads. Writes, formats, status detection, and reads of 1-5 blocks all appear to function properly with the streaming running, but the large read fails almost every time unless we stop the music first. We're going to try breaking apart the read into multiple single-block read(s) and see if that helps, but in the meantime are there any specific cautions with memory card access during SPU streaming?

Re: Memory card and SPU streaming?
Written by Rob Vawter on Fri May 30 17:29:10 1997

To Ian Schmidt:

I haven't heard of this problem before, so I don't know of any specific cautions. It does seem probable that the trio of callbacks associated with SPU streaming are interfering with the mem card event timing. You really can't do much about the Cd callback, except lower the sampling rate of your VAGs. You can increase the buffer size of your SPU stream buffers, causing fewer callbacks to occur (both the SPU DMA and the SPU IRQ). Lowering the pitch of the SPU stream would also cause the SPU IRQ and SPU DMA to be called less frequently.

However, all this said, I think your best solution is to decrease the block size of your mem card reads.

Memcard, formatting & guidelines
Written by Robert W. Calfee on Fri May 09 17:22:13 1997

Hi all!

I've discovered an interesting feature of "firstfile" and perhaps some other library functions. Although the developers guidelines expressly forbid the formatting of memory cards except with the consent of the user, the same apparently does not apply to the Sony library. Using the firstfile call on an unformatted card will blindly, quietly, and automatically format the card without confirmation or consent. This will cause your title to suffer rejection based on rules 12.4 - 12.6 (or whatever).

It is therefore important to determine that the card is formatted before handing it off to any Sony library call.

(and from now on I'm recommending passwords over memcard saves to all who will listen) rwc

>

Re: Memcard, formatting & guidelines
Written by Mike Fulton on Fri May 09 18:07:50 1997

To Robert W. Calfee:

That's interesting information about "firstfile" auto-formatting. I'll pass it on to see if that's being done intentionally or not. (I think it was rule 12.2.2 that you were referring to.)

However, you should not have gotten as far as calling firstfile() before determining that the card was not formatted. You may want to check that you are handling all card events properly.

Mike

PAD crash in Lib 3.7

Written by Dave Marshall on Wed May 07 18:35:40 1997

When using Lib 3.7, StopPAD() will crash if it is executed before a call to InitPAD(). This does not happen with 3.6

dave

Re: PAD crash in Lib 3.7
Written by Mike Fulton on Wed May 07 20:44:17 1997

To Dave Marshall:

You're not supposed to call StartPAD or StopPAD before calling InitPAD.

With earlier library versions, these functions simply set a flag which controlled whether or not the controller reading code actually fed data into the buffer.

However, as of library 3.7, there's more stuff going on under the hood. There's an interrupt handler patch being installed and removed by these functions, so even though you could get away with calling these functions out of sequence before, you can't now.

Mike

Parallel Port Information

Written by Byon Garrabrant on Tue May 06 00:33:40 1997

How can I learn more about using the PARALLEL I/O PORT on the back of the PSX?

Re: Parallel Port Information
Written by Mike Fulton on Tue May 06 22:50:14 1997

To Byon Garrabrant:

I'm not sure how much released information we have regarding the parallel port, and the only products I know of (off the top of my head) that work with it are developer-oriented.

I'll look into it, but it will probably take at least few days.

Do you mind sharing what you've got in mind? Via EMAIL if you like.

Mike

Memory Card & Kanji

Written by Alex Tyrer on Tue Apr 29 20:38:22 1997

Hello-

I'm not sure if this is the place to post this message, but here goes. I need some clarification...

I've just about finished our memory card module for saving our game between sessions, but have now noticed that I may have to support Kanjii (some kind of Unicode?) 16 bit encoding for the "title" field in our file's block header.

While this doesn't appear to be that difficult in and of itself, simply converting one encoding to another and back, I'm getting mightily worried when I read later in the documentation (Power Point slides from some conference) that our title MUST display all kanjii from other titles' saved game files. Is this totally necessary for a US released title? We were just planning to use ASCII characters, save them in 16bit Shift-JIS code, then convert them back when reading back in.

There appear to be over 6000 characters mentioned in the Kanji font docs (which I find terse and confusing to say the least). Can we just use a sub-set? Are we supposed to make a small subset of all displayable characters? What about vertical writing?

How do I display Kanjii? It looks like Greek to me (pun intended!) - how do I know if my code is displaying it correctly? I don't read Japanese.

What do we need to do at a minimum?

Help!

Alex

Re: Memory Card & Kanji
Written by Mike Fulton on Wed Apr 30 01:38:37 1997

To Alex Tyrer:

IF your game displays memory card stuff saved by other games at all, then it must correctly display any Kanji characters. The Kanji characters are provided in ROM for you and there is a function which gives you a means of locating any SHIFT-JIS character within the font so that you can display it.

Please see our sample programs.

Also, before this goes any farther, you should probably get together with your account executive here at SCEA and figure out what needs to be done exactly.

Mike

Dual analog joypad button confunction

Written by Ian Schmidt on Thu Apr 17 21:56:14 1997

I recently got one of the prototype dual-analog with force-feedback controllers, and I've discovered an anomaly with it. Specifically, when the controller is in plain analog mode (LED is green), several of the buttons are mixed around, as follows:

Square registers as L1

Triangle registers as R1

R1 registers as Triangle
R2 registers as Square
L1 registers as R2.

All other buttons continue to maintain their usual mapping., and this does not occur in force-feedback mode (LED is red) or digital mode (LED is off). Is this a bug in the pre-production controller?

Thanks,
Ian Schmidt

Re: Dual analog joypad button confunction
Written by Mike Fulton on Fri Apr 18 00:39:03 1997

To Ian Schmidt:

You have a misconception about the force feedback option that we must clear up before we talk about anything else.

There is no relationship between the color of the LED and the force feedback device: The LED merely indicates which controller ID is being used (and accordingly, what button and/or analog data is being returned).

Forcefeed back is always an option regardless of the LED color.

In fact, it's our recommendation that you write to the force feedback device without paying even the slightest attention to what sort of controller is out there. Either the device will vibrate or it won't.

The reason for this is that there may be other types of controller in the future which support the force feedback option, and they may have different controller ID codes.

Regarding the button mapping when the controller is in GREEN mode, SCEI is aware of the difference. The last I'd heard, they did not plan to change it.

You can express your opinion about this by sending E-mail to:

"bill_rehbock@interactive.sony.com". Bill will forward your mail to the appropriate people at SCEI. Please make sure you specify your company name, the name of the project you are working on, and the publisher.

However, it is something that shouldn't really matter for games currently in development, since they have no reason to specify that the user should be using GREEN mode. We recommend that in your manual, you simply specify that only RED mode should be used, not GREEN mode.

Mike Fulton

Sony Computer Entertainment America

Re: Re: Dual analog joypad button confunction
Written by Ian Schmidt on Fri Apr 18 15:29:49 1997

To Mike Fulton

Thanks for your reply. My next question then is: does the dual analog joystick currently on sale (which also has ID 0x53, if I'm not mistaken) have the same mixed-up button configuration?

Ian Schmidt
Tiburon Entertainment, Inc.

Re: Re: Re: Dual analog joypad button confunction
Written by Mike Fulton on Fri Apr 18 18:40:54 1997

To Ian Schmidt:

No, the big analog joystick controller does not have a mixed up button configuration.

The 0x53 mode of the small analog controller was supposed to be compatible with the big dual analog joystick controller, but somewhere along the line, the button mapping for that mode got scrambled. I don't know if this is something that happened by design or if it's just a bug. I just know they've decided not to change it. I haven't heard what the reasoning behind that decision is.

Like I said before, please do feel free to voice your opinion. The controller is not for sale yet, so it's still not too late. Send mail to Bill Rehbock

Mike

Re: Re: Re: Re: Dual analog joypad button confunction
Written by Mick West on Fri Apr 18 20:38:00 1997

To Mike Fulton:

Does this mean that there is no way of detecting if the 0x53 controller is a big analog stick or the small handheld one in "Green" mode with confused buttons?

Are Sony planning to issue any guidelines as to a standard way of dealing with this problem? Obviously we can just say "make sure the controller Led is Red" in the manual, but it seems rather weak to do that... especially as most people don't read the manual before playing games.

Mick.

Re: Dual analog joypad button confunction
Written by Mike Fulton on Mon Apr 21 18:35:54 1997

To Mick West:

Yes, you are correct that there is no way to detect a difference between the big dual analog joystick controller and the small dual analog controller.

Our recommendation is that new games have no reason to use the green mode, so their manual should indicate that only the red mode should be used.

I agree that it's a weak solution, but it's the best we can do unless SCEI changes their mind about fixing the controller.

Like I said before, please send EMAIL expressing your opinion to Bill Rehbock that he can pass along.

Mike Fulton

Combat cable TUTO1 matching sync
Written by Jim Buck on Fri Apr 04 22:27:43 1997

I was curious if it was absolutely necessary to rematch vsync periodically after the initial matching. What would be the repercussions of removing the call to sync_vsync_trger over a period of time? Since the reads block until a write from the other PSX is made, isn't there a virtual vsync rematch if one PSX were suddenly getting ahead of the other one anyway?

Re: Combat cable TUTO1 matching sync
Written by Mike Fulton on Mon Apr 07 19:46:52 1997

To Jim Buck:

First of all, can you please specify more precisely what sample program you're looking at? Unfortunately, we have so many things named "TUTO1" that it doesn't narrow it down enough. If you can specify the path on the CD or where else you got this sample from, I'll take a look at it.

Mike

Re: Re: Combat cable TUTO1 matching sync
Written by Jim Buck on Wed Apr 09 20:41:02 1997

To Mike Fulton:

I installed it to the following directory: PSX\SAMPLE\ETC\COMB\TUTO1... not sure where on the cd it would normally be.

It came from the files in COMBAT.ZIP that were put on the website on March 7 (well, the newsletter with the announcement was dated the 7th). They are only supposed to work with lib3.7 of the combat library.

Memcard and Exe modules
Written by Dave Marshall on Fri Mar 28 18:58:00 1997

I am trying to use the joypad and memory card libs in both a front end EXE module and a game EXE module. I am either not shutting down the joypad/card correctly or there is a bug because when I call StartCARD() in my game EXE module it overwrites part of my code. It writes joypad data into the

address that InitPAD() was given in the Front End EXE module. StartCARD() eventually does a SYSCALL 0 in the ROM somewhere. It is this call that writes the data into my code.

The following is the code that I use to init the card/joystick and shut it down. Let me know if I am missing something or doing it in the wrong order.

```
//startup code
{
EnterCriticalSection();
    /* OPEN ALL CARD EVENTS */
    //software events
ev0 = OpenEvent(SwCARD, EvSpIOE, EvMdNOINTR, NULL); ev1 = OpenEvent(SwCARD,
EvSpERROR, EvMdNOINTR, NULL); ev2 = OpenEvent(SwCARD, EvSpTIMOUT, EvMdNOINTR,
NULL); ev3 = OpenEvent(SwCARD, EvSpNEW, EvMdNOINTR, NULL);

//hardware events ev10 = OpenEvent(HwCARD, EvSpIOE, EvMdNOINTR, NULL); ev11 =
OpenEvent(HwCARD, EvSpERROR, EvMdNOINTR, NULL); ev12 = OpenEvent(HwCARD,
EvSpTIMOUT, EvMdNOINTR, NULL); ev13 = OpenEvent(HwCARD, EvSpNEW, EvMdNOINTR,
NULL);

/* INITIALIZE THE CARD SYSTEM */
InitCARD(1); StartCARD(); _bu_init(); ExitCriticalSection();

    EnableEvent(ev0);
    EnableEvent(ev1);
    EnableEvent(ev2);
    EnableEvent(ev3);
    EnableEvent(ev10);
    EnableEvent(ev11);
    EnableEvent(ev12);
    EnableEvent(ev13);

    InitPAD(&portbuffer[0][0],INPUT_BUFFER_LENGTH,
    &portbuffer[1][0],INPUT_BUFFER_LENGTH);
    StartPAD();
    ChangeClearPAD(0);
} //end of startup code

//shutdown code
{

StopCARD();
EnterCriticalSection();
/* DISABLE ALL CARD EVENTS */
DisableEvent(ev0);
DisableEvent(ev1);
DisableEvent(ev2);
DisableEvent(ev3);
DisableEvent(ev10);
DisableEvent(ev11);
```

```
DisableEvent(ev12);
DisableEvent(ev13);
ExitCriticalSection();
/* CLOSE ALL CARD EVENTS */
CloseEvent(ev0);
CloseEvent(ev1);
CloseEvent(ev2);
CloseEvent(ev3);
CloseEvent(ev10);
CloseEvent(ev11);
CloseEvent(ev12);
CloseEvent(ev13);
StopPAD();
} //end of shutdown code
dave
```

Re: Memcard and Exe modules
Written by Brian Dawson on Thu Apr 03 22:06:54 1997

To Dave Marshall:

As it turned out the problem Dave encountered was due to the order in which he initializing the Pad and Card. The proper order is InitPAD then InitCARD. For more details on the initialization order please refer to the 1.8 Tech. Ref. CD in DevConf\Winter95\A_Memcrd.PDF.

This is a PDF of the slides used in the March '95 Developers Conference Advanced Memory Card Seminar

Brian Dawson@SCEA Dev. Support

Re: Memcard and Exe modules
Written by Mike Fulton on Thu Apr 03 22:40:38 1997

To Dave Marshall:

I think the short answer here may simply be that within your game module, call InitPAD again before you call StartCARD or InitCARD. That will make sure that the buffer being used is the right one for the new module, and not the wrong one for the shell that's no longer there.

There's also some new undocumented stuff in library 3.7 that may apply. I don't know if it's undocumented on purpose or not and I've asked SCEI for some more info. I should probably hear something back from them by Monday or Tuesday.

In the meantime, try doing InitPAD first and see if that helps. I am not sure what the exact rule is for which should normally be done first, but I'll look into that and in the meantime try it and see what happens.

Mike

MemCard and Overlays

Written by Evan Ten Brink on Wed Mar 19 18:32:47 1997

I am starting to work on memory card support and was wondering if the code can be placed in an overlay and loaded only on startup and when saving. Is this a bad idea? What problems should I be aware of? Has anyone got it working? Thanks.

Re: MemCard and Overlays
Written by Brian Dawson on Thu Mar 20 01:32:53 1997

To Evan Ten Brink:

This should work fine. No caveats come immediately to mind... of course make sure you close your card events down when exiting the overlay module. I am going to try this here as well as look closer at Libcard to see what you may have to watch out for. If I come up with something I will post here as well as contact you.

Brian Dawson@SCEA

Combat Cable Docs

Written by Borden Covell on Wed Mar 12 00:21:16 1997

Can someone please direct me to the latest combat cable documentation. I'm using the new examples which call `_comb_control (5,0,0)` which is undocumented in the latest doc files that I have (from the 1.8 CD).

When I use the example code in my game the game dies when it tries to do an mdec decompression after these `_comb_control` calls, so I'd like to know more about them.

Thanks!

Re: Combat Cable Docs

Written by Mike Fulton on Thu Mar 13 03:47:06 1997

To Borden Covell:

_comb_control(5,0,0) just calls ResetGraph(5) to reset the graphics system.

It's equivalent to calling ResetGraph(0) except that that the usual ResetGraph printf() call is skipped. This includes a ResetCallback call.

The ResetGraph stuff is likely what's killing your MDEC decompression.

Can you specify what sample program you were looking at so I can take a look at it?

Mike

Re: Re: Combat Cable Docs
Written by Borden Covell on Thu Mar 13 17:50:14 1997

To Mike Fulton:

I am using code from the tuto2 example included in the latest combat cable examples on the Web Site.

Combat Cable

Written by David Minogue on Wed Mar 05 23:00:59 1997

How can I connect a combat cable between an DTL-H2000 development system (9-pin 'D' connector) and a blue debug station (8-pin card-edge)? Is there an available device or schematic for making this connection across the systems?

- David Minogue.

Re: Combat Cable
Written by Brian Dawson on Fri Mar 07 18:10:50 1997

To David Minogue:

David,
SCEA-R&D's Hardware Engineer will be contacting you to discuss this matter further.

Brian Dawson@SCEA Dev Support

Multi-tap Problems

Written by Andrew H. Marrinson on Wed Feb 19 22:55:41 1997

First, sorry to anyone who read the first version of this when I mis-posted it to the development environment forum.

I can't seem to get multi-tap working in my program. I run InitTap and StartTap as you would expect. What happens is that once in a while the pad buffer will look like a multi-tap is there, but only like once a second. The rest of the time it behaves as though no multi-tap is connected and just passes through whatever is on port A of the tap.

I checked the archive of old BBS messages, but there was nothing quite like this. I'm stymied.

This is with an H2500, H2080, and lib3.8.

Thanks,

-andy

Re: Multi-tap Problems
Written by Mike Fulton on Fri Feb 21 00:18:17 1997

To Andrew H. Marrinson:

I'm sorry Andy, did you say library 3.8? I'm sure that must be a typo since we just released the BETA 3.7. Can you clarify which library version you are using please?

Also, if you can post the code for initializing the controller (InitTAP, etc.)

By the way, the "InitTap" function was replaced by "InitTAP" in library 3.6 (I think it was 3.6).

Which are you using? The same is true for some other *Tap functions.

Mike

Re: Re: Multi-tap Problems
Written by Andrew H. Marrinson on Fri Feb 21 16:05:04 1997

To Mike Fulton:

Yikes, I should be more careful. I meant lib3.6. I combined the CD-ROM version (1.8) with the library version. Indeed, I am using InitTAP and friends. Here is my init code:

```
#define PADBUFLen 8
#define PADPERTAP 4
#define TAPOFFSET 2
#define TAPBUFLen (PADBUFLen*PADPERTAP+TAPOFFSET) #define MAXPAD 2 // pad
number is zero-based // Control Pad Buffers volatile unsigned char
padBuf[TAPBUFLen*MAXPAD+PADBUFLen];

long PSX_InitPad() // returns 1 if successful, 0 otherwise {
long result;
```

```
padBuf[68] = 0xff; // used when a multi-tap port is queried and no
padBuf[69] = 0x00; // multi-tap is connected
result = InitTAP((char*) &padBuf[0],TAPBUFLLEN,(char*)
&padBuf[TAPBUFLLEN],TAPBUFLLEN);
if (!result)
return(0); // Init error
result = StartTAP();
return(result);
}
```

Hm, I just noticed InitTAP is documented as having no return value. (Wish it were declared void.) I don't think that is my problem though, since it does fill the buffers, it just ignores the presence of the multi-tap most of the time. (Once every few seconds it suddenly sees the multi-tap for a VERY brief time.)

Again, sorry for the confusion over version numbers.

- andy

Re: Re: Re: Multi-tap Problems
Written by Mike Fulton on Tue Mar 04 02:41:52 1997

To Andrew H. Marrinson:

InitTAP has an undefined return value, but it is not defined as "void" in the library code.

Everything else looks OK, but I don't understand what the following lines of code from your example are supposed to do:

```
padBuf[68] = 0xff; // used when a multi-tap port is queried and no
padBuf[69] = 0x00; // multi-tap is connected
```

Mike

Re: Re: Re: Multi-tap Problems
Written by Andrew H. Marrinson on Tue Mar 04 21:34:41 1997

To Mike Fulton:

You ask why I have this code in my initialization routine:

```
padBuf[68] = 0xff;
padBuf[69] = 0x00;
```

Basically, in my read routine, if the caller requests controller 2-A for example, and there isn't a multi-tap in slot 2, I return a pointer to padBuf[68]. The caller then sees that the first byte is 0xff and concludes that no controller is connected.

I didn't want to just return a pointer to the appropriate slot, since the documentation doesn't say what will be in those slots if no multi-tap is present.

Hope that makes sense. (If I post my read pad code, it will be more obvious.)

Anyway, what you are saying is that my initialization looks right except for testing the result of InitTAP. But that isn't the problem, since if it was nothing would work (since I wouldn't be calling StartTAP), whereas I have everything but MultiTAP working.

So what is going wrong? Any suggestions what to look for? Do you think there is a problem with the library code or my program? Or could it be the 2080 box or some other part of the hardware?

- andy

Re: Re: Re: Re: Multi-tap Problems
Written by Mike Fulton on Sat Mar 08 00:05:41 1997

To Andrew H. Marrinson:

Did you say you were using a DTL-H2000 ISA card set or a DTL-H2500 PCI card?

There's been some problem that I've been looking into with the Multitap for several weeks. I had mentioned it to the engineers at SCEI, but they were unable to reproduce the results. But I have continued to mess around with it, and I think I finally may have figured out something.

Over the next several days I'll be trying out a few changes to the library code and hopefully this will do the trick. I'll post another message to you when I have figured out anything, but if you want to know the status in the meantime, or if you have additional information about exactly what you're experiencing, please let me know.

Mike

Re: Re: Re: Re: Re: Multi-tap Problems
Written by Andrew H. Marrinson on Fri Mar 14 15:55:07 1997

To Mike Fulton:

Sorry for the delay in my response. I hadn't been checking the web site on this recently because Brian Dawson has been talking to me over the phone about it.

Today I was finally able to borrow some time on a machine with an old 2000 card in it. I had been using a 2500. Anyway, the problem is the same on both the 2000 and 2500. I'll phone Brian and let him know as well.

My pad reading code is pretty straight-forward. If it is something in our game, I suspect it is some other aspect that is affecting the pad reading as a side-effect. I just don't know where to start looking.

- andy

Re: Re: Multi-tap Problems
Written by Christopher Deel on Thu Feb 27 21:12:20 1997

To Andrew H. Marrinson:

I am having the same problems with the MultiTap and Library 3.6. It even happens with the sample controller code from the developer CD.

Re: Re: Re: Multi-tap Problems
Written by Andrew H. Marrinson on Fri Feb 28 13:17:14 1997

To Christopher Deel:

Hm,. I had problems with the sample code until I recompiled with libtap instead of libgun (and changed the code accordingly). Perhaps there is a clue in that—I should have mentioned it earlier.

It is gratifying to hear that I am not the only one having these problems.

- andy

Boot from OSD fails during mem card events
Written by Rob Vawter on Wed Feb 19 17:59:04 1997

chaplin:

We're doing a PAL version of Virtual Pool, and when the program boots up from the memory card/CD player screen, it hangs during the memory card initialization. Is there a difference in how the memory card expects to be initialized for PAL boxes? Mind you, this is a blue box, not a production unit but I doubt there would be a difference between the production and blue boxes in this matter. Unfortunately, it doesn't even get to doing anything with memory card blocks. It hangs during the initialization of the memory card events! Basically, when the user boots up from the memory card/CD Player screen, then closes the CD tray, the program hangs during my memory card initialization. It's pretty standard code, taken from the sample programs and works perfectly for the NTSC versions.

vawter:This problem has been reported a few times before. In the past, it has always been a problem with initialization. Calling InitCard(1) (which is in the sample code between opening the events and enabling them) before controller initialization is the most likely cause of the problem. Let me know otherwise.

Memory Card

Written by Matt Whiting on Thu Feb 13 02:37:08 1997

We are trying to write more than 2k at a time out of 8k block and we get failure on write. Do we need to write to 2k at a time or is this a bug? Thanks for your prompt response

Re: Memory Card
Written by Eric W. Brown on Tue May 20 18:32:06 1997

To Matt Whiting:

Hi, we've had the same problem writing to the memory card. We are using one of the root counters for millisecond timing, and I was able to figure out that the problem only happened when the root counter was active. When I changed the program so that we shut down the root counter before writing to the memory card, the program worked.

I've also heard you should VSync(3) between the open and the write.

Eric W. Brown

eric@virtualmusic.com

Re: Memory Card
Written by Brian Dawson on Mon Mar 03 21:31:39 1997

Matt resolved this problem by padding out the size of his writes to be evenly divisible by 128.

I will be looking further into this and posting my findings as soon as practical.

Brian Dawson

Re: Memory Card
Written by Brian Dawson on Wed Feb 19 00:46:54 1997

To Matt Whiting:

Hey Matt, per our phone conversation I will expecting a code snippet. I'll look over it and get back with as soon as i have solution. Good talking with you.

Brian Dawson@SCEA Dev Support

Re: Re: Memory Card
Written by Brian Dawson on Thu Apr 03 22:17:10 1997

To Brian Dawson:

Through test I concluded that when writing to a block device such as the Memory Card, any writes not a multiple of the devices block size will not be carried out, yet will not return a negative value indicating a failure. Note that these findings apparently conflict with the current documentation, which implies that only the data over a multiple of 128 bytes will be discarded, not the whole of the data.

Brian Dawson@SCEA Dev. Support

New Dual thumb pad Analog Joypad

Written by Chris Reese on Wed Jan 29 20:01:35 1997

Is there any information about this new joypad that is available. We have plans to support it in our title and would like to know things like button configuration and if it supports both analog and digital modes.

Re: New Dual thumb pad Analog Joypad
Written by Mike Fulton on Wed Feb 12 02:29:56 1997

To Chris Reese:

See the anlgctr.zip file here on the web site for a sample program and documentation.

Re: Re: New Dual thumb pad Analog Joypad
Written by Mick West on Mon Mar 10 02:17:50 1997

To Mike Fulton:

Where exactly is the "anlgctr.zip" file then?

Mick.

Re: Re: Re: New Dual thumb pad Analog Joypad
Written by Peter Alau on Mon Mar 10 17:27:01 1997

To Mick West:

That file is now located within the Sample3.zip file that is located under the programmer tools.

- Peter

Re: Re: Re: Re: New Dual thumb pad Analog Joypad
Written by Mick West on Mon Mar 10 19:33:56 1997

To Peter Alau:

You mean I have to download 11 MB, just to get a teeny little anlgctr.zip file? Shurley shome mishtake?

Still, needs must, thanks for the pointer.

sigh..... time to upgrade to a T1.....

Mick.

Re: Re: Re: Re: Re: New Dual thumb pad Analog Joypad
Written by Mike Fulton on Mon Mar 10 22:10:54 1997

To Mick West:

I'm sorry, but I didn't realize that the controller sample had been folded into such a big file, and that this particular archive file wasn't setup for downloading the individual files within.

I just raised a Red Flag about fixing the situation.

In the meantime, I've attached the program to this message for you to download.

Mike Fulton

Re: Re: Re: New Dual thumb pad Analog Joypad
Written by Mike Fulton on Mon Mar 10 05:08:01 1997

To Mick West:

I don't know exactly, but I was told it had been posted. I'll have the person who posted it respond to you.

Mike

Re: New Dual thumb pad Analog Joypad
Written by Rob Vawter on Fri Jan 31 20:59:18 1997

To Chris Reese:

I think this might be what you are looking for:

The following announcement was sent to us from Sony Computer Entertainment Japan:

New Model Analog Controller(Revised Version:11/20/96)

November 20, 1996

Sony Computer Entertainment Inc.

R & D Vice President

Shin'ichi Okamoto

“PlayStation Expo” was held from November^{3rd} till 4th at Makuhari Messe. We, SCEI exhibited the new model analog controller with vibration device for reference. Following is the release schedule and technical information of this new controller:

[1] Overview

The new controller is a standard controller with two 2-axis(horizontal and vertical directions) analog levers and vibration device to support even wider range of games. All software functions except for the vibration device are either equivalent to the analog joystick mode or analog controller mode with additional buttons allocated to the analog levers.

[2] Design Overview

The new controller is a standard controller plus 2 analog levers. One lever is placed between “SELECT” and downward directional buttons, and the other between “START” and “X” buttons.

[3] Release

As of today, the detail schedule for both tool and mass production has not been determined. Please contact SCE Business Affairs if you have a plan to support the new model controller for titles that will be released by the end of this year.

[4] Technical Information

4-A Controller ID

Three controller IDs below can be selected with a mechanical switch:

- 0x41: Standard Controller
- 0x53: Analog Joystick
- 0x73: Analog Controller

4-B Vibration device

The vibration device can be set ON/OFF by only a program software. From safety point of view, when there is no data transmission with the controller for more than 140 msec, the vibration device will be set off.

4-C Updated Library (LIBAPI)

As stated in the section 4-D, a new LIBAPI function for sending data to the controller is required for the ON/OFF operations. The updated LIBAPI library is planned to be released at the same time as the controller with the vibration device for programming.

Since data can be sent in parallel with receiving data(e.g. status data

of the controller button) at every V-Sync interruption, it will not affect the CPU time available for an application. Moreover, sending ON/OFF data to the controller without the vibration device will not cause any problem.

4-D Specification of the New Library Function

SendPAD Sets controller send buffer

Syntax

```
void SendPAD(bufA, lenA, bufB, lenB )
```

```
char *bufA, *bufB;
```

```
long lenA, lenB;
```

Arguments

bufA, bufB Buffer for sending data

lenA, lenB Length (in bytes) of the buffer for sending data Explanation

This function registers sending data buffer for controller

At every V-Sync interruption, the target device ID and data set to the buffer is sent. The buffer becomes invalid once it is sent.

=====

Byte Description

0 Flag to specify validity of the buffer data

0x00 Invalid, 0x01 Valid Other values Undefined

1 Target Device ID

2 - Data to be Sent(Sending data length = Received data length)

Return Value

None

Remarks

Buffer Setting for New Analog Controller with Vibration device:

=====

Target Device ID 0x03(Vibration), others:Vibration Off

Data to be sent (First Byte) 0x01:Vibration ON, 0x00:OFF

Others: Undefined

(Second byte and after) Always 0x00,others: Undefined

Warning:

Vibration device can be set on only within one V-Sync interval;
between two V-Sync interruptions.

Thus to keep the vibration device on, following must be set at
every V-Sync interruption.

- Target device ID
- Vibration: ON
- Buffer data flag: VALID

4-D Maximum Number of Concurrent Connections

Up to 2 controllers without multi-tap and 8 controllers with two
multi-taps can be concurrently connected. This limit is equivalent
to the standard controller.

Due to electricity limitation, only two controllers can set the vibration
device on at the same time. This maximum number may be changed to the
larger number upon release.

4-F How to operate Vibration device with multi-tap

To operate the new analog controller with vibration device when it is
connected to a multi-tap, you need to use data transmission function
SendTAP, in LIBTAP that will be released at the same time as the
development hardware with the vibration device. SendTAP has the same
functional specification as SendPAD.

Following shows the details of send buffer used for SendTAP function:

=====

Byte Description

0 Flag to specify validity of the buffer data

0x00 Invalid, 0x01 Valid Other values Undefined

1 Target device ID of Controller A

2-6 Data to be sent to Controller A

7 Target device ID of Controller B

8-12 Data to be sent to Controller B

13 Target device ID of Controller C

14-18 Data to be sent to Controller C

19 Target device ID of Controller D

20-24 Data to be sent to Controller D

To send a data in a buffer;

- Write 0x01 into the byte 0
- Call SendTAP function

Upon V-Sync interruption immediately after the function call, the data in a buffer is sent.

The target device ID (0x03 for Vibration device) and contents of the data to be sent(0x01 for ON) are the same for SendPAD function.

4-G Calibration

In Japan, the current acceptance criteria for the master titles that support the analog controller requires to include at least the built-in calibration function described below, and the same criteria applies to the analog controller with the vibration device as well:

=====

0 position calibration

Setting the play(space for movement)

Sensitivity

Detection of max. stick movement value

4-H Life of Vibration Device

The life of the vibration device depends on the number of ON/OFF operations.

It is highly recommended not to set the vibration device ON and OFF many times as it extremely shorten the life.

***** End of document *****

LibCD in 3.6 causes many errors

Written by Eric W. Brown on Tue May 20 23:22:40 1997

Our code works great with the 3.4 libraries. When we link with 3.6, we get dozens of CD errors; the error messages indicate that it is retrying the operation over and over. We tried a fix (posted on the old BBS) of putting a call to VSync(3) after any place the speed could be changed. We also tried inserting CDSync before all read operations. Neither of these seemed to make any difference. Are there any known problems like this? We are using the root counters for our own millisecond timing; would this be causing the problem? NOTE: I have a little more information on this. The errors appear to be happening inside of CDReadSync. The text of the error messages is "CdRead: retry..." followed by a number of lines "DiskError: com=CdlPause,code=(mm:nn)" where mm:nn is numbers like 43:80 or 06:10, usually repeated several times.

Re: LibCD in 3.6 causes many errors
Written by Brian Leake on Wed Jun 04 04:41:49 1997

To Eric W. Brown:

I also found a number of problems when moving to the 3.6 libs. I changed my .CTI file to use

XA attributes, for example -

File myfile.bin

XAFileAttributes Form1 data

Source c:\data\myfile.bin

EndFile

You'll need to specify the CD as XA, ie -

Disc XA_PSX CD.EMU

Try this, it might help with your read errors.

- Brian.

Re: LibCD in 3.6 causes many errors
Written by Mike Fulton on Wed May 21 02:14:38 1997

To Eric W. Brown:

You don't say if your code is getting back CD errors, or if the printf output from the library is simply reporting internal errors.

That's important because in library 3.5-up, the CD library code is changed somewhat from 3.4, and one of the changes is that it's more verbose in printing debug information. CD errors of one type or

another happen all the time, and it's generally no big deal provided you are handling errors properly, which the library does.

It's unlikely your timer is having any effect on the CD code, unless you are possibly spending so much time within a timer interrupt that you are missing a CD interrupt somewhere along the way.

One thing you need to make sure of is that you are not doing any CD-related calls except CdReadSync until the read is completed. Or if you've setup a CdReadCallback routine, don't do anything until it's been called.

If you do anything else, you might step on the library's toes.

The way CDRead works is that it sets up CdReadyCallback and then issues a low-level read command. After each sector becomes available, the callback checks for an error, and if none has happened, it transfers the data to the appropriate position within the specified buffer.

If an error does occur during the read of a given sector, then the callback routine will pause the disc, reset the current location to the sector where the error occurred, and then reissue the read. It will retry that sector up to 10 times before giving up and then returning an error to CdSync or the CdSyncCallback routine.

After the last sector has been transferred, then the callback issues a command to pause the disc.

I'd have to see your CD code itself to get a good idea of what the problem might be, but chances are we can figure it out without too much trouble. Why don't you give the SCEA developer support line a call and ask for me so we can talk about it.

Mike Fulton @ SCEA

Small width videos

Written by Ron Maxwell on Wed May 07 15:26:42 1997

Does anyone know how to display a video that has a width smaller than 16 pixels? I tried displaying a video that was 8 pixels wide and the decompressor returns a 16 pixel slice that has the data doubled.

Any help would be appreciated.

Thanx,
Ron Maxwell
Virtual Music Entertainment, Inc.
ron@virtualmusic.com

PS Could you email me your response as well.

Re: Small width videos
Written by David Minogue on Thu May 08 13:25:07 1997

To Ron Maxwell:

Another possibility would be to not use the MDEC at all; just stream uncompressed video.

Even if your video is 240 pixels high at 30Hz, your data rate would be only:

$8 \times 240 \text{ (pixels/frame)} \times 2 \text{ (bytes/pixel)} \times 30 \text{ (frames/sec)} = 115,200 \text{ bytes/sec}$

Since you have 300KB/sec at double speed, that leaves plenty of bandwidth for interleaved XA audio. Of course, there's more work involved in creating the data and rewriting the player. And the quality won't improve much since the compression rate in the MDEC would be very low with such a small video.

- David.

Re: Small width videos
Written by Mike Fulton on Wed May 07 21:07:31 1997

To Ron Maxwell:

The MDEC has to deal with things that are increments of 16 pixels wide and tall.

However, you can compress something larger, and then when it's time to transfer the decompressed data to VRAM, don't transfer the whole thing.

Mike

Re: Small width videos
Written by Dave on Wed May 07 17:53:42 1997

To Ron Maxwell:

Str files have to be a multiple of 16 bytes. I suggest you encode your video with some extra padding to make it 16 pixels wide and then decompress into a vram buffer prior to copying the required area to the display buffer. You may be able to combine two or more source images into the target image to get the width up to a multiple of 16 bytes, and then select which image you display from the vram buffer. This will be a more efficient use of disk space and may result in reduced seek times. Also you can play back multiple images at the same time.

DaveC

CdGetSector
Written by bd on Thu May 01 17:38:48 1997

In the following code snippet:

```
...
char sector[2048];
char lastByte;

CdGetSector(sector, 512);
lastByte = sector[2047];
...
```

Is last byte always correct? Is CdGetSector always synchronous?

Thanks!

Re: CdGetSector
Written by Mike Fulton on Thu May 01 20:04:30 1997

To bd:

To be absolutely safe, you should do CdDataSync(0,0) immediately after the CdGetSector, in case there are changes to the way CdGetSector works. It's documented as being non-blocking, and maybe someday it will be changed to work that way even though it doesn't now. In the meantime, CdDataSync will just check a register, see it's not set, then immediately return, so it's not really taking much time.

Also note that instead of specifying a constant of 512 in your CdGetSector call, you should make sure you have specified the right value for the current CD mode. I mean, if you've specified mode 1 and you're reading 2340 bytes per sector instead of 2048, then you need to read more data.

Mike

Playing Multichannel ADPCM
Written by Borden Covell on Thu May 01 16:14:12 1997

Two questions about playing multi-channelXA:

1. Is there a way to combine music tracks that are not the same length into one ADPCM file?
2. What is the meaning of the member file in the CdIFilter structure?

Re: Playing Multichannel ADPCM
Written by Rob Vawter on Thu May 01 17:47:35 1997

To Borden Covell:

- 1) Yes, MovPack will pad out the lengths of the shorter tracks with dummy data (zero data I believe)
- 2) The file member of the CdIFilter structure is always equal to 1. I'm not sure what it was originally designed for.

Re: Re: Playing Multichannel ADPCM
Written by Borden Covell on Fri May 02 16:22:23 1997

To Rob Vawter:

Thanks. I was hoping you would say that...

Now, I've thought of another question.

I know how to detect the end of DA play (check status of CD) but what is the minimum code required to detect the end of the ADPCM data. Does it require reading data from the CD subsystem and checking the sector number in the header or is there a simpler (read less overhead) way?

Thanks!

Re: Re: Re: Playing Multichannel ADPCM
Written by Rob Vawter on Mon May 05 14:03:37 1997

To Borden Covell:

here are two ways to detect the end of XA audio playback:

1) The less precise, but easier to code polling method (See `psx\tuto\repeat.c` for an example, specifically the functions `cdRepeatXA()` and `cbready()`). Basically, an end position is set, and when polling detects that the cd is beyond that point, the oven timer goes off. Please remove the code checking if the cd is before the start point of the track however, as that can very occasionally cause a skip in the replay of XA. The skip is caused by the following: as the head searches for the start of the XA track, it may end up just a hair before the start of the track during the poll; then the `cbready()` function has decided a new seek is in order just as the start of XA playback begins, so the first nibble of the XA track is played back and then the track restarts. Bad. Also, since you are only polling, the end of the track may be reached, and if you are really unlucky AND happen to have another XA track following the first, playback of the second may begin before the reseek of the beginning of the desired track, giving you a slight glitch of noise. Also bad.)

2) The more precise, but more annoying to keep track of `CdReady` callback. (See `psx\tuto\scee\cd\xaplayer\xaplay.c` for an example. With this method, at least one channel of the XA data must be left blank or filled with data. Each sector as this portion of the XA track is read through, a `CdldataReady` will be returned. Let's say you have an XA track which is mono 37.8, so you can interleave up to 16 channels. So, you interleave 15 and leave the 15th slot empty. Now, everytime 1 sector of XA audio is played, you will get a happy little `CdldataReady` result. The annoying part of this process occurs if some of the XA channels are different lengths. For example, if you played the longest channel out of your 15, eventually, all other channels will run out of XA audio. When they run out, they are padded with zero data. Your `CdReady` callback will be returning more `CdldataReady` results per sector. First one, then two per audio sector, until you are receiving 15 per audio sector after all other XA channels have run out. You need to keep track of how many sectors in length each of the XA channels are for this method to work properly, but it is a very precise method.)

Enjoy.

Strange Sector Header

Written by Graeme Ing on Wed Apr 23 00:20:52 1997

When I do a `CdGetSector()` I'm getting 12 bytes at the beginning of my data. I can't figure out what these 12 bytes are. The code works fine under LIB3.5, it breaks when I installed LIB3.6. What has changed? The only mode bit I set is `CdldataModeSpeed`, and I want it in the default mode of 2048 bytes/sec.

Graeme

Re: Strange Sector Header
Written by Mike Fulton on Wed Apr 23 00:58:20 1997

To Graeme Ing:

It sounds to me like you've specified "CdIModeSize1" to request 2340 bytes of data per sector.

This would give you 12 bytes of data which includes the sector's timecode, followed by 2048 bytes of your regular data, followed by another 280 bytes of data which is normally used for error correction code, ISO9660 file identification, etc.

The CdIModeSpeed1 flag is used internally by CdRead() with newer versions of the libraries, but you wouldn't be doing your own CdGetSector() calls in that case (or at least you shouldn't be doing them).

Can you outline what you're trying?

Mike

Re: Re: Strange Sector Header
Written by Graeme Ing on Wed Apr 23 02:18:41 1997

To Mike Fulton:

Interesting. I'm using CdControl(CdIReadN,x,x) to read my data sector by sector and using the callback to pull each sector out of the buffer. Standard stuff. I'm not using CdRead() anywhere, but I *DO* perform a CdSearchFile() to reliably locate the start sector of the data I want. Once I know the offset of the data into the track, using CdSearchFile(), I do my own seeks and reads directly. Is CdSearchFile() invoking the high-level code that could be setting CdIModeSpeed1? This seems very likely. If so, do you recommend that I set the Mode back to 2048 *after* the CdSearchFile(), or that I use another method to locate the start of my data.

Thanks,
Graeme

Re: Re: Re: Strange Sector Header
Written by Mike Fulton on Wed Apr 23 23:48:41 1997

To Graeme Ing:

CdSearchFile uses CdRead internally.

The basic rule is, don't interleave high-level CD calls like CdRead and CdSearchFile with low level calls like CdControl, etc. It's OK to use both types of call, but make sure you don't have them interleaved together.

Solution!

Written by Graeme Ing on Wed Apr 23 16:12:38 1997

As my previous message surmised, CdSearchFile() does indeed set the sector size to > 2048!! The solution: Simply do the CdSearchFile() *before* setting the mode flags. This of course is more elegant because CdSearchFile() may play with the other flags too, such as CdIModeSpeed.

Graeme

LibPress: A Memory Devouring Monster..?

Written by Alex Tyrer on Tue Apr 22 00:10:33 1997

Hello-

I finally got our movie w/sound streaming all working (thanks for the helpful pointers), however when I linked the code onto the front-end of our program I noticed a whopping increase in memory usage. The problem doesn't seem to be the code, but rather these mysterious globals lurking in the 'TABLE' module of libPress (section of our map file)...

```
8006C5E8 DCL_DVLC <<--- these
8006C9E8 DCC_DVLC <<--- four
8006CDE8 CF_DVLC <<--- big
8007CDE8 CF2_DVLC <<--- tables
8007D5E8 gControlEvents
```

Together that's nearly 70,000(!*?!) bytes and it's killing our memory budget. Are all four of these tables really necessary? Are they used for DEcoding or ENcoding or both? Can we strip any of them? Or maybe load the data from a file into dynamically allocated memory? I'm allocating the sector ring buffers, VLC and image double buffers dynamically so it obviously came as a shock when these 'mystery' static buffers were linked in to our program. 70k of static global data is a _LOT_ of memory! Before this we only used about 9k globals. Any advice/suggestions would be most welcome! Alex b-(

Re: LibPress: A Memory Devouring Monster..?
Written by Mike Fulton on Tue Apr 22 01:19:38 1997

To Alex Tyrer:

These tables are absolutely necessary for movie playback. They contain the huffman decoding tables that are used to decompress the first stage of the encoded movie data.

However, there is no requirement that you keep these tables around in memory all the time.

They are only required when you are decoding MDEC data.

My suggestion is that you load your game into memory on the dev system, then write out the area of memory containing these tables back to a PC file. This gives you a raw binary file with the table data.

Then you can use the table area to load other data that is used when not playing movies, and swap the two sets of data in and out (that's assuming you have other data which you can dump and reload this way).

Optionally, you probably have a few requirements for large scratch buffers, primitive buffers, or something like that, where the memory isn't being used when playing movies. Simply use a pointer to access these items, rather than using a statically allocated buffer, or instead of using malloc(), and then set the pointer to point to a position within these tables. (Don't forget to zero the memory out.) In most cases, this is the easiest route.

Mike

Re: Re: LibPress: A Memory Devouring Monster..?

Written by Alex Tyrer on Tue Apr 22 02:47:20 1997

To Mike Fulton:

Thanks for the info (now at least I know what the data is). I'd have put the Huffman/String table dictionaries on the front of the STR data file (or embedded in the stream if it was needed to change dynamically - where I assumed it was until now). That's how we used to compress our art in ROM way back in the 16bit (S*G* G*n*s*s) days. It's how pkzip/lzh work (mostly), and allows for smaller tables (you throw it away when the codes get too big and start over with a new dictionary). All I see in the memory area of these globals are acres of very similar repeating codes, with some smaller denser areas of variability. Is this data produced algorithmically or empirically/statistically? If it's algorithmically then presumably the code that generates it could be included in the libs at some future date. As to the suggestion solution - it will work provided the tables continue to be contiguous, and we'll have to figure out the size of these items from the map file manually - they aren't declared/documented anywhere I've looked. Are they likely to change contents/size? A simple addition to the library would be to provide a couple of interface functions much like the ones we already have to modify the quantization matrices (ie DecDCTGetEnv & DecDCTPutEnv) to at least get the address & size of the tables (maybe combine them all into an anonymous structure so we'll _know_ that the linker can't separate them). Ah well... Thanks again! Alex

Re: LibPress: A Memory Devouring Monster..?

Written by Mike Fulton on Tue Apr 22 18:32:58 1997

To Alex Tyrer:

First of all, please use paragraphs! It makes it hard to read when everything is bunched together.

The Huffman tables are not embedded in every movie for two reasons. First, they are fixed and do not change from one movie to the next, so why duplicate this data again and again? If you're going to read the tables off the CD anyway, why not do it from a separate file that you only need once?

Second, the movie file format and the way the streaming libraries read data isn't really designed to have a big chunk of data at the start of a movie that doesn't go into the streaming ring buffer.

The Huffman tables used by LIBPRESS were decided on by compressing dozens of different movies and taking the tables that provided the best compression on average. As far as generating them on the fly goes, you can't do that for decompression.

Please keep in mind that the MDEC stuff uses standard Huffman encoding. The schemes with floating code tables like LZH are a bit more efficient in terms of compression ratios. They are also slower to decompress than Huffman encoding, which uses a static lookup table. Since the Huffman stage of the decompression process is handled by the R3000 and not the MDEC hardware, we need to use the fastest possible algorithm.

The tables occupy contiguous memory locations starting from the address of label "DCL_DVLC". From there, you have 69680 bytes which are used by the tables. The tables are all defined together in the same object module; the linker cannot separate them.

The Huffman tables are treated much differently from the quantization matrices. Those are actually stored on the MDEC chip itself, not within RAM. There's also more reason to play around with them, since you can control the movie quality by adjusting the quantization tables. (It is something of a black art, to be sure, but you can do it.)

I've created a data file (HUFFMAN.DAT) which contains the raw data for the Huffman tables. You'll find it in the attached archive. You can load this file to restore the tables when you need to decompress a still image or movie.

Mike

Name: Huffman.zip
Content-type: application/x-zip-compressed
Size: 7605 bytes

Re: Re: LibPress: A Memory Devouring Monster?
Written by Alex Tyrer on Tue Apr 22 23:26:40 1997

To Mike Fulton:

Thanks Mike.

I think we can work around our memory problems using this technique. It requires a little work to patch this memory into the rest of our contiguous managed memory heap, but that's just a SMOP (simple matter of programming).

In the interim I changed our game to boot the movie player which then Load/Exec's the game once it's finished or interrupted. I only lose 256 bytes for the parent's stack, plus the added couple of seconds of seek/load time, which is good enough for now.

Ta Ta for now,
Alex

str Problems

Written by barbm on Thu Apr 17 18:46:34 1997

I am trying to get an .str to run in an application. It runs fine as a stand-alone using the tuto1.c sample program. In the application, the code gets stuck in the while loop in strNext() waiting for a return code = 0. It never is. This code used to work before the 10-3-96 version of the psx libraries were installed. Can anyone help?

Re: str Problems

Written by Mike Fulton on Fri Apr 18 00:44:29 1997

To barbm:

I understand that you spoke with Mike Koziniak here at SCEA earlier and that we determined that the problem was due to some portion of memory that belongs to the streaming code becoming corrupted.

Please let us know when you've tracked it down and have it working.

Mike

Streaming Audio & Video

Written by Alex Tyrer on Mon Apr 14 23:52:26 1997

Hello peeps-

Notwithstanding that PSX audio isn't my strong point (yet at least anyway) hence my apologies if this seems like a bunch of dumb questions - I'm trying to stream movies with audio_and_ video. I've successfully built (I think) STR files from uncompressed AVI files, both with and without audio tracks using MovConv (new version). What I end up with, I presume, is an STR file containing both the sound and image data. I've got the movie up and running (playing _NONE_ sound-containing STR files), but am now confused as to how get at the sound data which I'm presuming has been interleaved into the STR file. Do I need to record a special Mode2 or XA form2 track for this STR file to reside on, or can I just add it with CDGen along with all our regular files? Do I need to run MovPack? I've looked everywhere for an example of this and can find none. Also the (pdf) docs briefly hint that ADPCM data can be encoded into the STR format but provide no clue as to how this might be decoded, but rather concentrate on the image compression algorithms. It seems like this topic (streaming AV movies from CD) encompasses CD streaming, MDEC decoding and SPU streaming, so I hope this is the correct forum to post my questions. I also don't have access to

the old BBS (assuming it's still up) - no dial out facilities, so I can't go look there for solutions to this (probably well worn) issue. Please help. I'm confused! Any clues? Alex Tyrer
HeadGames Inc. atyrer@capsinc.com [broke(?) due to InterNic DNS debacle]
alby@nanospace.com [should always work]

Re: Streaming Audio & Video
Written by Mike Fulton on Tue Apr 15 02:18:15 1997

To Alex Tyrer:

If you have an STR file on your PC that contains audio, then to make a PlayStation disc that will play correctly, you need to tell CDGEN that this is a FORM 2 file.

Basically, when you added in audio, MOVCONV or MOVPACK creates a file with 2336 bytes per sector. CDGEN only has to add 16 bytes per sector and then fill in the appropriate info.

With regular data files, there are 2048 bytes/sector and CDGEN has to add in 304 bytes and then fill in the appropriate info.

You don't normally need to run MOVPACK unless you have more than one stream of audio to interleave or if you have your own data to interleave.

When you have an XA audio track interleaved with movie data, it gets played automatically, because sectors containing XA audio are automatically sent by the CD controller to the SPU for decoding and playback. You don't really have to do anything (although you can screw up and do the wrong thing, like specify an XA playback channel that doesn't exist in the movie you're playing.)

To create a movie with audio, when you convert the movie in the first place, tell MOVCONV that the movie has audio and specify the desired sample rate, etc. Do this even if the AVI file you're converting doesn't have any audio. This will cause MOVCONV to either convert the existing audio, if any exists, or else reserve space for audio you'll be adding in later.

If your AVI file had audio, you're done. Burn a CD and have fun.

if you still need to add in the audio, then use MOVCONV to create an XA audio file from your WAV source file. If you don't have a WAV source, you'll want to convert whatever you have to this format before proceeding.

MOVCONV does not perform stereo-mono conversions, so make sure you convert to the same thing as your source file. It will convert sample rate. It will not convert 8-bit files. For the best results, pass it 16-bit 44.1k, either stereo or mono.

Once you have an XA audio file, use the "Interleave audio with video" option in MOVCONV to merge the two together.

The resulting file can be burned onto a CD.

SPU streaming is a whole other topic and doesn't have anything to do with playing back XA audio in a movie. It is also a subject about which you should ask questions in the "Sound Programmers" area.

Mike

CDGet sector
Written by bd on Wed Apr 09 21:27:16 1997

What happens if I get a DataReady callback and I don't transfer all the data within CdGetSector?
In other words, if I've called CDGetSector(ptr, 200) and then I wait for the next DataReady callback.
Will weird things happen?

Thanks!

Ben

Re: CDGet sector
Written by Mike Fulton on Wed Apr 09 21:59:12 1997

To bd:

I've always recommended that you suck out all of the data from the CD controller buffer. If there's some data you don't need, read it into a dummy buffer.

However, I'm not aware of any definite problems with only reading a partial sector. Go ahead and try it and see what happens. Just remember you have no random access into the sector buffer. You start at the beginning and read as much as you read.

Mike

Skipped Sectors - Ouch!
Written by Tom Rolfs on Fri Apr 04 02:14:03 1997

FYI:

Anyone using the CD low level interface (i.e. CdControl) should read the section headed "Special CD-ROM Notes" in the 3.6 and 3.7 Overview docs. Of particular note is the section on "skipped sectors". If you fail to follow the advice given there you are going to experience infrequent random bugs in your title especially if you do a lot of streaming (this is coming from someone who has already been burnt by the skipped sectors scenario).

Re: Skipped Sectors - Ouch!
Written by Mike Fulton on Fri Apr 04 02:30:22 1997

To Tom Rolfs:

Keep in mind that the skipped sector problem will only occur if you are spending so much time in interrupt code (i.e. callback routines) that the CD sector-ready callback and the resulting sector transfer doesn't get to happen in a timely fashion.

Mike

Re: Re: Skipped Sectors - Ouch!
Written by Tom Rolfs on Fri Apr 04 19:42:03 1997

To Mike Fulton:

Absolutely true. However, one of our titles in test was experiencing streaming hiccups about once every 144 hours on the blue playstations. We have streaming audio and our own data dmas to and from sound ram, CD interrupts and async rendering going on so I guess it is hard to absolutely guarantee that we are not going to miss an interrupt during the course of the game. Doing the header check makes everything that much more robust. Personally I would make it mandatory to do the check; it can't hurt

CdControl: blocking or not?

Written by Eric Yiskis on Wed Mar 26 22:45:52 1997

In the Run-Time Library Reference (v3.0 P.81) under CdControl, it says "This function is a non-blocking function, so it is necessary to use CdSync to detect actual transfer termination". However, in the table below, in the "Type" field, some are labeled "B" - blocking and some are labeled "N" - non-blocking. I assume the ones labeled blocking actually are blocking? I only ask because under emulation, one of my CdControl(CdSetloc,...) calls seems to take a very long time (3 frames @ 30 frames per second). I have tried the CdControlF, and it didn't seem to make a difference. Is there a "non-blocking" version of this call? Also, this delay either doesn't exist or is much smaller on an actual CD, or perhaps it doesn't always show up?

Re: CdControl: blocking or not?
Written by Mike Fulton on Thu Mar 27 00:31:12 1997

To Eric Yiskis:

CdControl is not blocking per se, but if the previous CD operation you did has not finished yet, then CdControl will wait until that operation is finished before starting a new one. You cannot queue up a list of CD commands automatically like that.

So it may not be the CdSetloc, but rather it may be the previous operation.

The CdControlF function essentially skips the ACK/NAK handshaking with the CD controller. I don't recommend that you use it unless it's for the CdINOP command or unless you're absolutely sure that the previous command isn't still in progress.

Mike

MDEC

Written by Jeff Beckham on Wed Mar 12 19:10:11 1997

I'm having a hard time getting the MDEC to work. Specifically I'm trying to load and decode a 128x128 BS file created with MovConv v2.1e. I'm following the support examples pretty much line by line for the decoding, but the MetroWerks compiler spits out the following during runtime.

MDEC_in_sync timeout:

DMA=(1,0),DDR=(0x000e5258->0x0008d004)

FIFO=(1,1),BUSY=1,DREQ=(0,1),RGB24=1,STP=0

MDEC_out_sync timeout:

DMA=(0,1), ADDR=(0x000e5258->0x0008d004)

FIFO=(0,0),BUSY=0,DREQ=(0,0),RGB24=0,STP=0

I'm thinking that the problem lies in the MovConv process, that I'm not setting up the conversion correctly. Any ideas?

Jeff

Re: MDEC
Written by Mike Fulton on Thu Mar 13 03:07:31 1997

To Jeff Beckham:

This is a bit off topic for the LIBGPU/GTE section. Please move it to the LIBCD/LIBSTR section. It's not very likely that the source data is not compressed properly, but it may be getting corrupted. We need to see your decode routine so we can make a decent guess at the problem.

Mike

DECDCTTAB & DecDCTvlcBuild ??

Written by Kenton Leach on Wed Mar 12 02:35:31 1997

I downloaded the latest Movie samples from the files area (movie.zip) and there are some new structures and functions used that I can't seem to find in the most recent libraries or documentation. There is a function DecDCTvlcBuild() used that takes a DECDCTTAB as a parameter. Am I just looking in the wrong place or are these not yet released, thus making the tutoX.c samples in the movie.zip file un-compilable? If these can not currently be compiled, which samples should I look at for playing back STR files?

- Kenton -

Re: DECDCTTAB & DecDCTvlcBuild ??

Written by frank on Wed Mar 12 17:11:44 1997

To Kenton Leach:

These are mentioned in the text file you see when you click on the date stamp for lib37.zip, which is in the same area where you found movie.zip.

Frank

MDEC Bus usage - From SCEA BBS

Written by Mike Fulton on Tue Jan 07 01:43:52 1997

Forwarded from CD/Decompression message base on the SCEA BBS.

Ben Diamand writes:

I think I may be having some bus arbitration/timing conflicts with dctin/dctout and have some questions regarding said timing:

Is dctin finished when the first dctout is successful? Can I assume that the bus is NOT in use at the time I get a dctout callback? I'm asking because I seem to be missing cd interrupts while doing the dctout. I am emulating the cd samples in that I am giving my own cd code time when I'm in the dctout callback right before I get ready to decode my next slice, but I'm still missing data ready callbacks. Well, I'm not missing them, but I defer them until the dctout callback is called and at about slice number 5 (I'm in 24 bit mode, and I'm using 16x240 slices, just like the samples), I get a dataready callback and the code is still waiting for a previous slice of time from the dctout callback! This is fatal...

Could dctout be blocking from within the dctout callback? What would cause this? Are there any diagnostic routines that I can query the MDEC with, like, "How much data do you think you have?" and "Do you think you're still in a DMA? Which direction? For how many Bytes?", etc?!

oh, I'm using the 3.5 version of the libs... Thanks, sorry if I wasn't coherent, it's Monday!

Ben

Re: MDEC Bus usage - From SCEA BBS
Written by Mike Fulton on Tue Jan 07 01:44:58 1997

To Ben Diamond:

There are virtually no circumstances where you're going to see any real problems as a result of bus arbitration. The PlayStation has excellent bus arbitration, and by attempting to second guess it, you're probably causing much worse problems than the theoretical ones you were trying to avoid in the first place.

When your DecDCTOutCallback routine is called, you cannot assume that the bus is not in use. The only thing you can assume is that the last DecDCTOut function is finished. That's only one of many possible bus operations, after all.

The DecDCTin call initiates a DMA transfer from main memory to the MDEC. There is a dedicated DMA channel used for all Memory <> MDEC transfers, so a DecDCTout operation cannot begin until the DecDCTin operation finishes. That doesn't mean the bus itself is completely tied up. The

DMA transfer does not steal the bus entirely, but rather is interleaved with free cycles so that other operations not relating to Memory<>MDEC transfers can be going on concurrently. This is called "Cycle Steal" mode DMA.

DO NOT UNDER ANY CIRCUMSTANCES DEFER YOUR CD CALLBACKS! This is almost certainly why you are seeing problems. Process the callback for each sector when it occurs, regardless of what's going on in the rest of the system. If you do otherwise, you'll have all sorts of problems. Don't try to second guess the system's bus arbitration.

The main mistake you are probably making is worrying about what the bus is doing in the first place. The only concession I would recommend making to managing bus arbitration is that you may consider breaking up large DMA transfers into several smaller ones so that the bus is free for longer periods of time in between. For example, you mentioned that you were doing 16x240 slices for DecDCTOut, and that's a hell of a lot of data to transfer in one step. (For a 16-bit RGB output, that's nearly 8k at once!) So break each vertical strip down into 2 blocks of 16x120, 2 blocks of 16x80, or even all the way down to 15 blocks of 16x16. The only extra time will be the setup time for the DecDCTOut function, and that will be pretty trivial.

Mike

Forum: OS/ Architecture

Trashing Memory?

Written by Stephen Magladry on Sat May 31 01:20:40 1997

I am wonder if the contents of 0x80000000, the start of the OS is a constant value? If so, what?

I would like to use that memory location as a check to see if I am trashing memory.

Thanks before hand.

Stephen

Re: Trashing Memory?

Written by Mike Fulton on Mon Jun 02 19:26:34 1997

To Stephen Magladry:

Memory locations from 0x80000000 to 0x8000FFFF are RAM locations which are used by the system and library kernel. (It is not, however, "the start of the OS" as you indicated.)

The usage of these memory locations is completely undefined.

You should not refer to them. Do not use them for your own storage, and do not rely on them to either remain constant or to change.

I'm afraid you'll have to find another method to determine if you are trashing memory. Do you currently have a problem with this, or are you trying to prevent a potential problem?

For debugging purposes, an ASSERT call is probably one of your best choices to track down a bad pointer.

Mike @ SCEA

InitHeap

Written by bd on Tue May 20 19:21:36 1997

If my program does *not* call InitHeap from the entry point, will anything else besides the malloc and related functions be affected? Do any other PSX library functions rely on this function being called? Thanks!

Ben

Critical sections

Written by bd on Fri May 09 03:38:10 1997

What happens if I enter/exit a critical section more than once? i.e.

```
void func()
{
    int inInt = CheckCallback();
    if(inInt == 0)
    {
        EnterCriticalSection();
    }
    ...
    func(); //yes, this is recursive...
    ...
    if(inInt == 0)
    {
```

```
ExitCriticalSection()
```

```
    }  
    }  
    ...
```

Will the multiple calls to Enter/ExitCriticalSection hurt things?

Thanks!

Ben

Re: Re: Critical sections
Written by bd on Sat May 10 00:12:51 1997

To Mike Fulton:

Thanks! This is exactly what I needed to know.

To answer your question, I have code that operates within a VBL that messes with many variables and I have a very short section of non-interrupt code that also messes with some subset of these variables. I need to make sure that they're mutually exclusive. I had code that worked fine until I had multiple Enter/Exit calls, then I had weird timing related bugs. I'll be more careful. Thanks again.

Ben

Re: Re: Re: Critical sections
Written by Mike Fulton on Tue May 13 22:46:10 1997

To bd:

The easiest thing to do would be simply to have a semaphore that the non-interrupt code sets when it is going to mess with the variables in question. Then if the interrupt sees that the semaphore is raised, it will leave those variables alone. The non-interrupt code would clear the semaphore when it was done, allowing the next interrupt to do what it needed.

Mike

Exec and callbacks

Written by bd on Wed May 07 23:38:39 1997

Looking at some examples of Exec, I see that ResetGraph(3) is called, StopCallback is called, and EnterCriticalSection is called.

What I want to know is, why are these three function calls necessary. I ask because I am loading a piece of code as a child, but I'd like portions of the parent to continue to execute while the Exec call is running and while some of the child's startup code is running.

In particular, I want to be able to play some sounds, (via SPU calls), muck with cluts(LoadImage), and run a custom midi player(I need a root counter/vbl for this). I wanted to use a VBL to give all these pieces of code time.

Thanks.

Ben

Re: Exec and callbacks
Written by Mike Fulton on Thu May 08 00:52:19 1997

To bd:

As long as your parent's memory is not being overwritten, and you don't turn them off before loading a child, there's no reason that callback routines would not continue to be executed.

It's possible you could have some stack issues if you're not careful. You may find it easier if the child uses the same stack as the parent.

The ResetGraph function is resetting the graphics library and GPU state. Specifying "3" as the parameter keeps it from clearing the screen and altering the display environment.

As far as the other stuff goes, it may not be applicable to your particular setup, but it depends a lot on what you're going to have the child doing.

Also keep in mind that you don't necessarily have to use "exec" in order to do what amounts to code overlays.

Mike

Re: Re: Exec and callbacks
Written by bd on Thu May 08 03:12:09 1997

To Mike Fulton:

Regarding overlays; this just seemed like an easy way not to have to deal with all the bs in setting them up. But I need to make sure that none of the sony functions will get confused if I init them under one stack(the parent) but make calls through the child, since obviously offsets from the SP will be incorrect. Can you verify this? Or do I need to use a 'global' stack that both share. I'd rather not.

Thanks!

Ben

Re: Re: Re: Exec and callbacks
Written by Mike Fulton on Thu May 08 19:25:51 1997

To bd:

Unless you either a) use standard overlays; or b) use a function pointer table to access all library functions, your child process will have to carry around it's own set of library routines.

Frankly, I'm not sure I really understand your plan as far as concerns what will be in the child versus what will be in the parent, what the library requirements are, etc. Can you start at the beginning and explain what you're doing?

Mike

Are callbacks mutex?

Written by David Minogue on Mon Apr 28 21:04:55 1997

I am a little confused about the nature of callback functions. Can callback functions interrupt one another? Is it necessary to use SW critical sections to guard code within callbacks, or is there another method?

In my case, I am using a CdSyncCallback to avoid blocking withing my CdReadyCallback when I want to issue a CdControl() command. If the CD controller is busy on a previous non-blocking command (ie. CdIPause), then the new command is enqueued for execution upon the next CdSyncCallback. I want to be sure that neither callback interrupts the other, or the queued command may be lost. I tried using SW critical sections, but this causes my code to hang for an unknown reason. It even hangs if the critical section in the CdReadyCallback is completely empty. Putting a critical section in just CdSyncCallback doesn't seem to cause this problem.

If callbacks are always executed serially, then I have no problem; but if they can interrupt one another, then I need to figure out a way to effectively guard the critical sections.

Any clarifications of this subject will be appreciated.

- David Minogue.

Re: Are callbacks mutex?

Written by Mike Fulton on Wed Apr 30 01:33:14 1997

To David Minogue:

What does "mutex" mean?

Generally speaking, most callback routines are called only from an interrupt context, and cannot be interrupted.

Keep in mind that callback routines are subroutines called from the interrupt handler, not the interrupt handler itself.

I don't think this information really applies to what you want to accomplish, however. The problem is you've got is that CD commands are not queued the way you think.

When you issue a "non-blocking" command, what is meant by "non-blocking" is that the command is issued to the CD controller, but control returns to you before that operation completes.

However, if you issue another command before the first one completes, it doesn't get queued. Rather, it blocks until the first command is finished. Then the next command is issued. It doesn't matter if

the call is “blocking” or “non-blocking”. If there’s a previous operation still going on, it will block no matter what.

What exactly are you trying to do from within your CdSyncCallback routine?

Mike

Re: Re: Are callbacks mutex?
Written by David Minogue on Wed Apr 30 15:10:08 1997

To Mike Fulton:

Thanks for your response.

By mutex I meant Mutually Exclusive, which is just another way of saying critical sections (while in one, you can’t enter the other). I didn’t mean to be cryptic; it just happens sometimes.

I am trying to use the CdSyncCallback routine to avoid blocking when issuing a command to a busy CD controller, thereby keeping all of my callbacks as short as possible. If the CD controller is busy processing a non-blocking command (specifically CdIPause) when I want to issue a new command (specifically CdIReadN), then I set a flag to cause the new command to instead be issued within the CdSyncCallback which will be triggered by the completion of the pause command. Since I am keeping track of when the CD controller is busy, I should even be able to use CdControlF(), but I haven’t taken that dangerous step yet.

My concern was that there could be a ‘window of vulnerability’ (a point at which being interrupted could cause an error) after the CdReadyCallback decided to queue a CD command for execution in the CdSyncCallback and before the command had actually been queued. If the CdSyncCallback could be entered at this point, then a command queued for the CdSyncCallback could be lost. If interrupts are blocked, then this shouldn’t be a problem.

Is the CdSyncCallback called for all CD commands, or just the non-blocking ones? What if CdControlF() is used? How much time does it save?

Since CdControl() waits for the completion of the previously issued command, under what circumstance is it appropriate to retry the call? I noticed that several of the sample programs contain the code:

```
while (CdControl(...) == 0) ;
```

Since CdControl(CdIReadN,...) is described as a blocking command, does this mean that it will wait until the associated seek is complete (ie. on the order of 1 second)?

What are the exceptions which made you qualify your first statement (Generally speaking...).

If callbacks are called from the interrupt context, then what distinguishes them from event handlers? Or are they functionally equivalent? Despite several attempts by Sony to clarify this, I remain confused. Sometimes I think I get it, then a couple of days later I’m confused again. Is there something about which stack is active? Or the relative priorities? Are interrupts really blocked in both cases? If neither callbacks or event handlers can be interrupted, then what is the purpose/effect of the Sw[Enter/Exit]CriticalSection calls? Are there different limitations to what can/cannot be done within Callbacks and Event handlers, or is it all trial and error?

I guess that’s more than enough questions for one message. I hope they made sense.

- David.

Re: Are callbacks mutex?
Written by Mike Fulton on Wed Apr 30 20:35:36 1997

To David Minogue:

The main difference between an event handler and a callback is that a callback routine gets called in direct response to some operation which is initiated by your code, while an event is either something timer-based or something which may occur randomly (like a user inserting a new memory card).

Also, some "callback" routines aren't directly tied to a specific interrupt. Such as DrawSyncCallback, which is called only when the LIBGPU task queue is emptied, not for each GPI interrupt which may occur.

Aside from that, it's not so much a technical thing about dispatching. (As far as I know... I've looked at the library code for callbacks pretty extensively, but not so much at the event stuff.)

The way interrupts work on the PlayStation is driven by the MIPS R3000 architecture. There is no interrupt priority like in an Intel 80x86 or Motorola 680x0. But I'm no expert and I don't fully understand how a lot of this stuff works, especially for things like non-maskable interrupts (NMI's). I suggest you get this book:

MIPS RISC ARCHITECTURE
by Gerry Kane & Joe Heinrich
Prentice Hall

If you can't find it at a computer book store, call MIPS customer service at (800) 443-MIPS.

Here's a rundown of the various CD-related callback functions and what they're for:

CdReadyCallback : This callback is executed when a new sector has been read into the CD controller. This is normally called for data and XA sectors only, but can also be CD-DA if you've turned report mode on. The usual response is to transfer the sector data with CdGetSector, but note that the CD subsystem automatically transfers XA audio data to the audio subsystem without CdGetSector.

CdReadCallback : This is called when the last sector of your CdRead operation is finished. This is actually called from within the CdReadyCallback routine that is installed by CdRead.

CdDataCallback : This is called when a DMA transfer from the CD controller to RAM (i.e. CdGetSector) is completed. It doesn't really need to be a public function, since CdGetSector is actually blocking (contrary to older documentation). Mainly this callback gets used internally by the streaming libraries, which don't use CdGetSector.

CdSyncCallback : This is called when any operation initiated by CdControl reaches completion, except for "play" or "read" which use the callbacks described above. Note that some operations do little more than set registers or return data maintained by the library.

While it should be safe if you know the CD controller isn't busy, I really cannot recommend using CdControlF. I've seen too many mysterious problems suddenly go away when someone switched from CdControlF to CdControl or even CdControlB.

Basically, what CdControlF does is to skip past the steps that read back the CD controller registers to make sure the command was accepted. Perhaps even more importantly, it also skips the code that saves the status of the CD controller in the result array you've specified.

So if you did:

```
CdControlF( command, inputparams, result );
```

With CdControl, the CD controller status would be placed in result[0]. But with CdControlF, this step is skipped, so the value in result[0] is left unchanged (ie. it's not affected at all by CdControlF, not even zeroed).

If you wanted to use CdControlF with CdINOP to get the previous CD controller status, that would be OK, but even then you'd be better off in most cases using CdControl to get the most-current status.

Mike

Faster Sqrt / Atan

Written by Bruce on Wed Apr 16 16:39:43 1997

I have a function that requires both 'sqrt' and 'atan' calculations for determining angles and distance between points. Initially I tried using the Math Library and found both of these supported functions to be tremendously slow. I've already made a lookup table for the 'atan' calculation which works great. The question I have is whether it is necessary to table this and to come up with a 'sqrt' cheat? Or does the PSX have any hardware accelerated support for these types of math calculations that I am overlooking? And if it is necessary, does anybody have a recommendation for a 'sqrt' cheat that quickly approximates a sqrt (where I need the sqrt is when calculating distances between two points in 3d space --- $\sqrt{x^2 + y^2 + z^2}$ --- this calculation doesn't need to be exact, but just a good approximation...

Thanx All,

Bruce

Re: Faster Sqrt / Atan
Written by Dave Scheele on Wed Apr 23 18:26:48 1997

To Bruce:

Sometimes I can set up things so that I don't need to take the square root to compare distances. An extremely simple example of this is to say

```
if (DistanceSquared(pt1, pt2) > 144) { ... }
```

instead of

```
if (DistanceBetween(pt1, pt2) > 12) { ... }
```

Without too much effort I was able to do away with the majority of my squareroot calls in my last game by doing this.

Dave

Re: Faster Sqrt / Atan
Written by Mike Fulton on Thu Apr 17 17:58:56 1997

To Bruce:

I think Christer's answer has the answer you need. Check it out.

Mike @ SCEA

Re: Faster Sqrt / Atan
Written by Christer Ericson on Thu Apr 17 01:44:20 1997

Bruce,

There's no hardware support for calculating neither sqrt nor atan (nor any other trigonometric functions for that matter), so you're best off rolling your own code.

For the general sqrt, things pretty much depend on what accuracy you need, how much memory you can spare, and how fast you want it to be.

A good "middle-ground" sqrt approach is to write your value v as $v = 2^{2n}k$, $n \geq 0$, $4 > k \geq 0$, and thus $\text{sqrt}(v) = 2^n \text{sqrt}(k)$.

Using the GTE LZC (leading zero count) registers to find the highest bit set in your value, n is easily (and quickly) found. v is scaled appropriately to give k , for which the square root is looked up in a table (remember that k has a limited range so this can be quite a small table if your accuracy demands are low) and then we shift the decimal point over by n places to get the final result.

As for the 3D Euclidean distance, a popular approximation is:

$$\text{dist} = \max(dx,dy,dz) + \text{mid}(dx,dy,dz) / 4 + \text{min}(dx,dy,dz) / 4$$

(other constants give lower error, I could look them up if you're interested).

Christer,
Neversoft

Word Boundries
Written by Jeff Beckham on Wed Mar 05 05:36:42 1997

I've discovered this the hard way, but would like confirmation from someone. Can memory reads cross WORD boundries? Must WORD reads be WORD aligned?

Code sample that works:

```
u_long *Data;
long Offset;
long Value;
```

```
Data = SomeAddress; // (malloc'd, so probably WORD aligned) Offset = 0; Data += Offset;
```

```
Value = *Data; // (WORD read, 4 bytes)
```

Code Sample that does NOT work:

```
u_long *Data;
long Offset;
long Value;
```

```
Data = SomeAddress; // (same as above)
```

```
Offset = 3553;
```

```
Data += Offset;
```

```
Value = *Data; (again, 4 byte read, start from ODD address)
```

This second code snippet stops with a memory load error, I suppose because of the odd address, and then trying to read across a WORD boundary. This took me so long to track down, I don't want to even mention it, but I would like confirmation that adding an ODD offset, and then reading 4 bytes, will cause a bus error, and that I can stop looking through pages of code for some unknown bug. By the way, this offset comes from seeking into a data file and reading the LONG value there. I guess the way to fix this is to pad my files to start on even address? Or do they have to be divisible by 4?

Thanks for your time,
Jeff

Re: Word Boundries
Written by Mike Fulton on Mon Mar 10 04:55:36 1997

To Jeff Beckham:

I'm looking through my R3000 book and I haven't quite interpreted everything yet, but I think it's saying that it's only for certain instructions that it requires 16-bit alignment for 16-bit read/writes and 32-bit alignment for 32-bit read/writes. However, I haven't quite read enough of the book to be sure yet.

However, for the code sample you included, the value would NOT be read from an odd address. Here's what you said wouldn't work:

```
u_long *Data;
long Offset;
long Value;
```

```
Data = SomeAddress; // (malloc'd, so probably WORD aligned)
```

```
Offset = 3553;
```

```
Data += Offset;
```

Value = *Data; (again, 4 byte read, start from ODD address)

The problem is you're not taking into account the way C pointer arithmetic works. When you add a value to a pointer, it adds that value times the size of what the pointer is pointing to.

Since "Data" is a pointer to a type that takes up 4 bytes, doing

"Data + Offset"

will actually add (3553*4) to the address already stored in "Data". Keep in mind that doing

"Data + Offset"

is the same thing as doing

Data[Offset]

I suspect the REAL reason you're crashing is because you were not taking this aspect of pointer arithmetic into account, and as a result you are trying to access some location past the end of memory.

Mike

Re: Word Boundries
Written by Christer Ericson on Wed Mar 05 22:54:29 1997

Jeff,

You can only read/write longwords (4 byte entities) off addresses that are 4-aligned (multiples of four). You can only read/write words (2 byte entities) off addresses that are 2-aligned. Trying to do otherwise will result in an address error.

Obviously the best way of avoiding this is to align your data in the first place. However, if you're porting lots of code from a platform that does unaligned reads (eg PC), or if you don't have the extra space that aligning takes, then you'd want to use the lwl/lwr instructions to read your longwords as it'll be much faster than doing 4 byte-reads.

We've been using the following inline assembly routine for reading unaligned data off our data files:

```
inline int ReadUnalignedInt(register char *src)
```

```
{
register int res;
    __asm__(
        "lwl %0,3(%1); \n"
        "nop; \n"
        "lwr %0,0(%1)" \n
: "&r"(res) : "r"(src) : "%0");
return res;
}
```

If anyone knows a better way of doing inline assembly that allows instruction reordering I'd be interested to hear about it; as it is, the nop really begs to be replaced with something that can be done while the first load is executed.

Christer,
Neversoft

Re: Re: Word Boundries
Written by David Bollo on Thu Mar 06 17:14:20 1997

To Christer Ericson:

Is the nop between lwl and lwr actually necessary here? In my "MIPS RISC Architecture" book in the documentation for lwr, it says "The contents of general register rt are internally bypassed within the processor so that no NOP is needed between an immediately preceding load instruction which specifies register rt and a following LWR (or LWL) instruction which also specifies register rt."

Hope this helps...

David
EA Canada

VBlank
Written by Stephane Leblanc on Thu Feb 06 04:05:36 1997

I'm experiencing something quite curious. At a specific place in the game, when I'm placing the main character and the camera around the same coordinate, for an unknown reason, my code jump from 525 scan lines to 790 scan lines. I isolated the routine that was hanging and found out that it wasn't taking more time but that the VBlank interrupt was occurring in it. This routine performs the polygons calculations, sprites, etc... When the character is in another coordinate in the world, not far away from the other, VBlank occurs as well in this routine but doesn't affect much my calculation time. There's nothing special about this area of the level. I also find out that the difference between 790 and 525 is 265, a NTSC video scan. So my question is, could the interrupt server be hanging for a scan before returning control to my code?

Re: VBlank
Written by Rob Vawter on Thu Feb 06 22:04:46 1997

To Stephane Leblanc:

Perhaps you are just spending too much time in the previous Vblank interrupt doing polygon calculations. Is this area of the level more polygon intensive. It's unlikely that the interrupt server could be hanging for a scan.

What other interrupts are you using?

Re: Re: VBlank
Written by Stephane Leblanc on Wed Feb 19 15:30:36 1997

To Rob Vawter:

I found the problem. The drawing is taking more than 2 scans to complete. Would you have any drawing speeding technics to share ? About 90% of our textures is 4 bits 32 x 64 and our 8 bits are 32 x 32. We sort them by textures types and then by polygones types. The major drawing slow down is that I have a 512 x 240 8 bits sprite as background. I've tried using a 16 bits image and transfered it using DR_MOVE and only got a 10% gain.

Re: Re: Re: VBlank
Written by Mike Fulton on Fri Feb 21 00:24:56 1997

To Stephane Leblanc:

When you use DR_MOVE, the time it takes to transfer the image is added to your rendering time. Instead, as long as you're using a 16-bit image for your background, try using MoveImage immediately after you do your buffer switch.

In other words, if buffer 1 is displayed while buffer 2 is being drawn, when you switch to display buffer 2 and go to draw buffer 1 again, do your MoveImage call immediately after the switch.

The reason is that the move will take place in the background while you are doing your 3d calculations, building your primitive list, and doing whatever else you need to build a screen up to when you do DrawOTag. As a result, your code won't be delayed by the time needed to process the move.

Mike Fulton

Re: Re: VBlank
Written by Stephane Leblanc on Fri Feb 07 14:04:33 1997

To Rob Vawter:

Actually, there's other place in the game that are more polygon intensive and doesn't slow down. I'm using VSyncCallback and DrawSyncCallback. I'm also using some CD interrupts but there were turned off since I was using LoadPC to get my data. Anyway, I think I know where the problem lies, I think that the GPU take more that 2 VBlank to draw the list and that my routine is waiting for some buffer to get free cause I'm using a ring buffer for my polygones packet (saves a lot of memory). I'll let you know how thing turn out.

intr timeouts

Written by Erik Strickland on Wed Jan 29 01:47:45 1997

Hey, Guys.

This is a repeat posting of a question submitted to the bbs.

I have a launcher which successfully Loads and Execs Twisted Metal 2 and Gameday '97. If I load TM2 after Gameday, then during the TM2 startup, an endless stream of interrupt timeout messages get displayed.

They read:

```
intr timeout(0040:0049)
intr timeout(0040:0049). . .
```

Which interrupt is timing out?

How can my launcher (or gameday) prevent these timeouts. (Note: launching gameday also causes the same timeout in Crash Bandicoot and Blast Chamber).

Thanks for your immediate attention. (My `_second_` deadline to get out of test is Friday).

- Erik.

Re: intr timeouts
Written by Mike Fulton on Wed Feb 12 02:23:08 1997

To Erik Strickland:

(I think we already replied on the BBS, but I want people reading here to get the response too.)

The "intr timeout" message doesn't mean that the interrupt is timing out, exactly. What it means is that an interrupt request occurred while another interrupt was being serviced.

When an interrupt occurs, the CPU throws an exception as soon as the bus is available and goes to the system interrupt routine. This is a dispatcher function which steps through all the bits in the interrupt request register, and if it sees that bit is set, it knows that type of interrupt is pending. So it calls the appropriate routine, which would include any callback routines you may have installed. That routine is responsible for clearing the bit in the interrupt register to indicate that the interrupt has been serviced.

Once it's gone through all the bits in the interrupt register, the dispatcher routine checks the register once again. If it's non-zero, then it indicates that another interrupt occurred. It's not really expecting this, so it prints the "intr timeout" message.

The system interrupt routine will then return, but because the interrupt register is non-zero, an exception will occur again immediately and the system interrupt routine will start stepping through all the interrupt handlers again.

The danger is that the interrupt register is just a YES/NO flag for each type of interrupt. There is no way to determine if there were more than one interrupt of any particular type.

The "0040:0049" values indicate that there is one interrupt pending and three interrupts enabled.

The bits in the first value (0x0040) indicate pending interrupts. In this case, bit 6 is set, indicating that TIMER 2 has requested an interrupt.

The bits in the second value indicate which interrupts are enabled. In this case, it's TIMER 2, the DMA controller, and the vertical blank.

The reason there is a potential (but not definite) problem here is that we can't tell if more than one TIMER 2 interrupt request has been made.

The cause of this problem would be that other interrupt routines are taking longer to process than the amount of time between two (or more) TIMER 2 interrupts.

This could be because the routines are simply too long, or it could be because there are other things like DMA operations tying up the bus.

Mike

overlay groups

Written by Michael Koziniak on Tue Jan 07 01:10:26 1997

>Jeff Blazier Writes

>

>Michael?

>hey,

>is it possible to create two or more overlay groups? >something like (pardon the pseudo code, I don't have a lnk >file handy):

>

- group balls
- group test1 (over "balls")
- group other
- group test 2 (over "other")

>

>and if so, is there anything different I need to do for the >makefile or address.s or lnk file?

>

>thanks,

>Jeff Blazier

It is possible to have more than one overlay group. Things you will need to do:

Specify another LoadAddress in your address.s file(i.e. LoadAddress 2).

Specify another group in your .lnk file and the groups which you want to load over that.

Note: If you want to have the same overlay code loaded in 2 different overlay areas in memory then you will need to have 2 separate bin files, one for each group area.

Michael Koziniak
Development Support Engineer
Sony Computer Entertainment

How to use LoadExec()

Written by Michael Koziniak on Tue Jan 07 00:53:24 1997

From: Mike Fulton
Note: Moved from first class

Phil Weeks writes:

We experienced a multitude of problems using LoadExec while releasing our current product in the U.S. We are now submitting for Europe and Japan and are again experiencing the same problems that we had before. We have a frontend exe that calls game/demo exes, and will fail in the library routines. I'm interested if anyone else has been using LoadExec and having problems similar and came up with a 100% solution, or if someone knows of when the libraries quit working consistently with LoadExec. Phil

We're not aware of any problems with LoadExec, but you do have to be careful about several things when using it. This has always been true, but it's possible you may notice it more or less with different versions of the libraries simply because of things moving around in memory. Please realize that just because it crashes within the library doesn't necessarily mean the library is at fault. Other things could be stepping on the libraries.

- 1) Do you use LoadExec to replace the current program with a new one? This is bad because you have no chance of error recovery. It's better to use "Load" and "Exec" separately. Also, try to load from a "safe" area of memory so that if the load fails for any reason, you have code to return to.
- 2) Are you positive that all callback routines are shutdown before calling LoadExec? You need to shut down certain CD callbacks explicitly by calling the appropriate function with a NULL pointer (CdReadyCallback, CdReadCallback, CdSyncCallback).
- 3) If you've done anything with the MDEC, reset it using DecDCTReset(0) before calling LoadExec.
- 4) After loading new code into memory, make sure you flush the cache using the FlushCache call.
- 5) Keep in mind that the initial stack location that you set is only used up to the point when control is passed to the new program. The new program can and probably will set a new stack pointer. (And what startup code are you using?)

Mike

Forum: Graphic Art Tools

Kinetix 3D Studio R4 Web Links
Written by Jeff Yates on Fri May 30 21:29:30 1997

Hi:

For those of you who may be using Kinetix software with Net Yaroze, here are a number of web resources for 3D Studio R4, Animator Pro, Texture Universe, and more. Hope they are helpful.

Regards,

Jeff

WEB SITES

<http://ktx.com/> - Kinetix Web Site Home Page

<http://ktx.com/3dcareers/> - 3D Careers Web Site Home Page

PRODUCT INFO

<http://ktx.com/3dsr4/> - 3D Studio Release 4

http://ktx.com/animator_pro/ - Animator Pro

http://ktx.com/texture_universe/ - Texture Universe

DOWNLOADS

<http://ktx.com/javascript/downloads/3dsr4/fastcpu.htm> - Fast CPU Fix

<ftp://ftp.ktx.com/download/patches/3dsr4/vib3ds/vib3ds.zip> - Latest Vibrant In-line Graphics

Drivers

http://ktx.com/javascript/downloads/3dsr4/r4_vrml.htm - VRML 2.0 Exporter Plug-in

PRODUCT SUPPORT

<http://www.autodesk.com/support/techdocs/mm/td801206.htm> - 8-bit Character Mapping in 3D Studio R4

<http://www.autodesk.com/support/techdocs/mm/td801340.htm> - 64MB Memory Limit and 3D Studio DOS

<http://www.autodesk.com/support/techdocs/mm/td801196.htm> - Technical Specification: 3D Studio R4

<http://www.autodesk.com/support/techdocs/mm/fax890.htm> - Network Rendering with 3D Studio

<http://www.autodesk.com/support/techdocs/mm/fax892.htm> - Creating Simple 3D Studio GeomeTrees

<http://www.autodesk.com/support/techdocs/mm/fax954.htm> - Troubleshooting Reflection Maps

<http://www.autodesk.com/support/techdocs/mm/fax756.htm> - 3D Studio Texture Coordinates

<http://www.autodesk.com/support/techdocs/mm/fax822.htm> - 3D Studio: Printing to the HP PaintJet

<http://www.autodesk.com/support/techdocs/mm/fax747.htm> - Updating IPAS 3 to 4 and the file axp.c

<http://www.autodesk.com/support/techdocs/mm/fax746.htm> - How To Calculate KXP Spline Derivatives

<http://www.autodesk.com/support/techdocs/mm/fax743.htm> - Authorizing 3D Studio R4 c2

<http://www.autodesk.com/support/techdocs/mm/fax685.htm> - Exporting AutoCAD R13 objects into 3D Studio R4

<http://www.autodesk.com/support/techdocs/mm/fax823.htm> - Clean Booting 3D Studio

<http://www.autodesk.com/support/techdocs/mm/fax709.htm> - 3D Studio R4 Running on MS Windows NT or MS Windows 95

<http://www.autodesk.com/support/techdocs/anim/td801259.htm> - Bus Logic SCSI drives and Phar Lap Conflicts

<http://www.autodesk.com/support/techdocs/anim/fax824.htm> - Clean Booting Animator Pro

<http://www.autodesk.com/support/techdocs/anim/fax748.htm> - Distribution Policy for Autodesk Animation Players

Programming for the artist board
Written by Carl Stika on Thu May 01 17:54:06 1997

Hi,

I was wondering if there are going to be different lib's for the artist board. Right now the libraries only work for visual c++ 1.5, 16 bit compiler. But I would like to be able to write programs in visual c++ 4.??, 32 bit compiler. Are there going to be different libs for this. Is this the right place for this message, or is there somebody I can talk to about this.

Thanks, Carl.

Re: Programming for the artist board
Written by Mike Fulton on Thu May 01 19:58:26 1997

To Carl Stika:

There is a 32-bit version of the .LIB files, but currently no 32-bit DLL, so you have to link in the LIB instead of using a DLL.

This isn't on the web site yet, but should be there pretty soon, so keep your eyes open for it.

Mike

Re: Programming for the artist board
Written by gshaw on Tue May 27 20:35:59 1997

I've got the 32bit libraries working but the programs do not work under windows NT. Is there any plans to change this. Most of our artists work under NT so the card is fairly useless without this support.

Re: Re: Programming for the artist board
Written by Mike Fulton on Wed May 28 18:08:44 1997

To gshaw:

The main problem is that the libraries directly access the artist board hardware, which is not allowed under Windows NT.

In order to make it work under Windows NT, a device driver would have to be written and this is not a trivial task.

I'm not aware of any plans to write a Windows NT device driver and make a new version of the library that works through it, but I'll pass along the suggestion to SCEI.

Mike @ SCEA

Lightwave to RSD converter?
Written by Tim on Thu Apr 24 16:26:47 1997

Does anyone know if there is a .lwo to .rsd 3d file converter, the plugins that are available, do not save the texture information, does anyone have a tool they use or a converter we can obtain?

Thanks

Re: Lightwave to RSD converter?
Written by Brian Dawson on Sat May 03 00:09:11 1997

To Tim:

The Lightwave Plugin does export out to the RSD format. RSD by definition contains only the texture names thus the VRAM locations are not resolved until the TMD to RSD conversion stage. This is true for all of the PSX Plugins that output RSD. Currently I am not aware of a tool for handling Lightwave objects such as the one your requesting, though I will keep an eye out for a solution. Maybe one of the other developers has written such a tool they can share with you. I may be misunderstanding you exact need, if so please repost with additional clarification.

Brian Dawson
SCEA Developer Support

BMP2TIM 16-bit textures?

Written by David Minogue on Wed Apr 09 16:12:28 1997

Can BMP2TIM be used to convert true-color BMPs (24-bit) into 16-bit TIMs? I know that I can do it with TIMUTIL, but I would prefer to be able to do it in a batch file.

- David Minogue.

Re: BMP2TIM 16-bit textures?
Written by Mike Fulton on Mon Apr 14 18:10:13 1997

To David Minogue:

Here's more info on doing the 24-bit to 16-bit conversion with TIMUTIL.

- 1) Load TIMUTIL
- 2) Bring up the "Preferences" window and set your desired output options.
- 3) From Windows Explorer or File Manager, drag and drop all of your 24-bit BMP files to TIMUTIL.
- 4) Use the "Save All" menu option of TIMUTIL to write out your 16-bit TIM files.

I know it's not as convenient as using a batch file, but at least you don't have to do every single file individually.

Mike

Re: BMP2TIM 16-bit textures?
Written by Mike Fulton on Wed Apr 09 22:15:27 1997

To David Minogue:

BMP2TIM will create a 24-bit TIM if given a 24-bit BMP as source.

However, both BMP and TIM formats are well documented so you should be able to make a tool that does this yourself.

I've suggested this idea to SCEI as a feature addition to BMP2TIM, but it would probably be faster if you did something yourself in the meantime rather than wait.

Mike

Texture Mapping Tools

Written by Chris Reese on Wed Feb 12 18:35:38 1997

What is a good texture mapping tool for meshes used for Playstation development.

What are some opinions about:

MeshPaint
4D Paint
Detailer.

Re: Texture Mapping Tools
Written by kozman on Wed Feb 19 00:50:44 1997

To Chris Reese:

I haven't had any experience with these programs. I have heard the 4D paint is good though. That program works somewhat like photoshop. You can get a demo from their website at "<http://www.4dvisasia.co.nz/>". I have asked the various companies how many people use their programs. I will post a reply when I hear back from them.

Michael Koziniak

How do I use SGI and Alias tools?

Written by Simon Freeman on Wed Feb 12 03:19:34 1997

Hi,

I'm trying to use the tools at:

http://www.sepc.sony.com/dev/art_files/alias.zip and
http://www.sepc.sony.com/dev/art_files/sgi.zip

These are zipped, lzh'd, gz'd and tar'd.

Could someone please walk me through getting the respective files onto the SGI, either the files from your website or those from developer CD 1.8?

We have WEB access on PC and Mac, software on Mac and SGI for dealing with LHA (LZH), Stuffit and Winzip.

The files will uncompress but after being UNCOMPRESSED from the .Z extension the .tar file is junk.

Thanks for the help,

#6

Re: How do I use SGI and Alias tools?
Written by Dave Scheele on Mon Feb 17 15:44:46 1997

To Simon Freeman:

Wow, deja vu ... I was about to ask the very same question, practically word-for-word! I'm an SGI neophyte, so any help would be appreciated.

Dave

Re: Re: How do I use SGI and Alias tools?
Written by kozman on Thu Feb 20 22:23:15 1997

To Dave Scheele:

The Alias and SGI compressed files are in good condition and will decompress. The trouble is the decompression of these files should be done mainly on the SGI. The Alias.zip file is a zip which turns into a lzh file which contains a compressed tar file. While it is okay to unzip the file and then decompress the lzh, the final uncompression and installation of the resulting SO file needs to be done on the SGI using the installation instructions outlined in the install.rea so that the files are placed correctly within an Alias environment

The SGI.ZIP file is a collection of compressed tar files. To decompress to the tar level I would recommend using Winzip. Once you uncompress one of the .z files it will ask you what to name it. Name it anything and then untar the file with: tar xf foo_tar

Questions on this convoluted mess of different compressed files is welcome.

- mk

MimeWave Tool

Written by Chris Reese on Wed Feb 12 02:49:26 1997

We are using the MimeWave tool for Windows and I wanted to know where I might find the source code to the mimewave.cpe file.

Re: MimeWave Tool
Written by Michael Koziniak on Tue May 20 01:21:23 1997

To Chris Reese:

You have the source for this. Its part of the Graphic Artist Tools located in on the CD in PSXGRAPH\SRC\PREVIEW.

- mk

Automating Media Conversion

Written by Geoffry Meek on Thu Jan 16 22:26:52 1997

I am looking for a way to automate the process of converting many Windows (.AVI) video files to PlayStation format (STR). Obviously, we have been using MovConv which does the job, but what I would really like is an interface to convert the media using a procedural programming language (C, preferably). So, in short, I'm looking for a static C library or DLL which I can call to convert AVI files to STR files. Does something like this exist? Any help would be appreciated.

Thanks,
Geoff Meek
Virtual Music Entertainment, Inc.

Re: Automating Media Conversion
Written by Mike Fulton on Fri Jan 17 23:04:13 1997

To Geoffry Meek:

There is no C library or DLL for DOS or Windows, however the newest versions of Movie Converter support scripting. The best way to do what you want is to have your program write scripts for Movie Converter.

Look for the new 2.1e in the download area within the next few days.

Mike

Forum: Sound Programmers

InitMalloc and reverb

Written by bd on Thu May 29 00:45:42 1997

SpuSetReverbModeParam seems to return -1 unless I have previously called SpuInitMalloc. I don't want to call SpuInitMalloc. Is there any way to get reverb working without messing with SpuMalloc and its ilk? I'm using version 3.5 of the libraries. Thanks!

Ben

Re: InitMalloc and reverb
Written by Rob Vawter on Thu May 29 20:51:09 1997

To Ben:

Unfortunately, SpuSetReverbModeParam checks to ensure that the reverb mode work size you need is not being occupied by a waveform at that moment in time. This internal check fails if the spumalloclist has no entries. Try SpuInitMalloc() with a num of 1. Sorry.

Voice steal priority - how to control?
Written by Ian Schmidt on Tue May 13 21:41:33 1997

I'm having a problem where some of my background ambient samples spontaneously drop out when I'm playing a lot of 'foreground' sound effects in a short period of time. I figured this is a voice stealing problem so I had my sound artist set the priority of the ambience sounds to '0' and all the other sounds to like '5' in Sound Deli. This seems to have helped cause the stealing to happen less often, but it still happens, which has me quite confused. (I'm assuming that 0 is highest priority, since it's not documented anywhere that I've been able to determine). Is there a sure-fire way to prevent a voice from ever being stolen, other than reserving the voice(s) entirely and using libsnd calls to play them?

Re: Voice steal priority - how to control?
Written by Rob Vawter on Thu May 22 18:59:02 1997

To Ian Schmidt:

Even if you set the priorities of the sounds you want to hear at all times to 127, it could be possible to lose these sounds if you have too many priority 127 tones attempting to play at one time. In the case of 25 tones of priority 127 playing at the same time, the oldest (in terms of key on timing) voice will be keyed off and replaced. You can use SsSetReservedVoice() after SsInit(), but before any playback of sound begins to set aside the number of voices available to the voice allocation system of libsnd. Then, with the leftover voices (ex:

SsSetReservedVoice(22) leaves voices 22 and 23 outside of the voice allocation system of libsnd), you can use the SsUtKeyOnV()/OffV() functions or libspu functions to ensure that these very important sounds ALWAYS play.

Your old message:

I'm having a problem where some of my background ambient samples spontaneously drop out when I'm playing a lot of 'foreground' sound effects in a short period of time. I figured this is a voice stealing problem so I had my sound artist set the priority of the ambience sounds to '0' and all the other sounds to like '5' in Sound Deli. This seems to have helped cause the stealing to happen less often, but it still happens, which has me quite confused. (I'm assuming that 0 is highest priority, since

it's not documented anywhere that I've been able to determine). Is there a sure-fire way to prevent a voice from ever being stolen, other than reserving the voice(s) entirely and using libsnd calls to play them?

Re: Voice steal priority - how to control?
Written by Stan Weaver on Wed May 14 18:16:30 1997

Hi Ian,

Actually, 127 is the highest priority, 0 is the lowest.

Stan Weaver
SCEA

Program #s in .SEQ files
Written by Abby Schneider on Fri May 09 22:20:31 1997

How do you configure the MIDI files so that the right midi track sees the right program #? I was doing it with the sysex commands, and the .SEQ file did not maintain the same numbers.

mkvab
Written by Ron Nakada on Tue Apr 29 19:28:14 1997

Is it possible to give MKVAB a large list of vags without using *.vag as an argument? It doesn't seem to take a response file (i.e. @vags.lst) as an argument, which would be ideal.

- Ron

Re: mkvab
Written by Rob Vawter on Wed Apr 30 17:29:20 1997

To Ron Nakada:

No, sorry, it isn't possible to give MKVAB a large list of vags using a response file. The attached file is a more complete document for MKVAB's usage

Name: mkvab.pdf
Content-type: application/pdf
Size: 174812 bytes

Mixing levels

Written by Tom Rolfs on Tue Apr 29 00:42:42 1997

I presume that the mixer on the PSX does not handle overflow. I presume this as we are getting major distortion when a few SFX's are playing at the same time and at full volume.

Questions:

- 1) Anyone else experience this problem? If not then maybe this is not a mixing problem.
- 2) If you did have this problem did you solve it and if so are you willing to give away your hard earned knowledge to us freeloaders?
- 3) Is there already support on the PSX, either hardware or library, for level compression or level limiting?

In my own mind I think the problem is an easy one to solve but I might be overlooking the obvious - and certainly the not so obvious. My current thinking is to monitor the sum of the volume levels for each active channel and calculate a global attenuation modifier to compress the levels.

I have not started work on this as yet and probably won't for a while but I thought I would first ask the experts.

Tr.

SpuGetKeyStatus / sound polling

Written by Bruce on Fri Apr 11 20:24:00 1997

I'm attempting to play some 'one-shot' sounds where I just want to start them playing on given voices/channels and later, when a sound finishes playing, reclaim that voice it was using for another sound. In order to reclaim the voice, I want to know if the sound has stopped playing. Initially I thought 'SpuGetKeyStatus' would handle this, but it appears that once I turn a voice on (SpuSetKeyOn), 'SpuGetKeyStatus' returns SPU_ON regardless whether the sound has finished playing or not. Am I using this function incorrectly, or more precisely, is there a way to poll a voice to see if a sound has finished playing?

thanx

Re: SpuGetKeyStatus / sound polling
Written by David Ross on Thu May 08 19:54:39 1997

To Bruce:

We are also having the EXACT SAME problem. We do have a little more info and a demonstration program. The info is that the problem appears to be sample-rate dependent.

My sound programmers description:

I've been working with Sony VAGs and I have difficulties detecting the end of a sound effect. The problem occurs when I play certain 6K Hz sound effects. When I play the VAG and read the channel key status, the channel always reads SPU_ON. In an attempt to identify what the problem is, I played

the same sound at different sampling rates (11K hz and found that when the sampling rate was changed the problem went away. What I am curious about is what causes this problem? Any help would be appreciated.

I have included one of the offending sound samples (sampled at 6K Hz and 11K Hz) and a small Sony VAG player which may be of some help. To use the program, just select the sound number (0 = 6K, 1 = 11K) by using the d-pad up and down arrows and then press start to hear the sound. The 16 numbers at the top of the screen are the returned values from SpuGetKeyStatus() for the first 16 channels. Sound number 0 will always leave channel 0 set at SPU_ON and sound 1 will return a value other than SPU_ON (3 = SPU_ON_ENV_OFF?). Why won't sound number 0 ever change from SPU_ON????

- Dave Ross, dross@sculptured.com

PS: Uploading only supported by NETSCAPE BROWSERS!?!?!? So I can't upload this demo program???? How long will Netscape be the only browser supported???

Re: SpuGetKeyStatus / sound polling
Written by Rob Vawter on Mon Apr 14 18:46:51 1997

To Bruce:

You are probably using SpuGetKeyStatus() correctly and also probably have a problem with your VAG data, as T. Rolfs suggests. What is the process for creating your VAG data? I'm guessing .WAV -> .AIF -> VAG. Or maybe 8 bit .AIF -> 16 bit .AIF. Let me know.

To Scott Turner:

The latency exists for two reasons:

- 1) All key on functions in libsnd queue up and wait for the next SsSeqCalledTbyT() to execute. This design allows for more key on requests to take advantage of the libsnd voice allocation management system, which is based on tone priority, size of envelope, and age of keyed on voice.
- 2) Are you checking to see the exact value SpuGetKeyStatus() is returning? It may be returning SPU_ON_ENV_OFF after the first tick, due to the 16 bytes of zero data which make up a VAG body. Also, it may take a few more ticks to actually have an envelope, depending upon the sample that was used to create the VAG.

My questions for you (Scott):

What are you trying to accomplish using SpuGetKeyStatus() and SsUtKeyOnV()? if your goal is to set up a system whereby all voices get maximum use, I suggest that you switch to just using SsUtKeyOn(). That functions calls an algorithm which searches for a voice which is in the SPU_OFF_ENV_OFF state. If it finds no voices in that state (and also available to libsnd), it finds a voice with a lower tone priority than the priority of the tone you are attempting to key on. If no voices have a lower priority, but some have the same priority, the algorithm finds the smallest envelope and oldest voice and uses that slot. Its an excellent system. Feel free to give the hotline a call if you'd like to discuss your goals (i.e. if I have guessed wrong about what you are trying to do) or discuss this system further.

- Rob

Old messages below:

From Bruce:

I'm attempting to play some 'one-shot' sounds where I just want to start them playing on given voices/channels and later, when a sound finishes playing, reclaim that voice it was using for another sound. In order to reclaim the voice, I want to know if the sound has stopped playing. Initially I thought SpuGetKeyStatus() would handle this, but it appears that once I key a voice on, SpuGetKeyStatus() returns SPU_ON regardless whether the sound has finished playing or not. Am I using this function incorrectly, or more precisely, is there a way to poll a voice to see if a sound has finished playing?

From T. Rolfs:

We had a problem exactly like this however it only happen on certain sound effects. We concluded that a bad VAG body was being created but we never understood why. Have you tried playing different samples?

I did have a sample program that demonstrates this problem (assuming it is the same problem), but I seem to have misplaced it. [I did upload it to the BBS, maybe it is still there] Our fix was to edit the original sample slightly and reprocess until the problem went away.

From Scott Turner:

I'm having a similar problem. I am attempting to accomplish the same task that you are, reclaiming voices when one-shots are done playing. I turn on the voice with SsUtKeyOnV(), and then poll to see if it's done using SpuGetKeyStatus(). The problem is, most of the time after I key on the voice, SpuGetKeyStatus() will return SPU_OFF for 1 to 5 times before it finally returns SPU_ON. I don't remember reading about this latency problem in the documentation. Is there a way to fix this, and to add to Bruce's question, a way to poll a voice to see if a sound is playing, and not have it return a false reading for the first few times I check?

Re: SpuGetKeyStatus / sound polling
Written by Tom Rolfs on Sat Apr 12 00:37:53 1997

To Bruce:

We had a problem exactly like this however it only happen on certain sound effects. We concluded that a bad VAG body was being create but we never understood why. Have you tried playing different samples?

I did have a sample program that demostrates this problem (assuming it is the same problem), but I seem to have misplaced it. [I did upload it to the BBS ounce apon a time, maybe it is still there]

Our fix was to edit the original sample slightly and reprocess until the problem went away.

Re: SpuGetKeyStatus / sound polling
Written by Scott Turner on Fri Apr 11 22:54:17 1997

To Bruce:

I'm having a similar problem. I am attempting to accomplish the same task that you are, reclaiming voices when one-shots are done playing. I turn on the voice with SsUtKeyOnV, and then I poll to see if it's done with SpuGetKeyStatus. The problem is, most of the time after I key on the voice, SpuGetKeyStatus will return SPU_OFF for 1 to 5 times before it finally returns SPU_ON. I don't remember reading about this latency problem in the documentation. Is there a way to fix this, and to add to Bruce's question, a way to poll a voice to see if a sound is playing, and not have it return a false reading for the first few times I check?

Thank you

What is the format of a VAG

Written by Gil Colgate on Thu Apr 03 00:11:30 1997

Does anyone know the format of a VAG (not a VAB) both after aiff2vag is done with it AND what it's format is when it is inside a VAB file (it's different).

Re: VAG/VAB format
Written by Rob Vawter on Thu Apr 03 17:31:27 1997

To Gil Colgate:

The only difference between VAG format and VAB format is the header - there is a 48 byte header on a VAG that is removed when the VAG is stuffed into a VAB. Also, there is variable size header on the VAB before the first VAG. The VAG format is proprietary, but here are some nice docs to look at.

- Rob

Name: VAGForm2.doc

Content-type: zz-application/zz-winassoc-doc

Size: 33792 bytes

Name: VABHEAD3.doc

Content-type: zz-application/zz-winassoc-doc

Size: 29696 bytes

SPU transfers during streaming

Written by Ian Schmidt on Fri Mar 21 22:31:29 1997

Is it possible to safely do transfers of any kind to and from SPU RAM while doing SPU streaming?

Based on the diagram in the Library 3.6 beta reference, I tried doing some small (512 byte)

SpuRead(s) during my stream's transfer completed callback, and I got the data ok, but this caused the streamed sound to skip erratically. I tried saving and restoring the SPU write address, to no avail. Can this work?

Re: SPU transfers during streaming
Written by Rob Vawter on Fri Mar 28 19:34:02 1997

To Ian Schmidt:

The answer is “yes, maybe.” It depends on a variety of factors -

1) The number of voices are streaming. More voices take up more of the SPU DMA time. 2) The size of the stream buffer. 3) The pitch of the stream. #2 and #3 go hand in hand - the transfer to refill the first half of the buffer is triggered when the IRQ of the end of the first half of the buffer is reached. If the playback of the waveform data is faster (greater pitch value), this IRQ is triggered more often. If the stream buffer size is smaller, this IRQ is triggered more often.

I'm sorry I don't have any numbers for you at this time - I'll see if I can come up with anything. The possibility has been bounced around to make other transfers to and from SPU RAM have an option to be of lower priority and thus stopped midway to ensure proper streaming. I'm not sure how high on the priority list of coding this is, I'll look into that as well.

Also, I have a few questions for you:

1) What code are you using to detect the end of the SpuRead()? SpuIsTransferCompleted()? A callback? Are you saving the SPU streaming system transfer callback before setting a new transfer callback and then restoring the old one?

2) Are you just using SPU RAM as a storage area because you are out of main RAM? SPU DMA is fairly slow, so this really isn't all that great an option.

Sorry for the slow response - I've been inundated of late.

- Rob Vawter

Re: Re: SPU transfers during streaming
Written by Ian Schmidt on Tue Apr 01 22:42:11 1997

To Rob Vawter:

I'm using SpuIsTransferCompleted(). Does this disturb the callback? I'd like to use part of SPU RAM basically as a cache for some things that I can't afford to keep in main RAM and I don't want to hit the CD to get, since as slow as SPU DMA is it's still far better than the CD. For reference, I'm streaming 2 voices at 22 kHz with 16k in-SPURAM buffers for each voice.

Thanks,
Ian Schmidt

Re: Re: Re: SPU transfers during streaming
Written by Tom Rolfs on Fri Apr 04 19:55:19 1997

To Ian Schmidt:

If you are going to do your own sound ram DMA you will have to remove the SpuSt callback first and then restore it when your transfer is done, otherwise your transfer will cause the SpuSt callback to get called and it will think that an audio transfer happened and so moves on to the next block of audio, this is why you get the skipping.

Note: I would recommend making your DMA's non blocking, i.e. don't use SpuIsTransferComplete(), use the Spu DMA complete callback. This way you are not waiting around in an interrupt too long.

Re: SPU transfers during streaming
Written by Rob Vawter on Mon Apr 07 22:28:57 1997

To Ian Schmidt:

The thumbs down symbol is the rating of my post from 4/4. I answered this on Friday, but the second I got in my car, I realized I'd left you a less than perfect answer, so now I'm changing it.

SpuRead() checks to see if a transfer finished callback exists. If a callback does NOT exist, an internal variable which keeps track of whether the SPU DMA is in use gets set to CONTINUE. If a callback does exist, this internal variable stays set at FINISHED. Now, since that is your case, two bad things will happen. #1: SpuIsTransferCompleted() will immediately return, pretending that the SPU DMA is done, which it is not. #2: When the SPU DMA is actually done, the transfer callback for your streaming will get called and streaming will proceed incorrectly.

So, Mr. Rolfs is correct; you should save the pointer to your old transfer completed callback function and set up a callback to handle the SpuRead(), which then restores the streaming transfer callback.

Please keep in mind that the info about safe SpuRead() or SpuWrite() limitations still holds - the amount of data you can transfer will depend on the pitch of the streams, the sizes of the half buffers, and now also, to a much lesser extent, the length of time spent in your SpuRead() or SpuWrite() callback function.

- REV

Spu Transfers don't work

Written by Edgar Hahm on Fri Mar 21 00:56:35 1997

I can't get any SPU transfer code to run properly under the PSY-Q debugger - not even the example code from the developer CD's.

The problem occurs when I call SpuWrite(). After any call to SpuWrite(), SpuIsTransferCompleted(SPU_TRANSFER_WAIT) never returns. I can call SpuIsTransferCompleted(SPU_TRANSFER_WAIT) before starting the transfer and it returns just fine.

Has anybody run into this problem before?

- Paul

Re: Spu Transfers don't work
Written by frank on Fri Mar 21 18:35:16 1997

To Edgar Hahm:

Which samples are you trying? I tried a couple of them out, tuto2 and tuto3 in the sample\sound\tuto directory, and they worked fine in the debugger.

Can you get the samples to run in non-debug mode? If not, be sure you have either a cd-rom drive or a blue terminator plug hooked to your dev board.

Frank

Re: Re: Spu Transfers don't work
Written by Edgar Hahm on Fri Mar 21 23:04:29 1997

To frank:

Oh! I should add my SPU transfer code runs just fine on a blue playstation with a Datel adaptor.

The code still DOESN'T work when I run it with run.exe and testmess.exe.

I'll look into whether it will run in debug mode.

- Paul

Re: Re: Spu Transfers don't work
Written by Edgar Hahm on Fri Mar 21 23:00:34 1997

To frank:

I've tried all of the examples in sample\sound\lowlevel, and they all end up waiting indefinitely for an SPU transfer to wait. The "balls" demo hangs on SsVabTransCompleted(SS_WAIT_COMPLETED).

It happens on both of our PSX development stations. I can step through the reset of the code no problem - PadInit(), InitGraph(), etc all work fine.

Another problem is that the text I/O is VERY slow - like less than 300 baud.

Paul

Re: Re: Re: Spu Transfers don't work
Written by frank on Mon Mar 24 19:21:18 1997

To Edgar Hahm:

Let's make sure we're on the same page.

Could you tell me what version of the libraries, debugger, etc. that you are using? (If they all came off the same CD, which release of that will suffice) I don't know what you mean when you refer to sample\sound\lowlevel.

Are you unable to run the samples with the RUN.EXE utility, DBUGPSX, or both? In different messages you say different things.

Do you have either a blue terminator plug or a DTL-2010 attached to your dev system(s)?

The text I/O problem—is this on happening on both systems? There are a number of things that could cause this. For starters, check to be sure there are no IRQ conflicts between your dev boards and other devices in your system.

Frank

Re: Re: Re: Re: Spu Transfers don't work
Written by Edgar Hahm on Mon Mar 24 21:49:19 1997

To frank:

By the way, the SPU transfer code doesn't work with RUN.EXE or with DBUGPSX. The rest of the code (InitPad, InitGraph, etc) works with both.

- Paul

Re: Re: Re: Re: Spu Transfers don't work
Written by Edgar Hahm on Mon Mar 24 21:46:53 1997

To frank:

I installed the tools and libraries from from the PlayStation Programmer Tools CD, Release 1.8.

sample\sound\lowlevel is a set of 6 example programs that demonstrate the use of LibSpu. They're not on Release 1.8, so presumably they came from an earlier release. I've attached 3 of them in case you'd like to take a look.

I have a blue terminator plug but it is not attached. Should I attach it? I'm afraid I don't know what a DTL-2010 is. Do I need one? I do have a CD emulator board attached. Regarding the hardware conflicts, it looks like there may be some funky stuff happening on my system. I'll look into it.

Thanks very much for your help.

- Paul

Re: Re: Re: Re: Re: Spu Transfers don't work
Written by frank on Mon Mar 24 23:09:01 1997

To Edgar Hahm:

The DTL-2010 is the external Playstation CD-ROM drive, aka the "black box."

In order to run the sound samples, you MUST have one of the following installed:

a) a "black box"

or

b) a blue terminator plug

The only reason to remove the blue plug (which is installed when you receive your boards) is in order to install a DTL-2010. If you don't have an external CD-ROM, put the plug back in. This should solve your problem.

Frank

Re: Re: Re: Re: Re: SPU Transfers don't work
Written by Edgar Hahm on Tue Mar 25 02:22:24 1997

To frank:

I installed the blue terminator and it works! Thanks a lot!

Regarding the slow character I/O, I don't see any conflicts on my machine. In hardware manager there is an "Unknown Device" with a yellow exclamation mark over it that I presume is the PSX dev board (the reason that I presume this is that the only other person with one of these board has the same thing show up in hardware manager).

I understand that the normal rate for text I/O is 1 character per frame. I seem to be getting on the order of 1/4 of that - like 6 seconds to do a full line of text.

- Paul

SPU Streaming Problem

Written by David Minogue on Tue Mar 18 18:55:37 1997

(Message from Richard Frankish):

the SPU stream problem we're having occurs if the game is left in the attract loop for a long time. The announcer samples are loaded from CD and then streamed from a main memory buffer to SPU RAM. The buffer enables (near) seamless joining of speech sections that may be many tracks apart on the CD.

We migrated to 3.5 libraries to try to fix the occasional CD error we were getting. Now the SPU stream just hangs occasionally with the announcer continually repeating the current contents of the SPU stream buffer. We can detect this because we have a watchdog timer that counts game iterations between successive 'buffer transfer complete' callbacks. Setting the stream environment command to SPU_ST_STOP (the usual way of stopping an SPU stream) has no effect. SpuStQuit() fails because the stream is stuck in SPU_ST_TRANSFER mode and it needs to be in SPU_ST_IDLE to quit the stream. The rest of the audio still works.

We've also tried 3.6 and 3.7 libraries are now going back to 3.4.

Rick

Re: SPU Streaming Problem
Written by Rob Vawter on Wed Mar 19 17:09:08 1997

To Richard Frankish:

I'm looking into this for and should be able to come up with some probable causes and solution in the next two days.

- Rob

Re: Re: SPU Streaming Problem
Written by David Minogue on Tue Apr 01 15:33:45 1997

To Rob Vawter:

Rick is still struggling with this streaming problem after 2 weeks. It's holding up the game.

Have you found anything yet? Can I get you any more information, code, whatever?

BTW: the problem was not eliminated by going back to lib 3.4.

- David.

multiple vabs

Written by Joseph Wilkerson on Mon Mar 10 23:06:46 1997

When I try to load a second VAB SsVABOpenHead always returns a -1.

Re: multiple vabs
Written by frank on Tue Mar 11 02:06:49 1997

To Joseph Wilkerson:

Please send more info next time, such as the actual code you are using, whether or not switching the order of the VABs matters, whether or not you are only using libsnd, what version of the libraries you are using etc. This will make it easier for me to pin down the problem.

Paraphrased from August 2nd, 1996 post in the Sound Programmers conference on the BBS:

Here are all the reasons SsVabOpenHead can fail: 1) SPU still has a transfer in progress. Did you use SsTransCompleted()?

2) The VabID passed is out of range (>=16).

- 3) No available VabId slots remain (16 VABs already open; definitely not the cause of your problem).
- 4) VabID passed is already in use.
- 5) VAB header format invalid ("VAB" not found)
- 6) Too many programs contained in the VAB - caused by parse error of the VAB header. Most likely to be caused by overwriting the header info or by very very old VAB formats.
- 7) SpuMalloc call failed. Either not enough room left in SPU RAM (512K-reverb work area - 0x1010 [spu decode data region+ additional loop info region]) or SPU RAM too fragmented by other transfers of VAGs or SpuInitMalloc() called by user with too few spumallocs possible.

Sound problems can also be caused by failing to have a black CD drive or blue plug in the boards and by improper initialization order. However, these seem to only manifest themselves in a failure of SsTransCompleted().

Are you using libtap?

- Rob Vawter

LFO

Written by bd on Thu Mar 06 01:58:20 1997

Is there some example code that uses the LFO features of the SPU? I'd like to emulate some kind of midi modulation effects.

Thanks!

Ben

Re: Pitch LFO sample code
Written by Rob Vawter on Sun Mar 09 22:15:26 1997

To bd:

Currently, there is no sample code for the Pitch LFO functions in libspu. I'll work on some this week and post it.

SsUtKeyOnV - No sound playback
Written by Rob Vawter on Tue Feb 18 18:09:29 1997

Forwarded from the BBS:

Harmon: I am using SsUtKeyOnV to play sound effects and sometimes no sound is played. I know this isn't a very detailed description, but that's all there is to it. Calls to the function work most of the time but there is the occasional time when a sound that has been playing fine will just not play. Why would this be happening?

Vawter: There are a variety of reasons. Let's make sure I understand the situation - Let's say program 1, tone 1 has been previously keyed on and now it doesn't key on. If future key ons of this same tone

don't work, then the VAB header was probably overwritten. If future key on attempts do work, then depending on the value returned by the call (voice or -1) the following could be occurring: 1) Consecutive calls to the function could be happening with too little time between. SsUtKeyOnV() is not reentrant. 2) Bad program, vabId, or voice # input. Unlikely.

I think most likely #1. Would this fit in with the behavior?

Switching XA audio tracks

Written by Rob Vawter on Wed Feb 05 17:13:05 1997

Some of you may be interested in switching XA audio tracks, as was discussed briefly in the Sept/Oct Conference Sound Session. Here's a nice little thread which discusses some problems and solutions...

Villeta: We are running tuto6.c from the samples (CD) and we are getting a sound glitch when we change filter by any amount except positive one. Do you know what's going on? Please let us know either way.

Burrowes: I have to admit that I haven't tried this before. The XA stuff I've been doing involves switching the filter very quickly during GAPS in the XA source (silence). The one game that I know of that does do seamless switching during music playback does simple branching which I'm sure simply moves from the 1st track through to the 8th (changes of positive 1 as you described above.) Maybe they found the same problem you did.

I can see why this happens...

The XA blocks (sectors) are arriving at the XA decoder's buffer at just the right speed to keep the buffer full and yet not overrun. Let's say the blocks below represent XA Blocks (not sectors) arranged on the disc so that they arrive at the decoder in time...

cd reading --->

[1] [2] [3] [4] [5] [6] [7] [8] [1] [2] [3] [4] [5] [6] [7] [8] [1] [2] [3] [4] [5] [6] [7] [8] [1] [2] [3] [4] [5]

[6] [7] [8]

If we are decoding stream 1 and switch to stream 4 you can see we have 2 choices...

1) Read the next stream [4] block which is now EARLY thereby killing the buffered block currently being decoded, or... 2) Skip the next stream [4] block and read the following stream [4] block thereby allowing the XA decode buffer to run out of data and wait for the next block!

Either way you're gonna have a problem! Now, when you simply skip 1 stream forward (from 1 to 2 for example) either... 1) the very next [2] block is read, overwriting the decode buffer BEFORE the decoder has had a chance to start on the [1] block it had already loaded, or... 2) the next [2] block is skipped and the following [2] block is ready allowing the decoder to run out of data for a VERY small period of time.

I suspect option 1 is what is happening (or something close to it.)

Bad news... I'd be surprised if there was any way around this.

Mr. Tanaka and Mr Suzuoki:(Two Japanese engineers)At first, the noise is occurred if you switch the XA filter during the playing. Because,

1A) The XA decode refers the previous data, so it cannot decode (or it makes the wrong data) right after the filter is switched.

1B) Since XA audio is ADPCM data, sampling data is recorded as a differential value from the previous sample. Therefore we cannot change the ADPCM channel without noise. Some techniques of encoding will reduce it, but it cannot be deleted completely.

2A) If the DC level of new channel is different from previous channel, the noise is occurred even if both data are "silence".

2B) The noise may also be because of mismatch of DC level.

Sampling wave data is sometimes added some offsets (DC offset). When one audio data is changed to another that has different DC offset, the noise may be happen. In this case, the noise will be deleted if we adjust the DC level of each channel.

Conclusion A) So, please mute the volume while you switch the filter. I think 1 frame (16msec) is long enough to mute the volume.

Conclusion B) But this noise does not last long time, so we can conceal it using CdlMute/CdlDemute command when channel is changed. Anyway changing channel causes the unexpected jump of the wave form, and it is not strange that the noise happens at that moment.

Lib 3.6

Written by Chris Reese on Wed Jan 15 21:37:12 1997

I have just recently installed the 3.6 libraries and have successfully compiled our game code that was running under the 3.4 libs. Everything seems to run fine except the music. I have been getting crashes with music and sounds running in the game. If I don't play music or sound then everything is fine. I am about to dive in and see if I can find out where exactly it is crashing but before I went too far I just wanted to know if there is something I need to setup differently for the sound driver in lib 3.6 that may have changed from lib 3.4.

Re: Lib 3.6 libsnd?
Written by Rob Vawter on Thu Jan 16 01:17:22 1997

To Chris Reese:

Did you get the version of lib 3.6 from the BBS, the website, or the newest CD (which may still be winging its way toward you)? The BBS version was unfortunately bugged. I'm also assuming that you are talking about MIDI music in libsnd and perhaps sound effects using SEQs. Let me know. If you are using the BBS lib 3.6, switch to the website or CD version. If there is another problem, let me know and I'll look into it.

Re: Re: Lib 3.6 libsnd?
Written by Chris Reese on Fri Jan 17 03:11:03 1997

To Rob Vawter:

I did d/l the 3.6 libs from the BBS. I just d/l the libs from the WEB here and everything seems to work fine now. Thanks for the info.

Forum: Sound Artist

DevCon Topics

Written by Stan Weaver on Tue May 20 23:26:59 1997

Hello all,

Well, the Developers Conference is coming up and I'm wondering if anyone has any topics in particular they'd like hear discussed. Any burning issues? Now's the time to speak.

All the best,

Stan Weaver
Audio Tools Technician
SCEA

VABs for general MIDI

Written by Steven Taylor on Fri Apr 04 00:47:00 1997

I was wondering if it was possible to get VABs that have programs that correspond to the general MIDI format. Something that would have basic instrumentation for piano, drums, saxophone, etc. Perhaps a third party product?

Re: VABs for general MIDI
Written by Stan Weaver on Sat Apr 26 19:14:52 1997

Hi Steven,

For you and anyone else who is interested, a General MIDI VAB has been uploaded to this web site. It can be found in Files/Sound Artist Tools in both Mac (gm-sea.hqx) and PC

(genmid.zip) versions. Download them at your convenience, but remember that this VAB has severe limitations due to memory constraints.

All the best,

Stan

Re: VABs for general MIDI

Written by Rob Vawter on Mon Apr 07 16:48:43 1997

To Steven Taylor:

Chris Stevens at LoudMouth Inc. has some nice pre-built VABs for the PlayStation. His number is (503) 726-8447. Alos, sometime in the next week, I hope to upload the general MIDI VAB that comes with the Yaroze kit.

- Rob

Sound Deli Corrupts VABs

Written by Jay Halderman on Tue Feb 11 20:38:54 1997

I have run into a problem where my VABs are corrupted when I perform either delete line or insert line operations on programs and the VAG list. For example, if I have a program that needs deleted I will delete the line, but, since my project requires program numbers to stay the same for the rest of VAB, I then insert a line where I deleted it. Instead of inserting a line there, Deli creates a duplicate program from somewhere else within the VAB. Even inserting a line without deleting one prior causes this corruption. Anybody run into this and knows what causes it? Better yet, anyone have a good solution?

BTW I'm using Deli 2.3E.

Re: Sound Deli Corrupts VABs

Written by Rob Vawter on Wed Feb 26 16:15:07 1997

To Jay Halderman:

The bug involving line insertion and deletion has been fixed for Sound Deli 3.0E, which will be available sometime within the next week. Look for it in the files area soon.

For those who haven't experienced it before, here are the symptoms: insertion of a line (program) caused a duplicate line to also be inserted somewhere earlier in the VAB. Deletion of this extraneous line also deleted the line inserted, thus bringing you full circle to the point where you need to insert a new line ad nauseum.

Re: Re: Sound Deli Corrupts VABs
Written by Jay Halderman on Thu Mar 06 19:25:05 1997

To Rob Vawter:

Any progress on a fixed version of Sound Deli? I have another big VAB coming up, and don't want to go through the same problems as last time. I had already been using export marked files only to circumvent the bug, but had little luck using this technique with the VAB you are familiar with, leading to my contacting you. i.e. -I could really use the new version soon.

Thanks,

jhalderman@activision.com

Re: Re: Re: Sound Deli Corrupts VABs
Written by Rob Vawter on Tue Mar 11 03:41:52 1997

To Jay Halderman:

Oooof. Word is that 3.0E will be released "sometime within the next few weeks" I'll push hard for it. Sorry.

Greetings!

Written by Stan Weaver on Thu Feb 06 22:03:01 1997

Welcome to SCEA's Developer Support web site Sound Artist Forum!

My name is Stan Weaver and I'll be here to answer questions regarding the development of sound resources for the PlayStation from the artist's perspective. That means not only helping you with whatever problems you may be having, but also assisting you in any way possible to ensure the best results for your project. So, if you want to ask something specific or just talk shop, please, don't hesitate to post!

Best of luck,

Stan Weaver
Audio Tools Technician
Sony Computer Entertainment America
stan_weaver@interactive.sony.com
(415) 655-3561

Re: Greetings!
Written by Mark Miller on Fri Mar 14 00:04:27 1997

To Stan Weaver:

Hey, I found this in a read me somewhere else on the board... Where can I find this Excel table?
Could you send it to me?

“ Table which multiplies loop size by 28 (Microsoft Excel 4.0 sheet) - This is the table for calculating the sample rate conversion coefficient in order to make the loop size of the sound source with loops into a multiple of 28 using Sound Designer II. (In order to use this sheet, Microsoft Excel 4.0 or table calculation software which can read Excel format data is needed.) - Since the necessary coefficient (SR Convert Ratio or new fs) is calculated when loop start, end or fs are entered, please do the SR Convert and reset the loop start (new l.start) and end (new l.end) accordingly.”

Thanks,

MSM

Re: Greetings!

Written by Mark Miller on Thu Mar 13 23:47:57
1997

To Stan Weaver:

Hi Mark Miller again. Forgot my password, so I created a new account. Well, after a long wait, I finally found the offending AIF file. As you will recall, this does not convert correctly in the current AIFF2Vag, but converts fine in 1.6...

Yours,

MSM

BTW, Congrats on the new job. Say hi to Buzz for me...

Re: AIFF looping
Written by Stan Weaver on Fri Mar 14 17:49:46 1997

Hi Mark,

I took a look at your AIFF in Sound Designer and it seems that the problem stems from a lost loop end point. I'm not sure what waveform editor you're using, but I know in SD you have to leave a tail after loop end (even just one sample in length) so that it can be recognized as a complete loop. Everything else seems like it's OK. I've attached a self-extracting archive which contains your AIFF, a VAG of that file, a new version of your AIFF which I've made some

alterations to, and the VAG from that file. The loop points in my version of your AIFF are by no means optimal sound-wise, but it illustrates the point. I do find it rather odd that your file converts fine in v1.6. I'll have to look into that more.

All the best,

Stan.

PS - Thanks! And you may see a personal response from Buzz himself as he often prowls this web site as well! Take care.

Re: Greetings!
Written by Mark Miller on Sun Feb 16 04:05:22 1997

To Stan Weaver:

Hi, Mark Miller here. Just downloaded the latest (2/14) release of the sound tools. It seems that AIFF2VAG does some strange things with looping files. Is this a known problem or something that I am doing. I know that this is a very vague comment, but I have not had time to experiment (deadlines etc). I was getting crap at the end of a looped file that would not stop until I shut down AIFF2Vag and booted Sound deli. The file behaved the same way in sound deli. I will attach it for you enjoyment.

The same file converted fine in AIFF2VAG V1.6 which is the version that I have been using without problems for quite some time. V1.8 put clicks into all of my loops, so call me conservative....

Yours, Mark Miller

BTW Did you ever work at Sega? Your name rings a bell....