## Subject: Re: Parallel Processing
Posted by Russell Ryan on Thu, 05 Jul 2012 15:29:21 GMT

View Forum Message <> Reply to Message

On Thursday, July 5, 2012 11:16:27 AM UTC-4, Stefan wrote:
> Am Donnerstag, 5. Juli 2012 17:13:02 UTC+2 schrieb fawltyl...@gmail.com:
>> On Thursday, July 5, 2012 4:52:23 PM UTC+2, Stefan wrote:
>>> Hi
>>>
>>> Now that I have successfully implemented multi-threading another problem occured:
>>>
>>> To invoke multiple processes I start in a loop with
>>>
>>> bridges[i]->EXECUTE, 'program,par1,par2', /NOWAIT
>>>
>>> where 'bridges' is an object-array which holds the different child processes.
>>> Upon the execution of the last process I do
>>>
>>> bridges[i]->EXECUTE, 'program,par3,par4'
>>>
>>> And after that I destroy my bridges in a loop.
>>>
>>> Now I have a problem if the last process finishes before one of the previous processes since upon its completion it will directly move to the part where all bridges are destroyed and kills my program...
>>>
>>> Is there an easy way to tell IDL to wait for all my processes to finish and then destroy the bridges?
>>>
>>> thanks
>>> Stefan
>>
>> Use IDL_IDLBridge::Status to determine whether the job is finished. Pseudocode:
>>
>> ndone=0
>> running=replicate(1, njobs)
>> while ndone lt njobs begin
>>     for j=0,njobs-1 do begin
>>        if running[j] then begin
>>           query status of the j-th job with IDL_IDLBridge::Status
>>           if finished
>>              destroy bridge
>>              running[j]=0
>>              ndone++
>>           endif
>>        endif
>>     endfor
>>     wait, 1

>> endwhile
>>
>>
>> regards,
>> Lajos
>
> Ah, I see you put WAIT there if its not finished. I thought of using this in a loop but without putting WAIT you spend too many resources.
>
> Thanks! :)

Yeah, I put a wait in there too.  I've thought about "calibrating" the wait time.  It seems to me that if your process takes (say 30 minutes) then polling the bridges every 1 second seems like overkill. My gut-feeling is that this wait time should be roughly 1/3 or 1/4 the expected time per process, but of course you'd want to ensure that it's always waiting a minimum of ~0.5 s.

Russell