
Subject: Simple task manager for parallel processing in IDL

Posted by [wita](#) on Mon, 10 Mar 2008 15:37:53 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello all,

Some months ago I ran into a performance problem with ENVI on a multi-cpu machine. The solution was to switch of the threading in IDL, essentially turning my "quad core duo" into a "core solo" machine from the IDL perspective. To gain advantage of the multiple cpu's I created a workaround by spawning multiple IDL sessions which do the actual work.

The thing works through a process manager which should be provided with a list of tasks. The process manager then determines the number of cpu's and for each cpu it sets a lock file and spawns a process client that is given the task. This client then starts the actual process you want to carry out and when this finishes the client deletes the lock file. The process manager in the background monitors the existence of the lock files and when a lock comes available it again spawns a new client until all tasks are finished.

Some limitations:

- Only works for jobs which can be split in independent tasks, there is no inter-process communication.
- No effort to collect the output from different tasks
- No task farming for running on grids, it only works for a single machine with multiple CPU's. This could be added though by distributing tasks through a central tasklist (for example through a database server).

The code is available through the IDL code contribution library:

<https://www.itvis.com/codebank/search.asp?FID=537>

An example is included in the header of `cgi_process_manager.pro`. Any feedback on bugs is welcome, particularly on Windows machines as I do not have a multi-cpu windows machine around.

Now some of you may be wondering "why isn't this guy using the IDL_IDLbridge?". The point is that many of the work I do involves ENVI functionality and so far I had some problems to get ENVI to initialize on an IDL_IDLBridge session (although the IDLPATH settings on the bridge are OK):

```
IDL> r = obj_new('idl_idlbridge')
% Loaded DLM: IDL_IDLBRIDGE.
IDL> r -> execute, 'envi, /restore_base_save_files'
% IDL_IDLBRIDGE Error: Attempt to call undefined procedure/function:
```

```
'ENVI_IOW'.  
% Execution halted at: $MAIN$  
IDL> r -> execute, 'envi_batch_init'  
% IDL_IDLBRIDGE Error: Attempt to call undefined procedure/function:  
  'ENVI_RV'.  
% Execution halted at: $MAIN$
```

It needs some further testing whether ENVI is initialized properly on the bridge (despite the error messages) as some functionality seems to work properly.

Cheers,

Allard
