

**CS 580 Client-Server Programming**  
**Fall Semester, 2004**  
**Doc 24 Thread Pools**  
**Contents**

|   |    |
|---|----|
| Thread Pools.....   | 2  |
| Iterative Server .....                                    | 7  |
| Basic Concurrent Server.....                              | 8  |
| Concurrent Server With Thread Pool .....                  | 14 |
| Concurrent Server With Thread Pool & Thread Creation..... | 15 |
| How to reuse a Thread? .....                              | 19 |
| VisualWorks Example.....                                  | 20 |
| Java Example .....  | 23 |

**References**

Java Performance and Scalability Vol. 1, Dov Bulka, 2000

Past CS580 lecture notes

**Copyright** ©, All rights reserved. 2004 SDSU & Roger Whitney,  
5500 Campanile Drive, San Diego, CA 92182-7700 USA.

OpenContent (<http://www.opencontent.org/opl.shtml>) license  
defines the copyright on this document.

## Concurrent Server With Threads

```
import java.net.Socket;
import java.net.ServerSocket;
import java.io.*;
import java.util.Date;
import java.util.logging.*;
import sdsu.util.ProgramProperties;

public class DateServer {
    private static Logger log = Logger.getLogger("dateLogger");

    public static void main (String args[]) throws IOException {
        ProgramProperties flags = new ProgramProperties( args);
        int port = flags.getInt( "port" , 8765);
        new DateServer().run(port);
    }

    public void run(int port) throws IOException {
        Thread.currentThread().setPriority(Thread.NORM_PRIORITY + 1);
        ServerSocket input = new ServerSocket( port );
        log.info("Server running on port " + input.getLocalPort());
        while (true) {
            Socket client = input.accept();
            log.info("Request from " + client.getInetAddress());
            DateHandler clientHandler = new DateHandler(
                client.getInputStream(),
                client.getOutputStream());
            clientHandler.setPriority( Thread.NORM_PRIORITY);
            clientHandler.start();
        }
    }
}
```

```
class DateHandler extends Thread {
    private InputStream in;
    private OutputStream out;

    public DateHandler(InputStream fromClient,OutputStream toClient) {
        in = fromClient;
        out = toClient;
    }

    public void run() {
        try {
            processRequest();
        }
        catch (IOException error) {
            // Log and handle error
        }
    }
}
```

```
void processRequest() throws IOException {
    try {
        BufferedReader parsedInput =
            new BufferedReader(new InputStreamReader(in));

        boolean autoflushOn = true;
        PrintWriter parsedOutput = new PrintWriter(out, autoflushOn);

        String inputLine = parsedInput.readLine();

        if (inputLine.startsWith("date")) {
            Date now = new Date();
            parsedOutput.println(now.toString());
        }
    }
    finally {
        in.close();
        out.close();
    }
}
```

## Smalltalk Version

```
Smalltalk defineClass: #ConcurrentDateServer
  superclass: #{Core.Object}
  indexedType: #none
  private: false
  instanceVariableNames: 'serverSocket '
  classInstanceVariableNames: ''
  imports: ''
  category: 'cs580'
```

## ConcurrentDateServer class methods

```
port: anInteger
  ^super new setPort: anInteger
```

## ConcurrentDateServer instance methods

```
processRequestOn: anReadStream  
  | clientRequest |  
  clientRequest := anReadStream through: Character cr.  
  (clientRequest startsWith: 'date')  
    ifTrue: [anReadStream  
             nextPutAll: Time dateAndTimeNow printString].  
  anReadStream close  
  
run  
  | childSocket clientIOStream |  
  
  [childSocket := serverSocket accept.  
   clientIOStream := childSocket readAppendStream.  
   clientIOStream lineEndTransparent.  
   [self processRequestOn: clientIOStream] forkAt: 45]  
  repeat  
  
setPort: anInteger  
  serverSocket := SocketAccessor newTCPserverAtPort: anInteger.  
  serverSocket  
    listenFor: 4;  
    soReuseaddr: true  
  
shutdown  
  serverSocket close
```

## Thread Pools Iterative Server

```
while (true)
{
    Socket client = serverSocket.accept();
    Sequential code to handle request
}
```

### When usable

TP = Time to process a request

A = arrival time between two consecutive requests

Then we need  $TP \ll A$

## Basic Concurrent Server

```
while (true)
{
  Socket client = serverSocket.accept();
  Create a new thread to handle request
}
```

### When usable

Let  $TC$  = time to create a thread

Let  $A$  = arrival time between two consecutive requests

We need  $TC \ll A$

Often this is good enough

Thread consume resources

- Memory
- CPU cycles

A program has a limit of

- Threads it can productively support
- Sockets it can have open

We need to insure we don't create too many threads

## Thread Creation Time - VisualWorks

Time in milliseconds

| VisualWorks Smalltalk |                                  |       |        |         |
|-----------------------|----------------------------------|-------|--------|---------|
|                       | Iterations or Number Created (n) |       |        |         |
|                       | 100                              | 1,000 | 10,000 | 100,000 |
| Empty Loop            | 0                                | 0     | 0      | 1       |
| Integer add           | 0                                | 0     | 0      | 3       |
| Collection create     | 0                                | 0     | 4      | 34      |
| Thread create         | 1                                | 5     | 45     | 380     |
| Thread create & run   | 1                                | 5     | 152    | 1268    |
| Thread create & run   | 0                                | 3     | 82     | 1048    |

| Java                |                                  |       |        |         |
|---------------------|----------------------------------|-------|--------|---------|
|                     | Iterations or Number Created (n) |       |        |         |
|                     | 100                              | 1,000 | 10,000 | 100,000 |
| Empty Loop          | 0                                | 0     | 1      | 7       |
| Integer add         | 0                                | 0     | 1      | 7       |
| Vector create       | 0                                | 5     | 10     | 42      |
| Thread create       | 2                                | 62    | 185    | 1771    |
| Thread create & run | 40                               | 402   | 2626   | 24909   |
| Thread create & run | 51                               | 351   | 2507   | 24767   |

Run on

- Macintosh PowerBook with 1.25GHz PowerPC processor
- OS 10.3.3
- VW 7.2nc
- Java 1.4.2\_03 with HotSpot Client

## VisualWorks Details Statements Timed

|                     |  |
|---------------------|--|
| Loop                | n timesRepeat:[ ]                          |
| Integer add         | n timesRepeat:[ x := 3 +4]                 |
| Collection create   | n timesRepeat:[x := OrderedCollection new] |
| Thread create       | n timesRepeat:[ [x := 3 +4 ] newProcess]   |
| Thread create & run | n timesRepeat:[ [x := 3 +4 ] fork]         |

### Code used

Transcript clear.

#(100 1000 10000 100000) do:

[:n |

| x |

ObjectMemory garbageCollect.

time := Time millisecondsToRun:

[n timesRepeat: [[x := 3 + 4] fork]].

Transcript

print: n;

tab;

print: time;

tab;

print: x;

cr;

flush]

## Java Details - Program Used

```
import java.util.*;
import sdsu.util.Timer;

public class TimeTests {
    public static void main (String args[]) {
        Timer clock = new Timer();
        SampleThread x = new SampleThread();

        for (int n = 100; n < 200000; n = n * 10) {
            System.gc();
            clock.start();
            for (int k = 0; k < n; k++){
                x = new SampleThread();
                x.start();
            }
            long time = clock.stop();
            x.x();
            System.out.println("" + n + "\t" + time);
            clock.reset();
        }
    }
}
```

```
class SampleThread extends Thread {
    int x;
    public void run() {
        x = 3 + 4;
    }

    public int x() {
        return x;
    }
}
```

## Java Code Timed

|                     |   |
|---------------------|---|
| Loop                | <pre>for (int k = 0; k &lt; n; k++){<br/>    }</pre>  |
| Integer add         | <pre>for (int k = 0; k &lt; n; k++){<br/>    x = 3 + 4<br/>    }</pre>                                  |
| Collection create   | <pre>for (int k = 0; k &lt; n; k++){<br/>    x = new Vector();<br/>    }</pre>                          |
| Thread create       | <pre>for (int k = 0; k &lt; n; k++){<br/>    x = new SampleThread();<br/>    }</pre>                    |
| Thread create & run | <pre>for (int k = 0; k &lt; n; k++){<br/>    x = new SampleThread();<br/>    x.start();<br/>    }</pre> |

## **Warning about Micro-benchmarks**

Micro-benchmarks are

- Hard to do well
- Misleading

Better to measure the performance of your system

## Concurrent Server With Thread Pool

Create N worker threads

while (true)

{

Socket client = serverSocket.accept();

Use an existing worker thread to handle request

}

### When usable

TP = Time to process a request

A = arrival time between two consecutive requests

Then we need  $TP \ll A * N$

## Concurrent Server With Thread Pool & Thread Creation

Create N worker threads

while (true)

{

Socket client = serverSocket.accept();

if worker thread is idle

    Use an existing worker thread to handle request

else

    create new worker thread to handle the request

}

### When usable

Number of requests we can handle at a unit of time

$$TP / N + 1/TC$$

where N is not constant

## **What to do with the new Worker Threads?**

Client requests are not constant over time

Requests can come in bursts

Threads consume resources

Don't want a large pool of threads sitting idle

### **Common strategy**

Have a minimum number of threads in a pool

When needed add threads to the pool up to some maximum

When traffic slows down remove idle threads

## Threads & Memory Cache

Threads require a fair amount of memory (why?)

Virtual memory divides memory into pages

A page may be in

- Memory
- Memory Cache
- Disk Cache
- Disk

Access to a page is faster if it is in memory

Last thread that completed is likely to be in memory or cache

Reusing last thread that complete can improve performance

## Which Should I use?

Which method to use?

Which values (number of threads, etc) to use?

Depends on your

- Application
- Implementation
- Hardware
- Performance requirements

## How to reuse a Thread?

Classic idea

Server places client requests in a queue

Worker repeats forever

- Read request from queue
- Process request

Queue

- Block on read if queue is empty
- Signals waiting threads when data is added

## VisualWorks Example

```
Smalltalk defineClass: #ThreadedDateServer
  superclass: #{Core.Object}
  indexedType: #none
  private: false
  instanceVariableNames: 'serverSocket workers workQueue '
  classInstanceVariableNames: ''
  imports: ''
  category: 'SimpleServer'
```

### ThreadedDateServer class methods

```
port: anInteger
  ^super new setPort: anInteger
```

### ThreadedDateServer instance methods

```
setPort: anInteger
  serverSocket := SocketAccessor newTCPserverAtPort: anInteger.
  serverSocket
    listenFor: 4;
    soReuseaddr: true.
  workQueue := SharedQueue new.
  workers := OrderedCollection new.
  5 timesRepeat: [workers add: self createWorker]
```

### createWorker

```
^[self processRequestOn: workQueue next] repeat] fork
```

## Example Continued

processRequestOn: aSocket

  | clientRequest clientIOStream |

  [(aSocket readWaitWithTimeoutMs: 10000) ifTrue: [^nil].

  clientIOStream := aSocket readAppendStream.

  clientIOStream lineEndTransparent.

  clientRequest := clientIOStream through: Character cr.

  (clientRequest startsWith: 'date')

    ifTrue:

      [clientIOStream

        nextPutAll: Time dateAndTimeNow printString;

        cr;

        nextPut: Character lf;

        commit].

  clientIOStream close]

  ensure: [aSocket close]

start

  [serverSocket isActive] whileTrue:

    [workQueue nextPut: serverSocket accept.

    workQueue size > 2 ifTrue: [workers add: self createWorker]]

shutdown

  serverSocket close

## **Issues not addressed**

How big to make the thread pool

Setting priority of threads

Using a thread for the server

Maximum number of threads to allow

Making values configuration settings

How to reduce the number of threads

Separating the Server from parsing the protocol

## Java Example

### SharedQueue

```
import java.util.ArrayList;
public class SharedQueue
{
    ArrayList elements = new ArrayList();

    public synchronized void append( Object item )
    {
        elements.add( item);
        notify();
    }

    public synchronized Object get( )
    {
        try
        {
            while ( elements.isEmpty() )
                wait();
        }
        catch (InterruptedException threadIsDone )
        {
            return null;
        }
        return elements.remove( 0);
    }

    public int size()
    {
        return elements.size();
    }
}
```

## DateHandler

```
import java.net.*;
import java.io.*;
import java.util.Date;

public class DateHandler extends Thread
{
    SharedQueue workQueue;

    public DateHandler(SharedQueue workSource )
    {
        workQueue = workSource;
    }

    public void run()
    {
        while (!isInterrupted() )
            try
            {
                Socket client = (Socket) workQueue.get();
                processRequest(client);
            }
            catch (Exception error )
            {
                /* log error*/
            }
    }
}
```

## DateHandler Continued

void processRequest(Socket client) throws IOException

```
{
  try
  {
    client.setSoTimeout( 10 * 1000 );
    processRequest(
      client.getInputStream(),
      client.getOutputStream());
  }
  finally
  {
    client.close();
  }
}
```

void processRequest(InputStream in, OutputStream out)  
throws IOException

```
{
  BufferedReader parsedInput =
    new BufferedReader(new InputStreamReader(in));
  PrintWriter parsedOutput = new PrintWriter(out,true);
  String inputLine = parsedInput.readLine();
  if (inputLine.startsWith("date"))
  {
    Date now = new Date();
    parsedOutput.println(now.toString());
  }
}
```

## DateServer

```
import java.util.*;
import java.net.*;
import java.io.*;

public class DateServer
{
    SharedQueue workQueue;
    ServerSocket listeningSocket;
    ArrayList workers = new ArrayList();

    public static void main( String[] args )
    {
        System.out.println( "Starting" );
        new DateServer( 33333 ).run();
    }
}
```

## DateServer Continued

```
public DateServer( int port )
{
  try
  {
    listeningSocket = new ServerSocket(port);
    workQueue = new SharedQueue();
    for (int k = 0; k < 5; k++)
    {
      Thread worker = new DateHandler( workQueue);
      worker.start();
      workers.add( worker);
    }
  }
  catch (IOException socketCreateError)
  {
    //log and exit here
  }
}
```

## DateServer Continued

```
public void run()
{
    Socket client = null;
    while (true)
    {
        try
        {
            client = listeningSocket.accept();
            workQueue.append( client);
        }
        catch (IOException acceptError)
        {
            // need to log error and make sure client is closed
        }
    }
}
```

## **How to Remove Extra Workers**

Let each worker above the minimum end after K requests

When system is idle have reaper thread remove some workers

Have get() on shared queue time out

When workers get time out on the queue, they exit