

CE 4333 Practical Computational Hydraulics using R¹

by

Theodore G. Cleveland

Department of Civil, Environmental, and Construction Engineering
Texas Tech University



¹The preferred programming tool of pirates, everywhere! Rrrrrr ...

Contents

1	Introduction and Getting Started	7
1.1	About R	7
1.2	Getting Started	8
1.2.1	Windows Users	8
1.2.2	Macintosh OSX Users	16
1.2.3	Linux Users	23
2	Algorithms and R for Computing	31
2.1	Tools	31
2.2	Programming	32
2.3	Interpolating Tabular Data – A useful algorithm	32
2.4	Linear Interpolation	32
2.4.1	Interpolation of Values in Two Pairs	34
2.4.2	Interpolation of Values in Two Arrays	35
2.5	Sorting	38
2.5.1	Bubble Sort	38
2.5.2	Insertion Sort	43
2.5.3	Merge Sort	45
2.5.4	Built-In R Sorts	45
3	Numerical Methods – Integrals, Derivatives, and Newton’s Method	49
3.1	Numerical Integration of Functions	49
3.1.1	Rectangular Panels	50
3.1.2	Trapezoidal Panels	52
3.1.3	Parabolic Panels	55
3.2	Exercise Set 2	58
3.3	Numerical Integration of Tabular Data	59
3.3.1	Reading from a file – open, read, close files	59
3.4	Integrating tabular data	60
3.5	Numerical Differentiation	64
3.6	Difference Approximations for Tabulated Data	64
3.6.1	Slope of a Secant Line	65
3.6.2	Disaggregation	66
3.6.3	Numerical Differentiation	67
3.6.4	Aggregation	69
3.7	Exercises	69
3.8	Finite-Difference Formulas	70
3.8.1	First Derivatives	70
3.8.2	Second Derivatives	70
3.8.3	Third Derivatives	71
3.9	Single Variable Quasi-Newton Methods	72
3.9.1	Newton’s Method — Using analytical derivatives	73

3.9.2	Newton's Method — Using Finite-Differences to estimate derivatives	76
3.9.3	Method Fails	79
3.9.4	Multiple Roots	80
3.9.5	Cycling	80
3.9.6	Near-zero derivatives	81
3.10	Related Concepts	81
3.11	Exercise Set 3	82
4	Tank Drain Simulation using Finite-Differences	83
4.1	Problem Description	83
4.2	Computational Approach	84
4.3	Problem Analysis — Development of the Analytical Solution	84
4.3.1	Application of Analytical Solution in \mathbf{R}	86
4.4	Problem Analysis — Development of the Finite-Difference Approximation	86
4.5	Application of the Finite-Difference Approximation in \mathbf{R}	87
4.6	Exercise Set 4	89
5	Simultaneous Linear Systems of Equations	91
5.1	Numerical Linear Algebra – Matrix Manipulation	92
5.2	The Matrix — A data structure	92
5.3	Matrix Arithmetic	93
5.3.1	Matrix Definition	94
5.3.2	Multiply a matrix by a scalar	94
5.3.3	Matrix addition (and subtraction)	96
5.3.4	Multiply a matrix	98
5.3.5	Identity matrix	99
5.3.6	Matrix Inverse	101
5.3.7	Gauss-Jordan method of finding \mathbf{A}^{-1}	101
5.4	Jacobi Iteration – An iterative method to find solutions	106
6	Simultaneous Non-Linear Systems of Equations	109
7	Numerical Methods – Multiple Variable Quasi-Newton Method	115
8	Pipelines and Networks	120
8.1	Pipe Networks – Topology	120
8.1.1	Continuity (at a node)	120
8.1.2	Energy Loss (along a link)	121
8.1.3	Velocity Head	122
8.1.4	Added Head — Pumps	122
8.1.5	Extracted Head — Turbines	123
8.2	Pipe Head Loss Models	123
8.3	Pipe Networks Solution Methods	124

8.4	Network Analysis	124
9	Pipelines Network Analysis	128
9.1	Script Structure	128
9.1.1	Support Functions	129
9.1.2	Augmented and Jacobian Matrices	130
9.1.3	Stopping Criteria, and Solution Report	131
9.2	Exercises	136
10	Pumps and Valves	139
10.1	Energy Loss (along a link)	139
10.1.1	Added Head — Pumps	139
10.1.2	Fitting (Minor) Losses	144
10.1.3	Extracted Head — Turbines	145
11	Pipeline Transients — Water Hammer	146
11.1	Analysis	146
11.2	Explicit Finite-Difference Model	147
11.3	An Example R Implementation	149
12	Open Channel Flow	153
12.1	St. Venant Equations	153
12.1.1	The Computational Cell	153
12.1.2	Assumptions	154
12.1.3	Conservation of Mass	155
12.1.4	Conservation of Momentum	156
13	Steady Water Surface Profiles	163
13.1	Steady Uniform Flow	163
13.2	Steady Gradually Varied Flow	163
13.2.1	Finite Difference Methods — Fixed Depth Change Step Method	163
13.2.2	Coding the algorithm in R	164
13.2.3	Example 1 — Backwater curve	166
13.2.4	Example 2 — Front-water curve	168
13.2.5	Finite Difference Methods — Fixed Space Change Method . .	170
13.3	Exercise Set 10	172
14	Finite-Difference Method for Unsteady Open Channel Flow	174
14.1	Governing Difference Equations – Lax Scheme	174
14.1.1	Continuity	174
14.1.2	Momentum	174
14.2	Mapping from the Physical to Computational Domain	175
14.3	Building the Difference Equations	177
14.3.1	Time Differences	177
14.3.2	Space Differences	177
14.3.3	Continuity	178

14.3.4	Momentum	179
14.3.5	Example 1 – Steady flow in a rectangular channel	179
14.3.6	Example 2 – Routing a storm hydrograph	185
14.3.7	Example 3 – Sudden gate closing in an aqueduct channel	185
14.3.8	Example 4 – Long waves in a tidal-influenced channel	185
14.4	Quasi-2D Methods	185
14.5	Exercises	185
15	Using Method of Characteristics for Boundary Conditions	186
16	Flow in Porous Materials	187
16.1	Storage	188
16.2	Permeability	188
16.3	Head Loss Models	188
16.4	Confined Aquifer Flow	190
16.5	Unconfined Aquifer Flow	193
17	Steady Groundwater Flow	195
17.1	Confined Aquifer Flow	195
17.1.1	Finite-Difference Methods – 1 Spatial Dimension	195
17.1.2	Finite-Difference Methods – 2 Spatial Dimensions	203
17.1.3	Generalizing the Boundary Conditions	214
17.1.4	Internal sources/sinks to handle a discharge (well) at a specific location	222
17.2	Unconfined Aquifer Flow	228
17.2.1	Finite-Difference Methods	228
17.3	Exercises	235
18	Unsteady Groundwater Flow	237
18.1	Confined Aquifers	237
18.1.1	Finite-Difference Methods – 1 Spatial Dimension	237
18.1.2	Explicit Formulation	238
18.1.3	Implicit Formulation	238
18.2	Finite-Difference Methods – 2 Spatial Dimensions	238
18.3	Unconfined Aquifers	246
18.3.1	Explicit Formulation	247
18.3.2	Implicit Formulation	248
18.4	Exercises	258
19	Flow Nets	261
19.1	Unconfined Aquifer Flow	274
19.1.1	Finite-Difference Methods	274
19.2	Exercises	274
20	Heat and Mass Transport	275
20.1	Concentration	275

20.2	Tracer Hypothesis	275
20.3	Advection (Convection) Process	275
20.4	Diffusion Process	275
20.5	Dispersion Process	275
20.6	Reaction Processes	275
20.6.1	Linear Equilibrium Isotherms and Concept of Retardation . . .	275
20.7	Advection, Dispersion, Retardation, Decay Mathematics	275
21	Analytical Solutions to ADR Equations	276
21.1	Impulse Solution	276
21.2	Ogata-Banks Solution	276
21.3	2D-Injection (Hunt) Solution	276
21.4	2D-, and 3D- Domenico Robbins Solutions	276
22	Numerical Solutions to ADR Equations	277
22.1	Explicit Methods	277
22.1.1	Centered Differences	277
22.1.2	Upwind Formulation(s)	277
22.2	Crank-Nicholson Methods	277

1 Introduction and Getting Started

In the 1990s Civil Engineering programs reduced programming courses in a effort to recover hours for other topics – a logical decision at the time, but with some consequences. The philosophy was that engineers would not need to be able to write computer programs, but instead just use them. Microsoft Excel and Lotus 1-2-3 were the dominant spreadsheet software programs (Borland QuatroPro was a close third), and with macro instruction capability, much legitimate engineering computation could be performed within these tools. In fact I developed Excel spreadsheets that could solve multi-dimensional diffusion problems (3D groundwater flow) using fully implicit finite difference methods. These spreadsheets were slow relative to MODFLOW, but you could watch the solutions evolve – ultimately the process was deemed a waste, because of the ever present “... there is no longer a need for engineers to be able to write programs.” Later on I developed spreadsheets to perform pressurized pipe network simulation, gradually varied flow simulation, and rudimentary water-hammer and St-Venant equation solutions. The spreadsheets were never really practical (yes they worked well, produced the same results as professional products, but were always intended a pedagogical tools), but they proved an important point – if you could teach a computer to follow an algorithm it made you a more self-help user of the professional tools.

In 2014 several of my students expressed desire to understand programming – they all know how to write code, but feel they don’t know how to build algorithms (and implement them). This workbook is an attempt to remedy that student self-identified weakness. I conducted several one-to-one classes (as special topics); they learned a lot, I learned even more. This book is a tribute to their interests.

The workbook plan is to introduce a programming tool – I have selected **R** because it has a rich development environment already available, graphics is almost trivial, then apply that tool to selected hydraulics problems of practical value. In the end the reader ends up with a toolkit that can either stand-alone, or more likely supplement professional tools they will eventually use.

R is freeware, but it is built and maintained by a consortium of programmers and statisticians. They have evolved the environment to work on most of the main architectures (MacOS, Windows, Linux); there are even parallel processor and GPU builds available, and a company called RStudio provides the APIs to even run it server side. Much of the underlying code is C, C++, and well proven FORTRAN routines.

1.1 About R

R is a open source environment that runs on Windows, Linux/UNIX, and Mac OS X. The individual binaries are unique to each OS and architecture, but **R** “source” is interchangeable among machines. With very minor differences, an **R** script will run equally well on any machine.

R is a statistical analysis tool, it is also a programming tool and language, it is also a nearly “publication” quality graphics tool. Naturally all this capability comes at a cost (especially since the software is distributed for “free”) — learning to do more than simple calculations takes some time (not much), but the skill is highly perishable. You will need to keep notes, or copies of your **R** scripts for future reference. Relearning after some time away from **R** is pretty simple, so the modeler only has to pay the steep learning cost once.

The remainder of this essay shows how to obtain and install **R** on a Windows machine. Macintosh and Linux installs are accomplished in a similar fashion. For the truly insane, the entire environment can be built from source on any machine with PERL, gcc, and gfortran compilers (default in Linux, easy to obtain for other architectures).

1.2 Getting Started

The first step required (for using **R** as a programming tool) is to install **R** on your computer. The source of the binary builds is the same regardless of the underlying operating system – the Comprehensive R Archive Network (CRAN for short). The remainder of this chapter shows how to get the tool running on the three main operating systems in current practice.

1.2.1 Windows Users

The purpose of this section is to demonstrate how to get **R** running on a Windows computer. This document assumes the following:

1. You have internet connection.
2. You have sufficient user privileges to install software on your machine. (If you need someone else to install, I did my install by running the installer as a local administrator — obviously you need the password)
3. You have 60MB or so of vacant disk space on the system directory.

The step-by-step guide is presented as a series of screen captures. Obviously adjust inputs to fit your machine. The version in these screen captures is quite dated — use the most recent, stable version offered on CRAN (Comprehensive R Archive Network).²

²I have updated the screen captures for Windows 10 — so these should replicate the steps.




Figure 1. Google "R" (alternatively google CRAN).

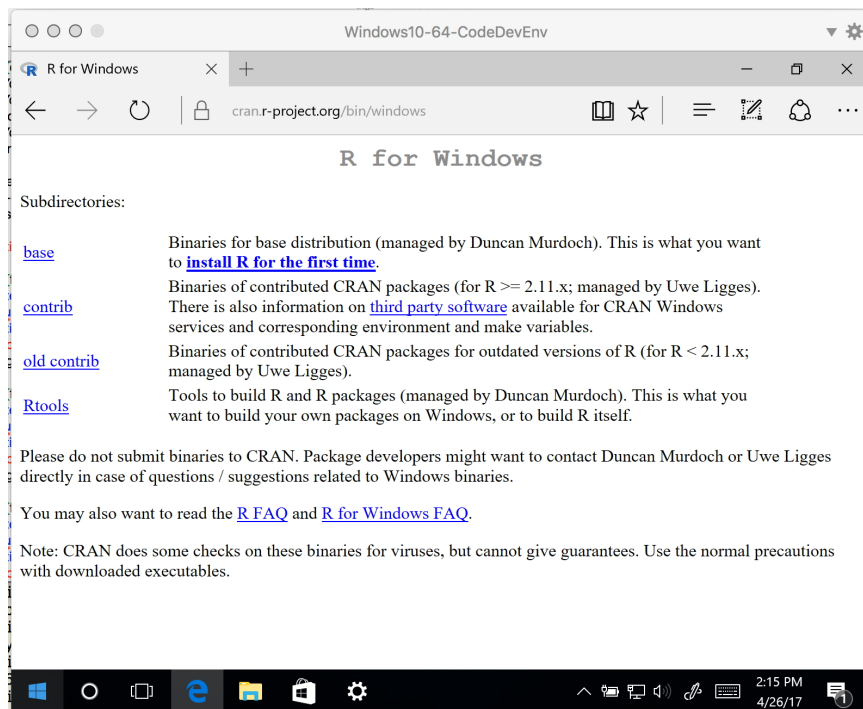


Figure 2. Taking the "Windows Link".




Figure 3. Choose “Install R for the First Time” – goes to the download page. We will next select “Download R ...” and it should download a windows installer. Choose “save” when prompted..




Figure 4. Download arrived. Now run the installer (you need install privileges – if its your personal laptop, then your regular user account should work)..

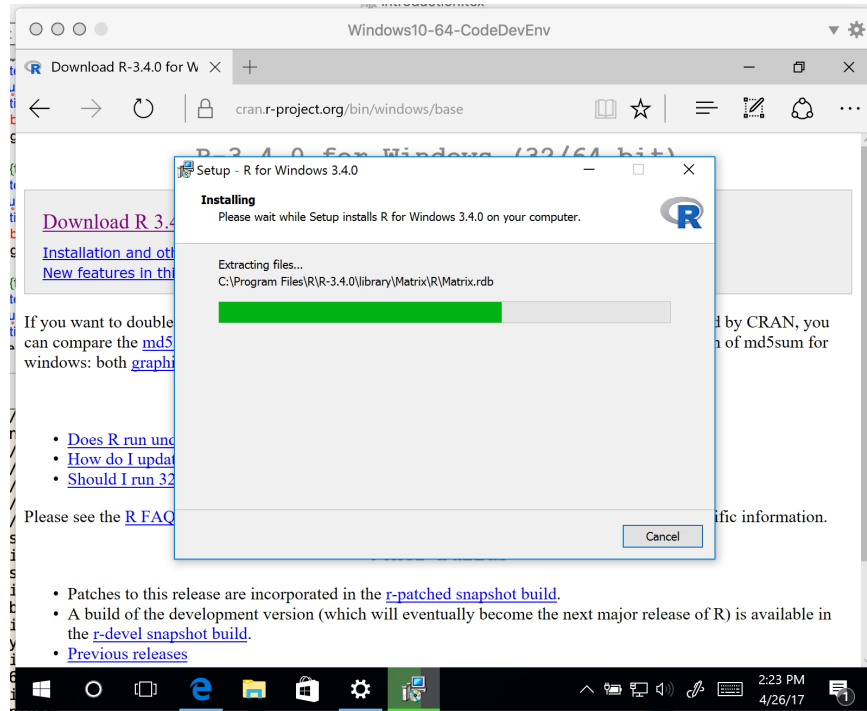


Figure 5. Installer run in progress. Accept the defaults – its just easier and works. Later you can re-install elsewhere in your filesystem..

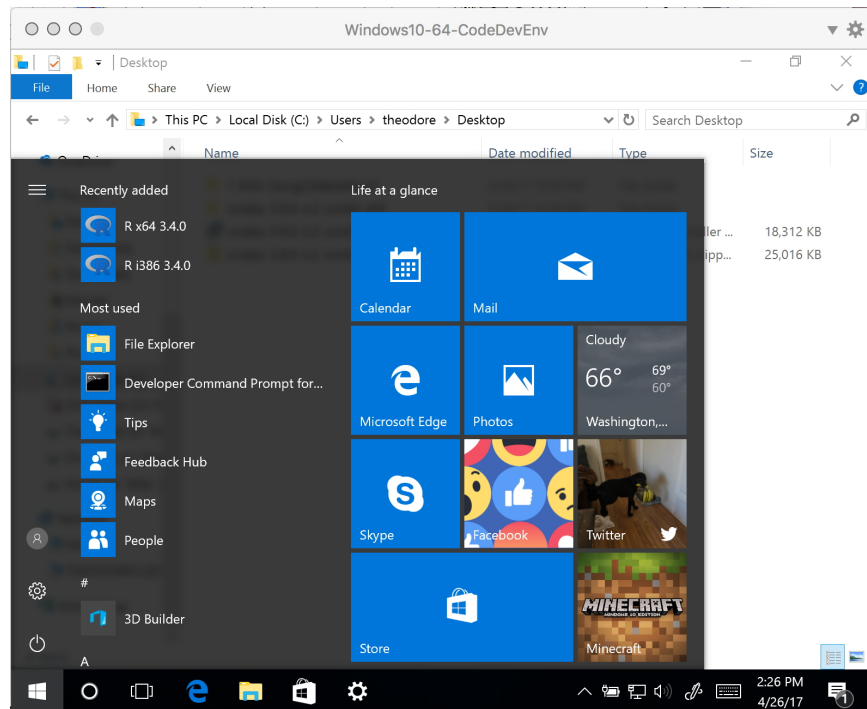


Figure 6. Installed “R base” packages. Notice it installs 32-bit and 64-bit versions. The next step is to verify the install by trying to run the program..

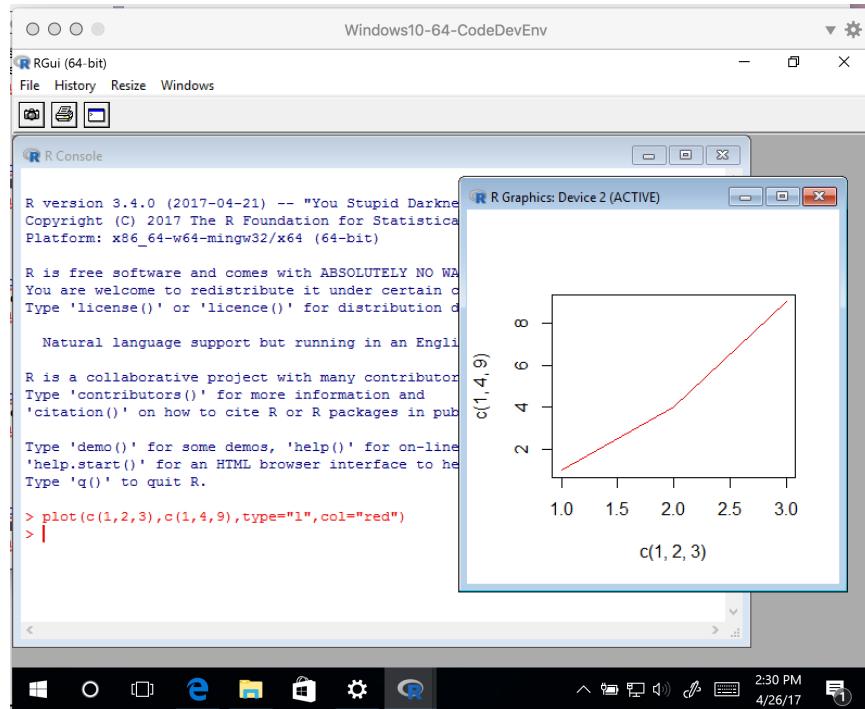


Figure 7. Run the program. Type `plot(c(1,2,3),c(1,4,9),type="l",col="red")` into the console window. A plot should be generated as shown. If this works, then your install is good and now we install **R Studio**.

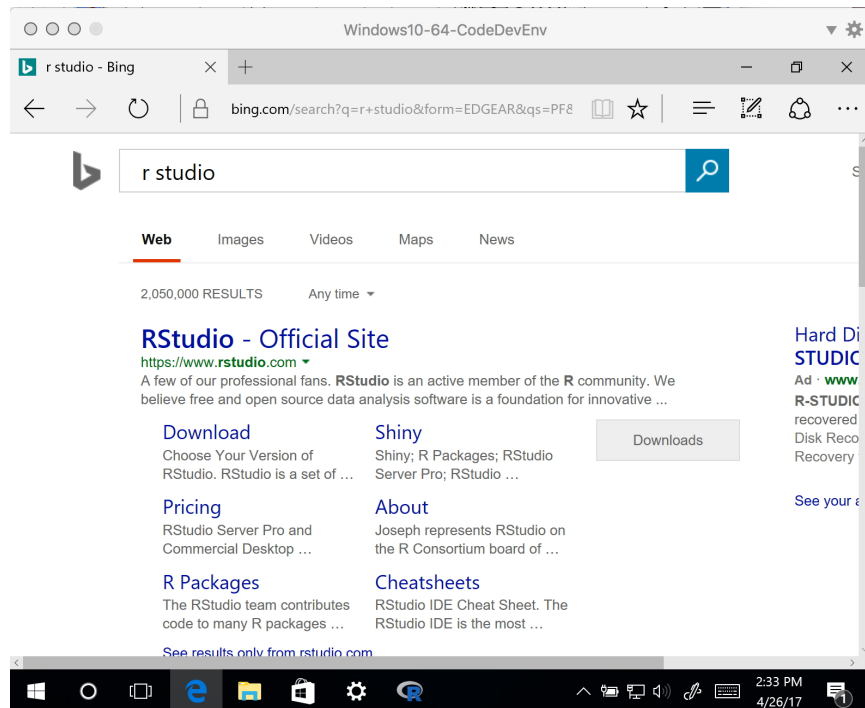


Figure 8. Search for **R Studio**. Choose the Download link..

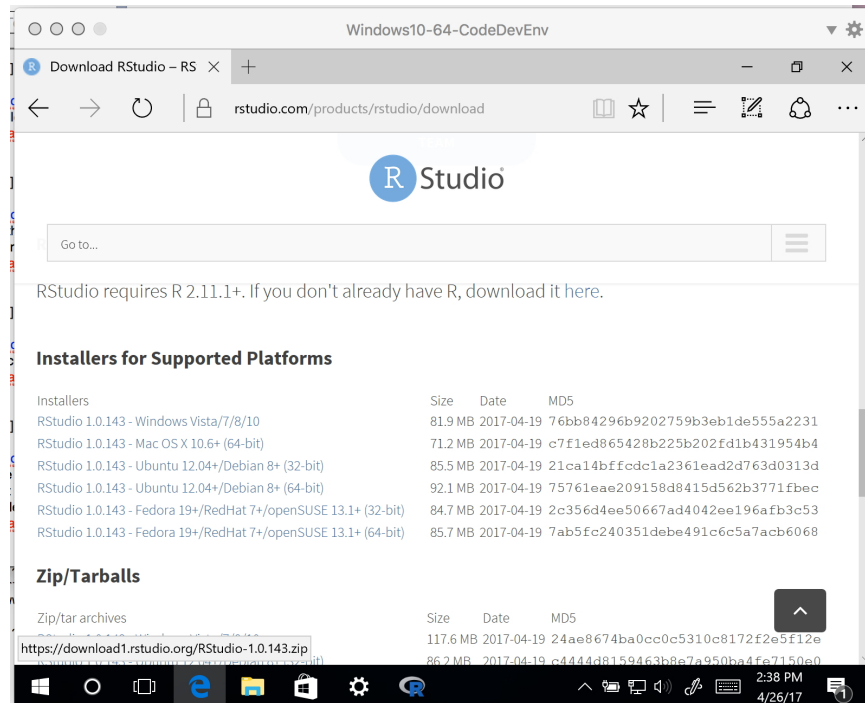


Figure 9. In the download link scroll down to the repository. You want to install the FREE Desktop version. If on a windows machine, it is the top most of the installers. Don't accidentally download the Zip/Tarballs – all that is source code and without the compilers you cannot make much use of it. Building from source is a challenge. Choose the windows installer and download, select "save" when prompted.

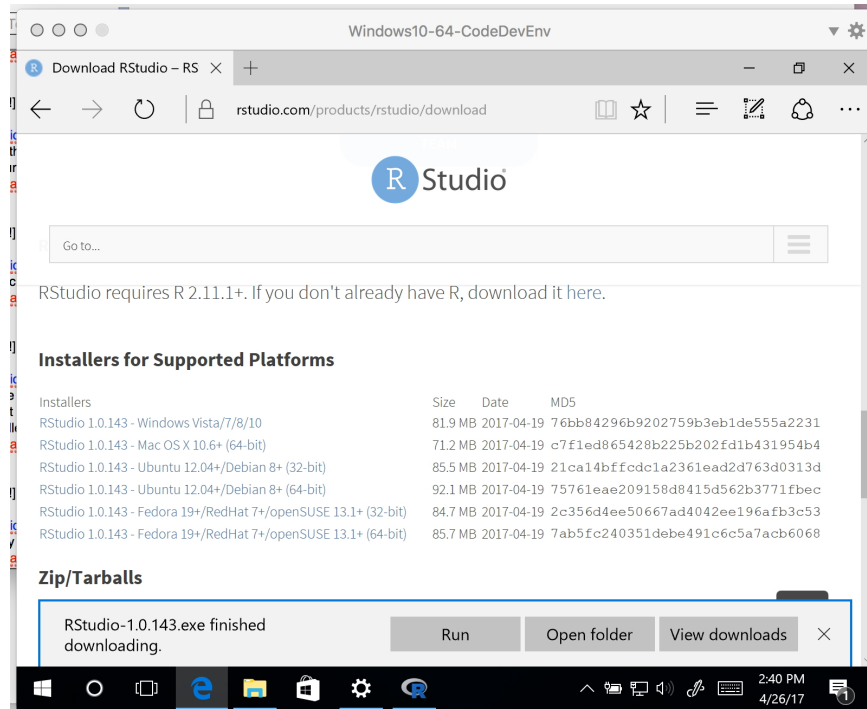


Figure 10. Download arrived, run the executable (it should be a .exe file). Accept the defaults during the install.

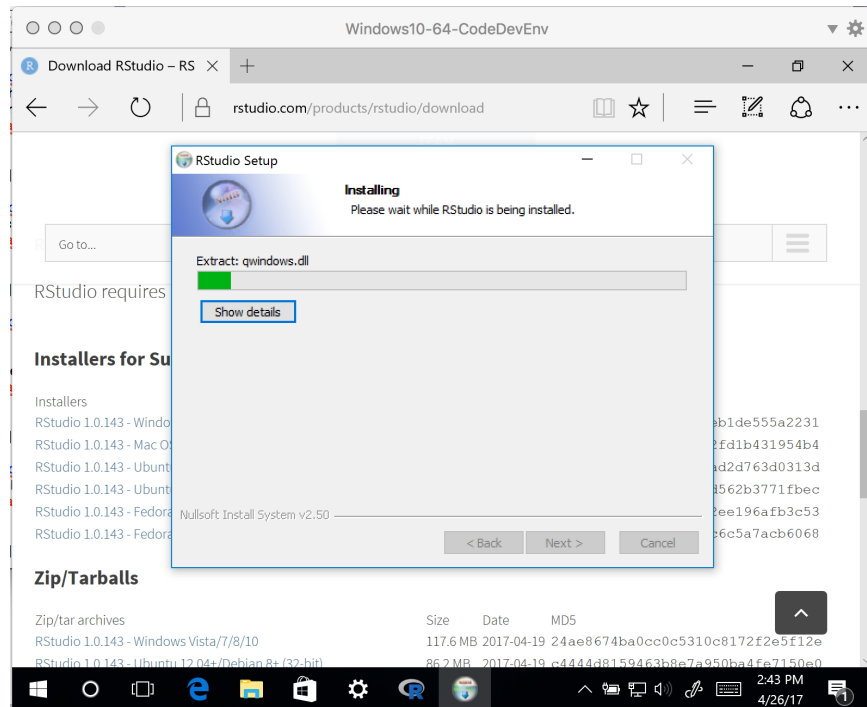


Figure 11. Installer in-progress. When it completes, you should now have **R Studio** and **R** installed. We will test the **R Studio** install using the same simple plot call.

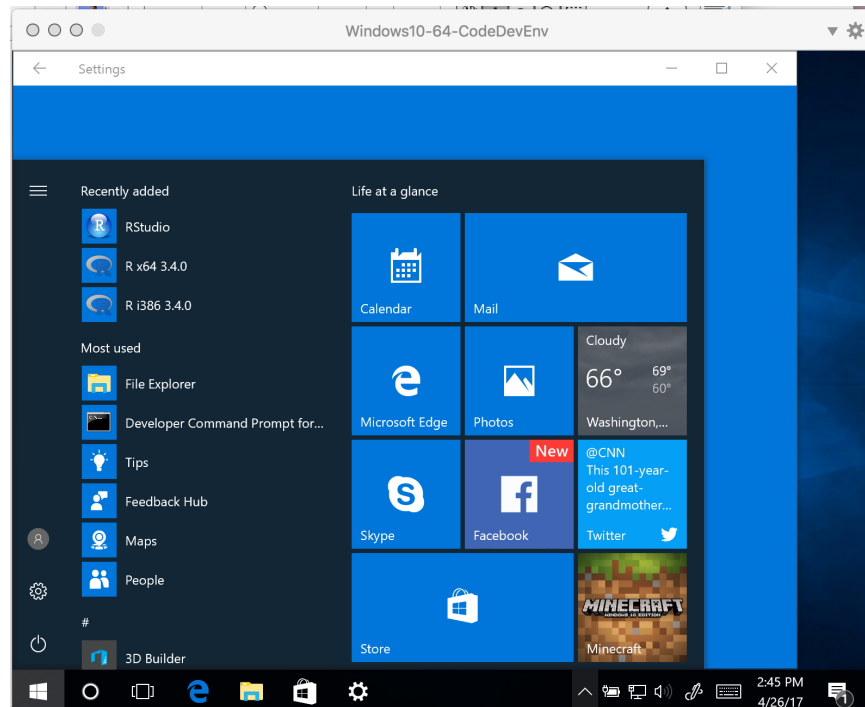


Figure 12. Here we see the program is installed, now run R Studio to verify the install.

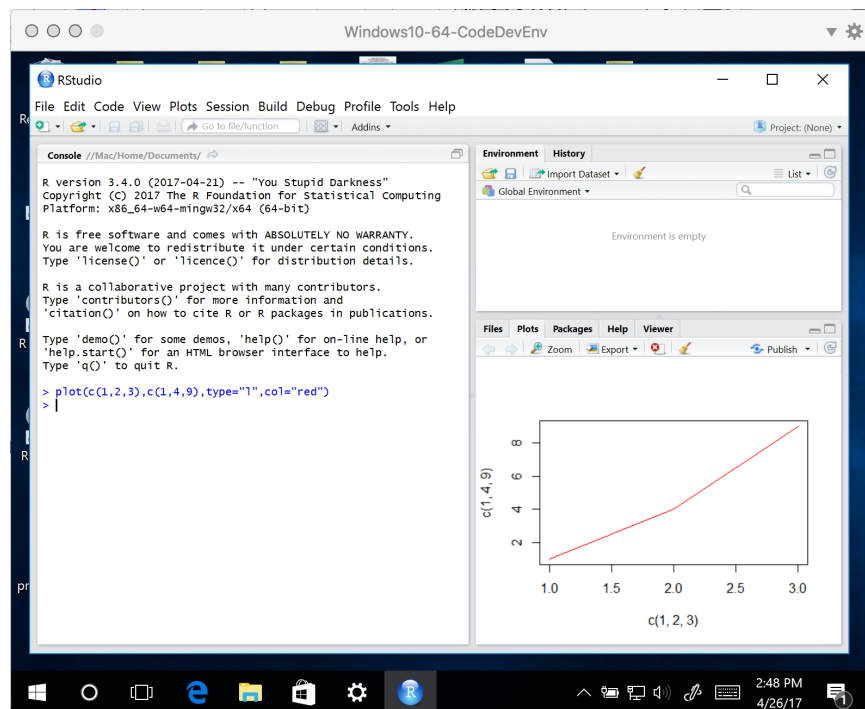


Figure 13. Type `plot(c(1,2,3),c(1,4,9),type="l",col="red")` into the console window. A plot should be generated as shown.

Yipee! It is running. You can install additional packages now or later. You should now have sufficient computation capability for the course.

1.2.2 Macintosh OSX Users

[Replicate Windows using MacOS screen captures] The purpose of this section is to demonstrate how to get **R** running on a Macintosh computer.³ This document assumes the following:

1. You have internet connection.
2. You have sufficient user privileges to install software on your machine. (If it is your personal machine, the install may request your password, but should install.)
3. You have 60MB or so of vacant disk space on the system directory.

The step-by-step guide is presented as a series of screen captures. Obviously adjust inputs to fit your machine. The version in these screen captures is quite dated — use the most recent, stable version offered on CRAN (Comprehensive R Archive Network).⁴

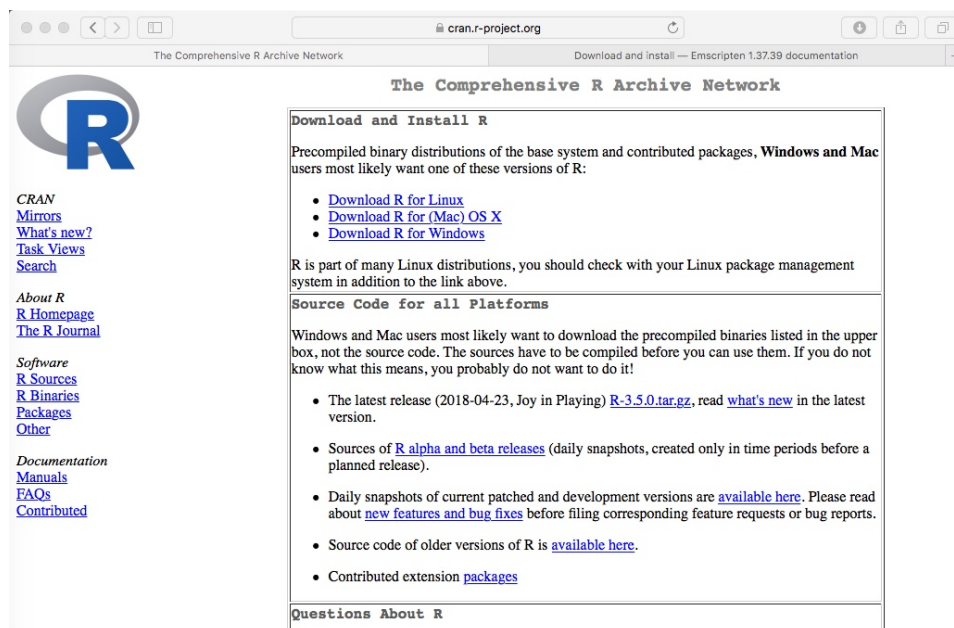


Figure 14. Google "R" (alternatively google CRAN).

³I assume no-one will be using a PPC-Based Mac. If so, the CRAN does have PPC builds of R, but R Studio is not available; you would have to build from the source code.

⁴I have updated the screen captures for Windows 10 — so these should replicate the steps.

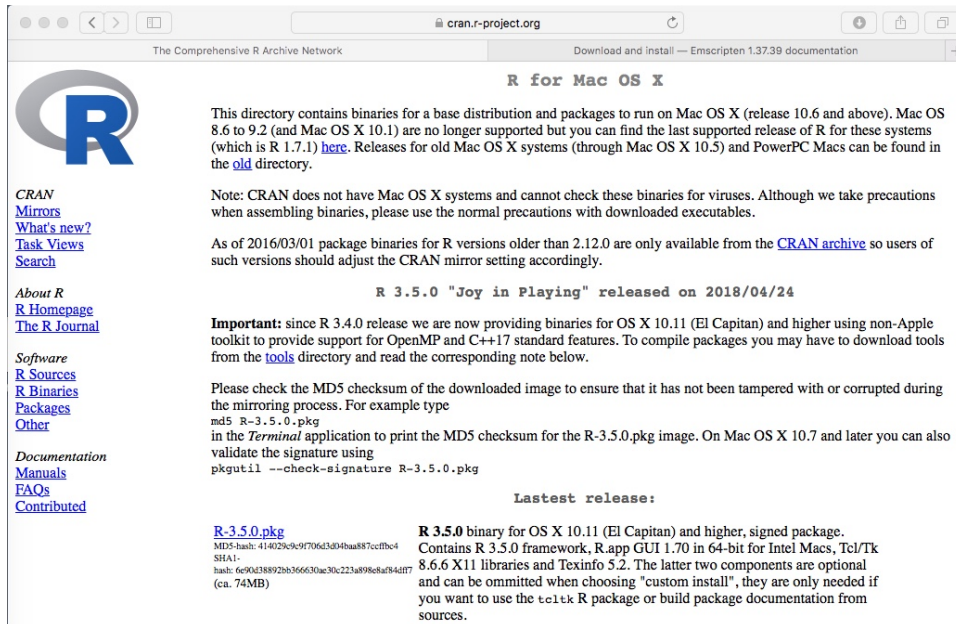


Figure 15. Select MacOS operating system link.

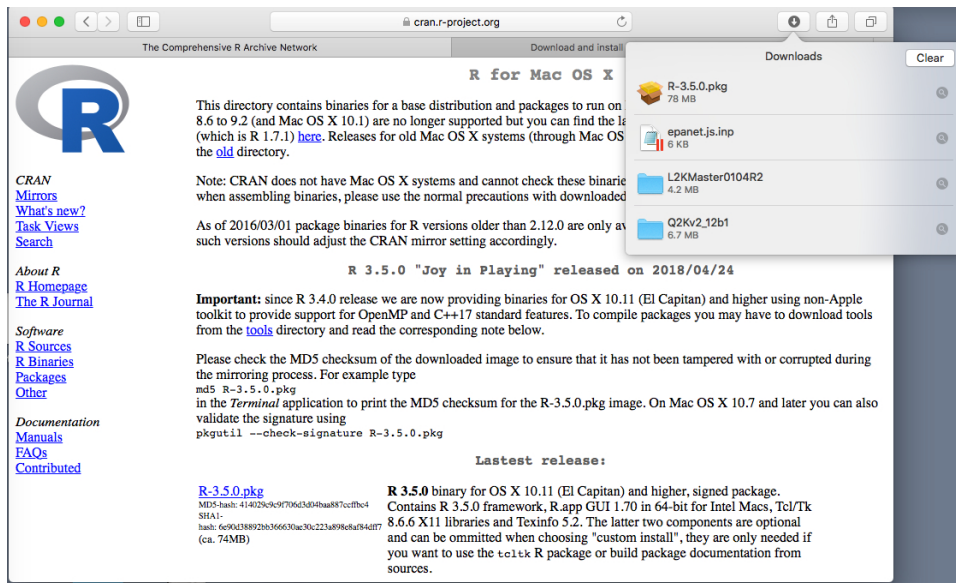


Figure 16. Download the R Installer.

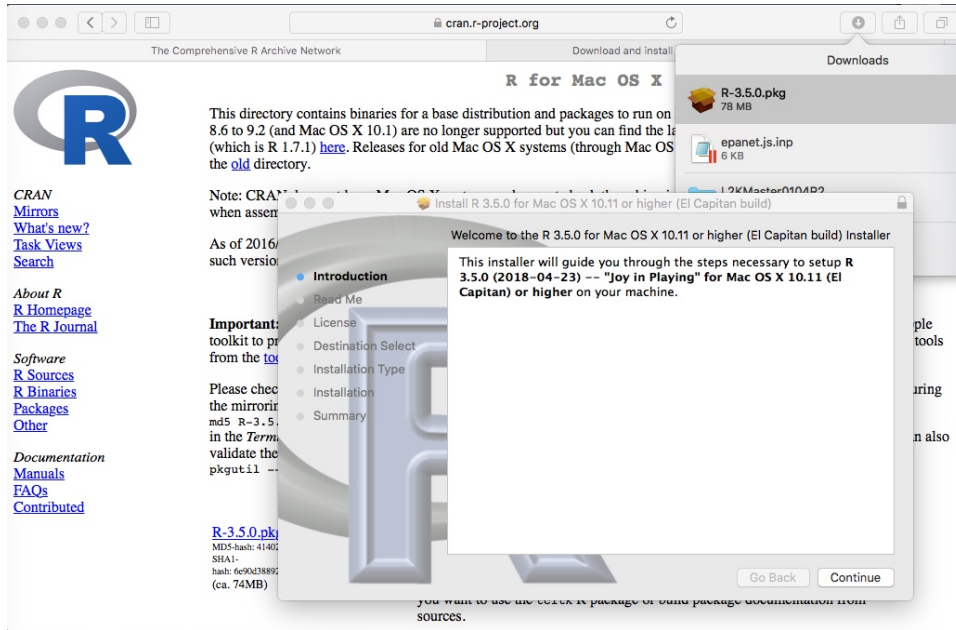


Figure 17. Run the installer, accept the defaults.




Figure 18. Successful install.

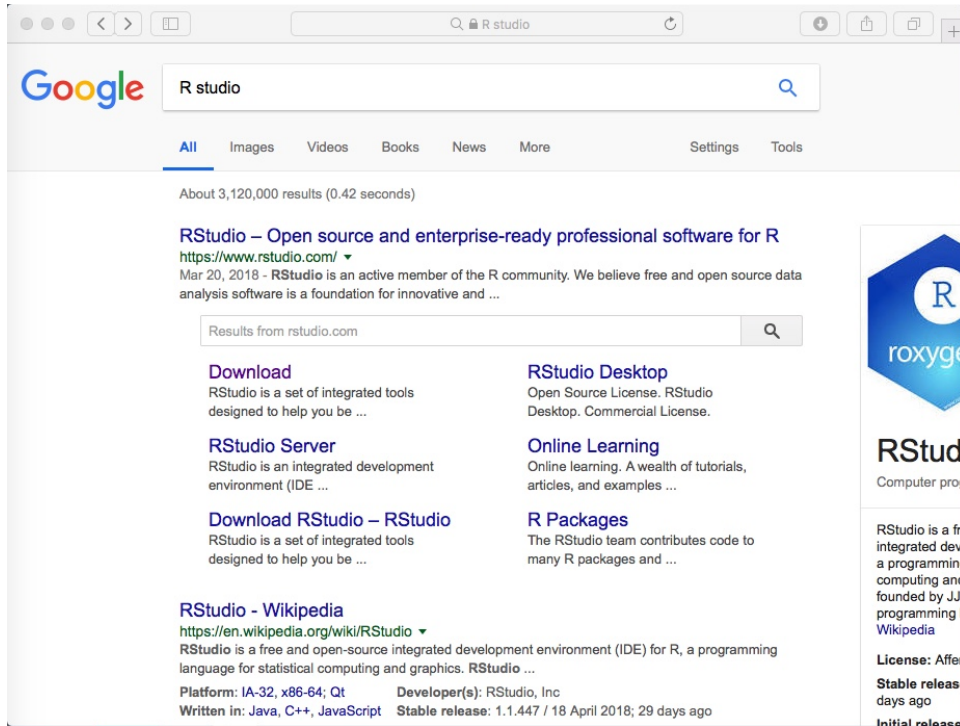


Figure 19. Google R Studio.

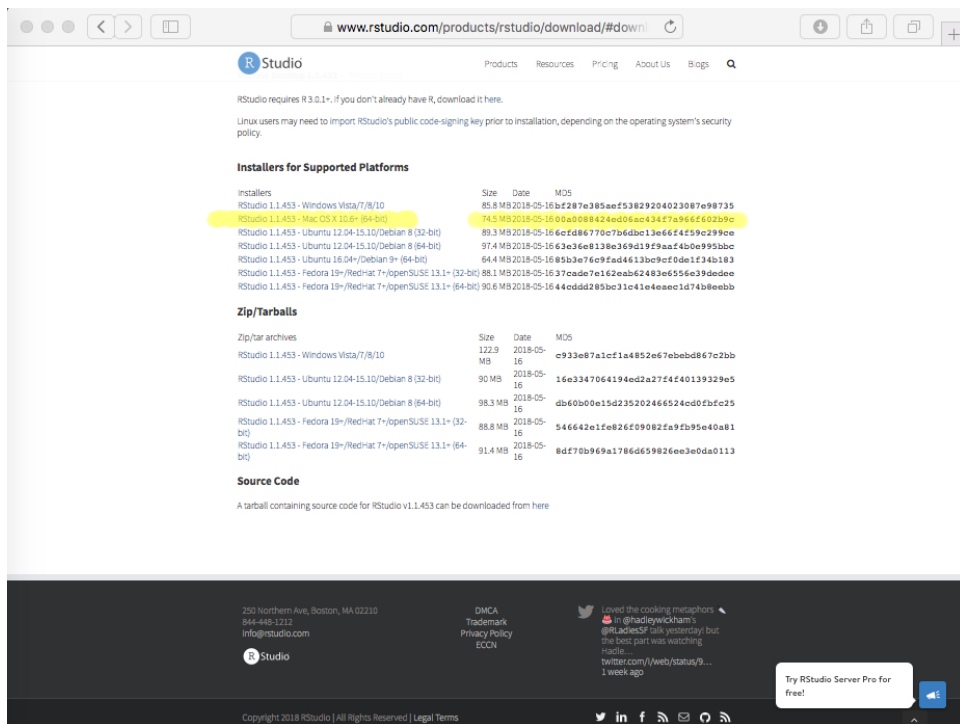


Figure 20. Select R Studio installer (for MacOS).

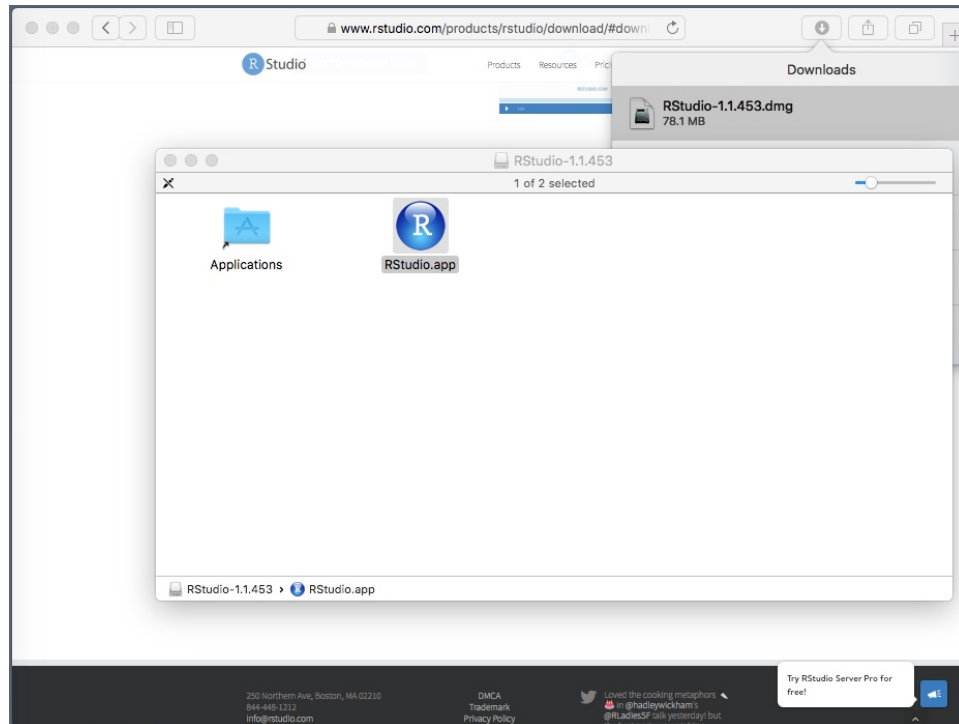


Figure 21. Mount the disk image. Then drag the R Studio icon on top of the Applications link – this action will install the program.

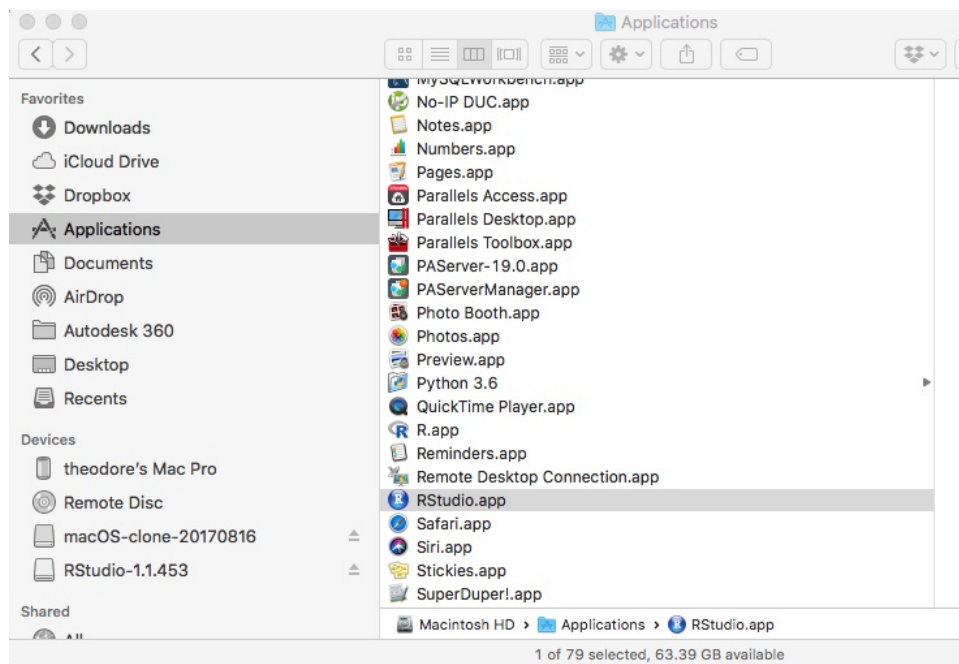


Figure 22. Start R Studio (in Applications directory, double click on the icon).

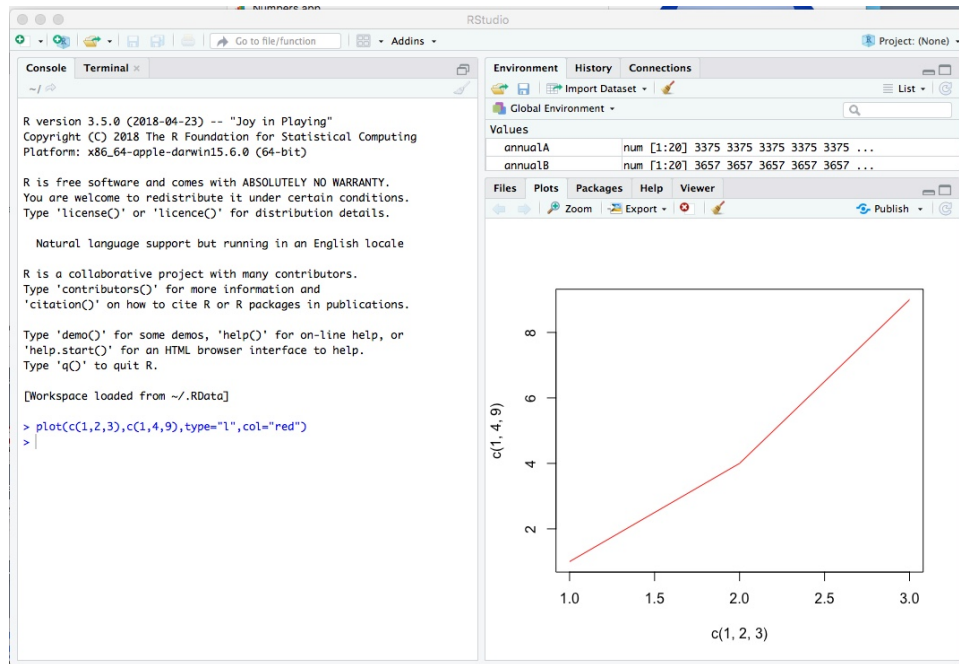


Figure 23. Type in some R script and viola – it produces a simple plot.

Yipee! It is running. You can install additional packages now or later. You should now have sufficient computation capability for the course.

1.2.3 Linux Users

The purpose of this section is to demonstrate how to get **R** and **R Studio** running on a Linux computer. This document assumes the following:

1. You have internet connection.
2. You have sufficient user privileges to install software on your machine. Generally, if it is your own machine then you have superuser (root) privileges. If it is some network machine maintained by someone else you probably don't have such privileges. The examples here use the `sudo <command>` to do the installs – on my machine the password I enter is my user password (and not the root password). Alternativley you can switch user `su` to root, and run the installs as root – this approach is considerably more dangerous in terms of wrecking your operating system that using the `sudo` approach.
3. You have 60MB or so of vacant disk space on the system directory.

The step-by-step guide is presented as a series of screen captures. Obviously adjust inputs to fit your machine. Use the most recent, stable version offered on CRAN (Comprehensive R Archive Network).




Figure 24. Terminal Prompt in Linux.

To get started we need to install **R**. The easiest method is to use the package manager – my Linux distribution is Ubuntu 16.XX. It is built from the debian distribution and uses the `apt` package manager. The package manager is pretty cool because when we request a package it finds the package and its dependencies, downloads everything needed and then we can install. In earlier times (the 1990's) using the Red Hat Package Manager (`rpm`) one would have to find the dependencies themselves (in all

fairness `rpm` would identify the dependencies and suggest where to find them!). So here we go, first open a terminal in Linux.

Next in the terminal window issue the command `sudo apt install r-cran-littler`.




Figure 25. Install R using the apt package manager (or rpm if using a red hat variant).

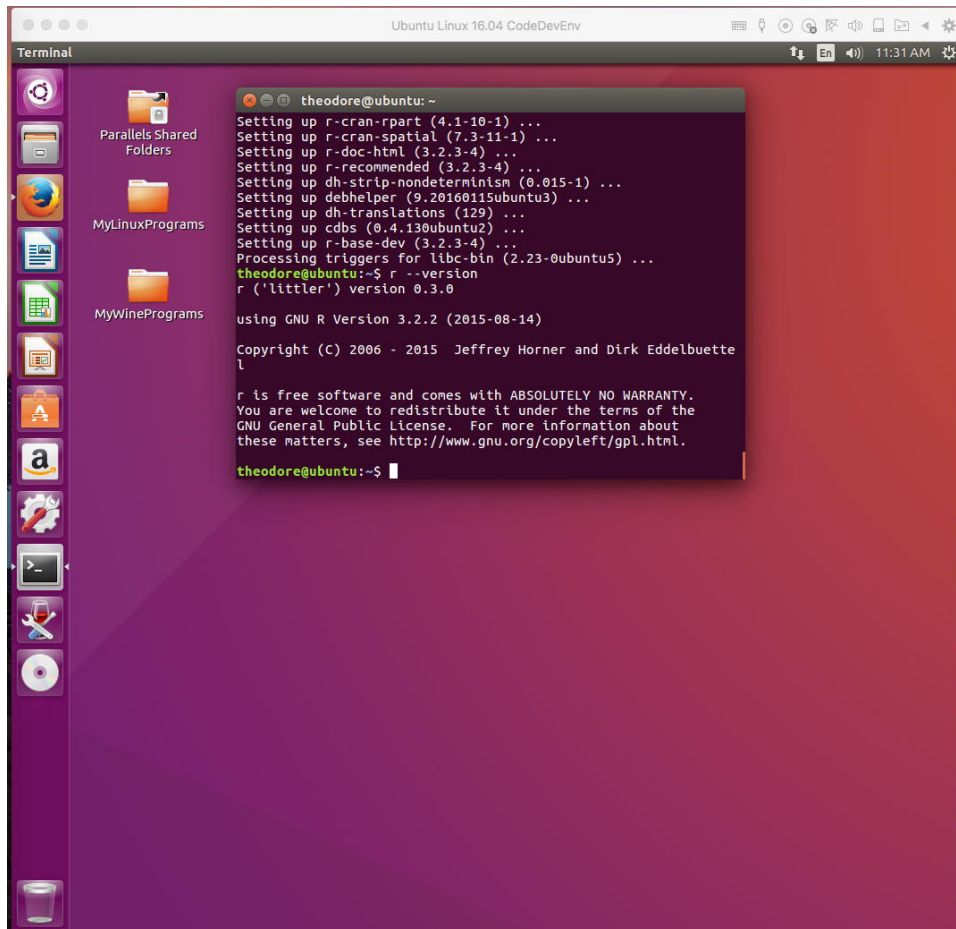
When you press return, the computer should ask for a password – use your user password; if you have install privileges this will work. If not, you will have to switch user to root and either add your user account to the group `wheel`, return to your account – or just install as root.

Next the install will begin and may take a few minutes. Usually there is a lot of installer messages that run across the screen (kind of like in “The Matrix”). The `apt` utility is downloading files and bringing them to resources on the system so that **R** can run. Eventually it should get to the end of the install and may look something like the next screen capture.

Our next task is to verify that the install was successful. Usually failure is obvious, but not always. I find the easiest way to verify that the operating system thinks the software is installed is issue the command to run the program with the version switch active. In this case, issue the command `r --version`. This command will try to run **R**, and will return the version number (or build number).

In the next screen capture when we run the command we see that the version number

is **R** version 3.2.2. The version numbering system is typical for all software – it identifies something like this is the second subversion, of the second revision, of the third stable release.



```

theodore@ubuntu:~$ r --version
r ('littler') version 0.3.0

using GNU R Version 3.2.2 (2015-08-14)

Copyright (C) 2006 - 2015 Jeffrey Horner and Dirk Eddebuette
l

r is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under the terms of the
GNU General Public License. For more information about
these matters, see http://www.gnu.org/copyleft/gpl.html.

theodore@ubuntu:~$

```

Figure 26. Verify install of **R** using `r --version`.

The next step I recommend is to go ahead and download a development environment. **R Studio** is an integrated development environment for **R**. We don't "need" it, but it enforces some discipline, gives us a place to store and modify **R**-scripts (little programs), and lets us see all of our work going on in one location.

To get **R Studio** we have to download it from the manufacturer, the copy we will use is free. Instead of `apt` this time I used Firefox to navigate to the **R Studio** website, then select download.

Next select the appropriate installer for your platform. Be sure you are selecting an installer and not the source codes for the program.⁵

We will download the 64-bit version (unless you have a 32-bit machine, then you need

⁵In theory we could build the program from source using the make utility and (hopefully) already installed compilers – this is for people with time, training, inclination and need. We are going to use the already built binaries!

32-bit software). We need to select our Linux platform – mine is Ubuntu/Debian. I also have a Red Hat/SuSe machine, so if I were using that machine, i would select its software.

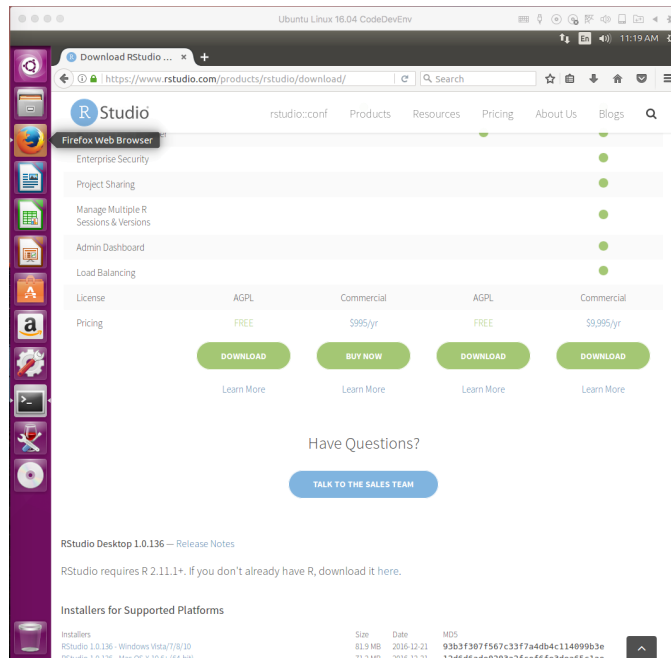


Figure 27. Browser search for R Studio; go to downloads. We are going to select the free (leftmost) column.




Figure 28. Download the version of R Studio; in this case 64-bit for Ubuntu Linux (what I am using). My computer asks if upon download if I want to run the installer, of course select YES.

Figure 28 is the selection page for selecting the installers. Once I select the package the computer starts the download, and asks me if I want to run the installer, as in Figure 29




Figure 29. Here we select install, and let the installer do its thing.

Selecting yes, the installer will attempt to install **R Studio** onto my computer. Once installation is complete, the program is ready for a validation (of install) run.




Figure 30. To launch **R Studio**, either select in the applications folder (or type in the terminal `rstudio`).

Figure 30 is a screen capture after installation using the Ubuntu program manager window. Both **R** and **R Studio** are available.

Either select **R Studio** to launch it, or type in a terminal window `rstudio` and the IDE should launch. The IDE itself is a bit complicated, but actually enables us to keep better track of our work and ultimately saves time. MatLab users should notice that the interface looks quite similar (at least it does to me) – its also the same concept.

Figure 31 is the result of launching the program. The left side of the IDE is an **R** console – exactly what would occur if we had just started **R**. The right side of the IDE is divided into an upper and lower panel. Each panel provides information about our programming environment at any instant – and the content is selectable from the icons at the top of each panel.




Figure 31. Upon opening you should have the following integrated development environment (IDE). The left panel is exactly what you would obtain if you just type `r` in the terminal window. .

The next few figures are a step-by-step example to introduce **R Studio** as well as test if it installed correctly. We could simply type in the **R** console within the IDE, but instead to get into the habit of saving our work, we will open a new scripting file.

Figure 32 shows the process of selecting FILE/OPEN to create a new file to store our scripts. We will type a few commands into the file and then run them.

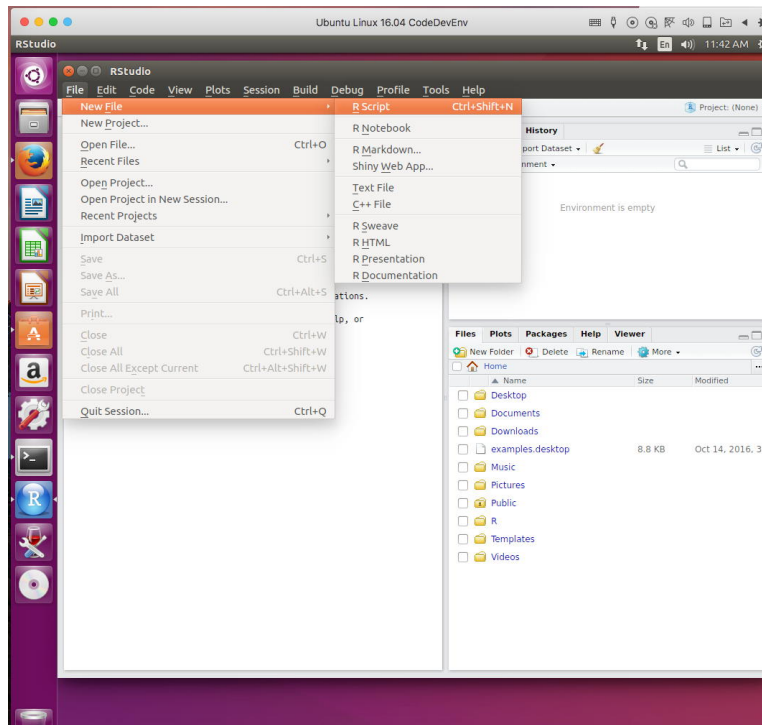


Figure 32. Open a new file – in this case as an R-Script.




Figure 33. Type in some R instructions, select the instructions and choose RUN.

So once the file is open the left side of the IDE divides into two panels, an upper and lower panel. The lower panel is the console environment, and the upper panel is a script editor.

Figure 33 shows the upper panel with the **R** commands shown in Listing 1

Listing 1. R code demonstrating a few commands

This fragment of code generates two vectors X and Y and then plots them.

```
##### Some R Commands #####
x <- c(0,1,2,3,4,5) # create a vector of 6 elements -- integers 0 to 5.
y <- x*x           # square x, put result in y.
plot(x,y,xlab="X-Axis",ylab="Y-Axis (X-squared)",lwd=3, type="l") # make and label a plot
```

To actually run the code we can either highlight the portion of the script file we wish the program to execute (that's what done here), or we can save the file and run the entire file. Often we will do both – highlight portions to develop a model, then save and run the file as needed. The term “sourcing” a [file] in R is jargon for running the commands contained within a file named [file]. The ability to save and reuse commands is really useful and is what makes **R** (or any other stored program software) really useful.

Figure 34 is the result of highlighting these three lines of code and running them (notice the little run icon above the script editor).

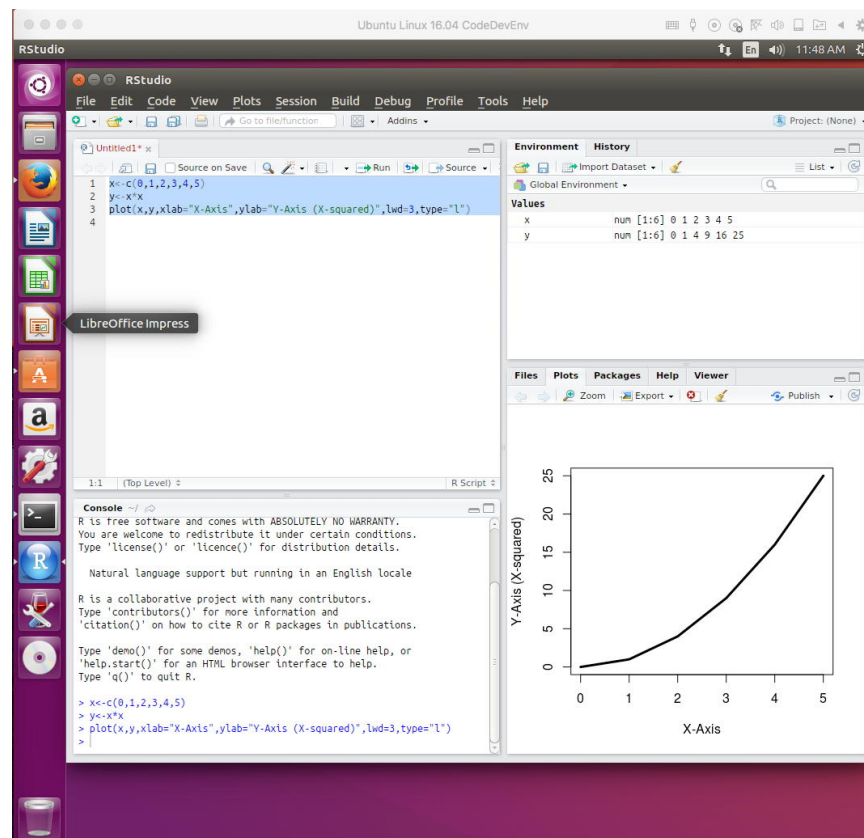


Figure 34. Completed script run. Notice the plot in the lower right panel. At this point you have a functioning programming tool..

2 Algorithms and R for Computing

An algorithm is a recipe. A useful definition is

An algorithm is a computational method or an ensemble of rules determining the order and form of numerical operations to be applied to a set of data $\mathbf{a}(a_1, a_2, \dots)$ in order to find a new set of values $\mathbf{x}(x_1, x_2, \dots)$ forming the solution of a problem.

An algorithmic procedure can be represented as

$$\mathbf{x} = f(\mathbf{a}) \quad (1)$$

From a mathematical perspective the main concern is that the algorithm is well posed:

1. A solution exists for a given \mathbf{a} .
2. The computation must lead to a single solution for \mathbf{x} given \mathbf{a} .
3. The results for \mathbf{x} must be connected to the input \mathbf{a} through the Lipschitz relation.

$$|\delta\mathbf{a}| < \eta \quad \text{then} \quad |\delta\mathbf{x}| < M|\delta\mathbf{a}| \quad (2)$$

where M is a bounded natural number, $M = M(\mathbf{a}, \eta)$ ⁶.

Certain common problems are not well posed as stated but with reasonable assumptions can be forced into such a state.

Thus an algorithm is a recipe to take input data and produce output responses through some relationships. If a well posed problem then each result is related to the inputs, and the same inputs (in an algorithm) produce the same results. By the recipe analogy, if you follow the same recipe each time with the same raw materials then the cake should taste the same when it is baked.

An important concept is that an algorithm operates on data (procedure-oriented); an object-oriented view is that an algorithm performs a task (generate response) based on states established by the data. Both points of view are valid and equivalent.

Most computational hydraulics models are built (by a quirk of history) in a procedure-oriented perspective.

2.1 Tools

A practicing modeler needs a toolkit — these tools range from the actual computation engine (EPA-SWMM, HEC-RAS, FESWMS, HSPF, WSPRO, TR-20, etc.)

⁶Think of M as a mapping function.

to analysis tools for result interpretation (**R**, **Excel**) to actual programming tools (**FORTRAN**, **PERL**, etc.) to construct their own special purpose models or to test results from general purpose professional models.

In this book **R** will be used for programming, analysis, and presentation.

2.2 Programming

Why programming?

There are three fundamental reasons for requiring a programming experience:

1. Teaching someone else a subject or procedure forces the teacher to have a reasonable understanding of the subject or procedure. Teaching a computer (by virtue of programming) forces a very deep understanding of the underlying algorithm.
2. You will encounter situations that general purpose programs are not designed to address; if you have a moderate ability to build your own tools when you need to, then you can. In all likelihood, you will “trick” the professional program, but you cannot invent tricks unless you know a little bit about programming.
3. Programming a computer requires an algorithmic thought process — this process is valuable in many other areas of engineering, hence the act of programming is good discipline for other problems you will encounter.

2.3 Interpolating Tabular Data – A useful algorithm

Material properties in physical systems are usually tabulated values. A frequent task is to interpolate in a set of tabular values to approximate the value between rows (or columns) in the table. Linear interpolation is the common technique used; and the tables can be stores as either separate files, or, if the tables are small enough, they can be directly imbedded into the code.

2.4 Linear Interpolation

Figure 35 is a sketch of a set of ordered pairs (x, y) .

These pairs (there are two in the sketch) represent values in a table, for instance x may represent water temperature, and y may represent vapor pressure at that particular temperature. Two adjacent values (in the table) are depicted in the sketch, and the pairs are ordered bases on the x -value.

Now suppose we want to estimate the value of y^* at some intermediate value x^* that lies between x_1 and x_2 . As a computational task, the problem statement is to




Figure 35. Sketch of two adjacent values from a table, plotted in Cartesian coordinate system..

“Estimate the value of y^* associated with the value x^* given the ordered pairs (x_1, y_1) and (x_2, y_2) .”

Linear interpolation simply uses the concept of similar triangles to scale the x and y distances between the ordered pairs to the intermediate location. Equation 3 is the result of application of similar triangles to the situation described by Figure ?? and the problem statement.

$$\frac{x^* - x_1}{x_2 - x_1} = \frac{y^* - y_1}{y_2 - y_1} \quad (3)$$

Next, apply algebra to solve Equation 3 for y^* , to obtain Equation 4

$$y^* = y_1 + \frac{(y_2 - y_1)(x^* - x_1)}{(x_2 - x_1)} \quad (4)$$

Now we can use 4 to estimate values between any two data pairs.

2.4.1 Interpolation of Values in Two Pairs

Figure 36 is a table of water properties from (CITE), that represents typically how tabular data are presented. The temperature column is arranged in increasing order and the other properties associate with temperature across a row.

Temperature	Density	Specific Weight	Dynamic Viscosity	Kinematic Viscosity	Vapor Pressure
	kg/m ³	N/m ³	N · s/m ²	m ² /s	N/m ² abs
0°C	1000	9810	1.79 × 10 ⁻³	1.79 × 10 ⁻⁶	611
5°C	1000	9810	1.51 × 10 ⁻³	1.51 × 10 ⁻⁶	872
10°C	1000	9810	1.31 × 10 ⁻³	1.31 × 10 ⁻⁶	1,230
15°C	999	9800	1.14 × 10 ⁻³	1.14 × 10 ⁻⁶	1,700
20°C	998	9790	1.00 × 10 ⁻³	1.00 × 10 ⁻⁶	2,340
25°C	997	9781	8.91 × 10 ⁻⁴	8.94 × 10 ⁻⁷	3,170
30°C	996	9771	7.97 × 10 ⁻⁴	8.00 × 10 ⁻⁷	4,250
35°C	994	9751	7.20 × 10 ⁻⁴	7.24 × 10 ⁻⁷	5,630
40°C	992	9732	6.53 × 10 ⁻⁴	6.58 × 10 ⁻⁷	7,380
50°C	988	9693	5.47 × 10 ⁻⁴	5.53 × 10 ⁻⁷	12,300
60°C	983	9643	4.66 × 10 ⁻⁴	4.74 × 10 ⁻⁷	20,000
70°C	978	9594	4.04 × 10 ⁻⁴	4.13 × 10 ⁻⁷	31,200
80°C	972	9535	3.54 × 10 ⁻⁴	3.64 × 10 ⁻⁷	47,400
90°C	965	9467	3.15 × 10 ⁻⁴	3.26 × 10 ⁻⁷	70,100
100°C	958	9398	2.82 × 10 ⁻⁴	2.94 × 10 ⁻⁷	101,300

Figure 36. Table of water properties in SI units (from CITE).

Now suppose we wished to estimate the density of water at 44° C. The two ordered pairs of temperature and density that surround 44° C are (40° C, 992 kg/m³) and (50° C, 988 kg/m³). So, to estimate the unknown density we can apply Equation 4 and obtain the following result

$$y^* = 992 + \frac{(988 - 992)(44 - 40)}{(50 - 40)} = 990.4 \quad (5)$$

We might want to do this a lot, so we could write a simplistic script in R and remember to load it into our environment when we need it

Listing 2. R code demonstrating the interpolation equation.

```
# EXAMPLE # ** Interpolating Between Tabulated Pairs
interpolate2pairs<-function(xstar,x1,y1,x2,y2){
# apply interpolation equation
# does not trap errors (divide by zero, etc)
  ystar <- y1 + (y2-y1)*(xstar-x1)/(x2-x1)
  return(ystar)
}
# In R Console
> interpolate2pairs(44,40,992,50,988)
[1] 990.4
>
```

For a single interrogation of a table we can stop here, but in many instances we have to interrogate a table a lot – we want some kind of program structure to handle the work so all we have to do is pass the temperature value and have the program return the density.

2.4.2 Interpolation of Values in Two Arrays

To accomplish repeated interpolation we will need to have: (1) an interpolating method (we have the beginning of one above in Listing 2), (2) the entire table so we don't have to enter the pairs, and (3) a way to automatically search the table so we don't have to look up values and supply them to the interpolator.

The table itself in this instance is relatively small, so we can simply assign values to some constant arrays in below in Listing 3.

Listing 3. R code assigning Liquid Properties.

```
# EXAMPLE # ** Assigning Constants
tempSI<-c(0.00,5.00,10.00,15.00,20.00,25.00,30.00,35.00,
40.00,50.00,60.00,70.00,80.00,90.00,100.00)
densitySI<-c(1000.00,1000.00,1000.00,999.00,998.00,997.00,996.00,
994.00,992.00,988.00,983.00,978.00,972.00,965.00,958.00)

# In R Console
> cbind(tempSI,densitySI)
  tempSI densitySI
[1,]      0      1000
[2,]      5      1000
[3,]     10      1000
[4,]     15       999
[5,]     20       998
[6,]     25       997
[7,]     30       996
[8,]     35       994
[9,]     40       992
[10,]    50       988
[11,]    60       983
[12,]    70       978
[13,]    80       972
[14,]    90       965
[15,]   100       958
>
```

Returning to our example, the value 44 lies between `tempSI[9]` and `tempSI[10]`, so we desire an algorithm that starts at `tempSI[1]` and determines if the search value is between `tempSI[1]` and `tempSI[2]`, if not, then increment the row counter and determine if the search value is between `tempSI[2]` and `tempSI[3]`, and so on.

Once we locate in the searched array where the value lies then the interpolation uses the lower and upper elements of the range to interpolate. In the case of our example, once we determine the 44 lies between `tempSI[9]` and `tempSI[10]`, then the interpolation equation is

$$y^* = \text{densitySI}[9] + \frac{(\text{densitySI}[10] - \text{densitySI}[9])(44 - \text{tempSI}[9])}{(\text{tempSI}[10] - \text{tempSI}[9])} \quad (6)$$

Listing 4 is an **R** script that implements the search and interpolation just described. The script assumes that the searched array (x) is ordered and increasing – not a trivial assumption! The script has some limited error handling to test if the search value actually lies in the total range of the array before beginning the search. Once these tests are passed, then the code searches in the x array for the search value x^* and finds the two rows that contain the value. Once the rows are located, the interpolation equation is used.

Listing 4. R code to Search and Interpolate.

```
# EXAMPLE # ** Search and Interpolate
getAvalue <- function(x,xvector,yvector){
  # returns a y value for x interpolated from (xvector,yvector)
  # xvector is assumed to be in a monotonic sequence
  # function performs limited error checks
  # NULL return is error indicator
  # T.G. Cleveland July 2007
  #
  xvlength <- length(xvector)
  yvlength <- length(yvector)
  # check that vector lengths same
  if(xvlength != yvlength){
    message("vectors xvector and yvector different lengths -- exiting function")
    return()
  }
  # check that x in range xvector
  if(x < min(xvector)){
    message(" x too small -- exiting function")
    return()
  }
  if(x > max(xvector)){
    message(" x too big -- exiting function")
    return()
  }
  #
  for (i in 1:(xvlength-1)){
    if( (x >= xvector[i]) & (x < xvector[i+1]) ){
      result = yvector[i]+(yvector[i+1]-yvector[i])*(x - xvector[i])/
        (xvector[i+1]-xvector[i])
      return(result)
    }
    # next row
  }
  # check if at endpoint
  if( (x >= xvector[xvlength-1]) & (x <= xvector[xvlength]) ){
    result = yvector[i]+(yvector[i+1]-yvector[i])*(x - xvector[i])/
      (xvector[i+1]-xvector[i])
    return(result)
  }
  # should never get to next line
  message("something is really wrong -- check the vectors!")
  return()
}
# In R Console:
> getAvalue(44,tempSI,densitySI)
[1] 990.4
>
```

Now we can load and run the `getAvalue` script and supply the two vectors plus the search value as shown in Listing 4

This look-up process is readily transferred to other cases, we do have to decide if the data will be coded as constants (as was done here) or read from a file – if the database is large the file read option is best. In terms of building a generic look-up tool several things actually happen in a particular order.

1. The function call loads in the table (of reading from a file, we would have to forward declare the vectors).
2. The function searches the first vector for the bounding location of the search variable.
3. Once the boundaries are located, the interpolation is performed – notice how the last boundary pair is handled.

Now we can combine the data assignment, the search and interpolate into a single function so when we want to evaluate in the future we only need the single function.

Listing 5 is an example of everything combined. Here I have embedded the `getAValue` script into the function so the whole function itself is portable (we don't have keep track of `getAValue`). This embedding can be replaced with a load from a library (but then we must keep track of the path).

The library way is preferable; if `getAValue` needs changing, we will have to change every instance of it in the code, if we miss one the code may still run and it could be years before we discover the error because a single instance of code fragment within a larger code was missed – its far better to only change a single instance of the function when maintenance is necessary.

Listing 5. R code to Return Water Density for Given Temperature.

```
# Script to return water density in SI units as a function of temperature
getDensitySI<-function(t){
# load the getAValue() function #####
getAValue <- function(x,xvector,yvector){
# returns a y value for x interpolated from (xvector,yvector)
# NULL return is error indicator
#
xvlength <- length(xvector)
yvlength <- length(yvector)
# check that vector lengths same
if(xvlength != yvlength){
message("vectors xvector and yvector different lengths -- exiting function")
return()
}
# check that x in range xvector
if(x < min(xvector)){
message(" x too small -- exiting function")
return()
}
if(x > max(xvector)){
message(" x too big -- exiting function")
return()
}
#
for (i in 1:(xvlength-1)){
if( (x >= xvector[i]) & (x < xvector[i+1]) ){
result = yvector[i]+(yvector[i+1]-yvector[i])*(x - xvector[i])/
(xvector[i+1]-xvector[i])
return(result)
}
# next row
}
# check if at endpoint
if( (x >= xvector[xvlength-1]) & (x <= xvector[xvlength]) ){
result = yvector[i]+(yvector[i+1]-yvector[i])*(x - xvector[i])/
(xvector[i+1]-xvector[i])
return(result)
}
# should never get to next line
message("something is really wrong -- check the vectors!")
return()
}
#####
# load the data vectors, tempSI and densitySI
tempSI<-c(0.00,5.00,10.00,15.00,20.00,25.00,30.00,35.00,
40.00,50.00,60.00,70.00,80.00,90.00,100.00)
```

```

densitySI<-c(1000.00,1000.00,1000.00,999.00,998.00,997.00,996.00,
  994.00,992.00,988.00,983.00,978.00,972.00,965.00,958.00)
# now call getAValue
result<-getAValue(t,tempSI,densitySI)
return(result)
}

```

The “library” approach is demonstrated in Listing 6; in this listing the path in the `source()` command is unique to my machine – your path is likely to be different. I find it is useful to contain all the various codes into a single directory and source that directory once to find the path, then change all the source calls to that path. In fact that path can be a string variable and the referencing can be automatic (as long as the files exist!).

Once the look-up function is built then we can interrogate the table many times; and even build a plot of the table – these features are demonstrated in Listing 6.

Listing 6. R code demonstrating use of `getDensitySI()`.

```

## In R Console
> # Example demonstrating use of functions
> # load in the functions (must exist -- use path on your machine)
> source('~/Dropbox/1-CE-TTU-Classes/UnderDevelopment/
  CE4333-PCH-R/6-RScripts/getAValue.R')
> source('~/Dropbox/1-CE-TTU-Classes/UnderDevelopment/
  CE4333-PCH-R/6-RScripts/getDensitySI.R')
> # Now use them
> getDensitySI(44)
[1] 990.4
> getDensitySI(54)
[1] 986
> getDensitySI(88)
[1] 966.4
> t<-seq(0,100,2) # make a temperature vector 0 to 100 in 2 degree increments
> d<-numeric(0) # forward declare d to store results
> howMany<-length(t)
> for(i in 1:howMany){
+   d[i]<-getDensitySI(t[i])
+ }
> plot(t,d,type="l",xlab="Degrees Celsius",ylab="Density (kg/m^3)")
>

```

The resulting plot is shown on Figure 37 below.

2.5 Sorting

Another frequent task in engineering hydraulics is the seemingly mundane task of sorting or ordering things. Here we explore a couple of simple sorting algorithms, just to show some of the thoughts that go into such a task, then will ultimately resort to the internal sorting routines built into R.

2.5.1 Bubble Sort

The bubble sort is a place to start despite it’s relative slowness. It is a pretty reviled algorithm (read the Wikipedia entry), but it is the algorithm that a naive programmer might cobble together in a hurry, and despite its shortcomings (its really slow and inefficient), it is robust.




Figure 37. Plot of density versus temperature generated using the `getDensity()` function..

Here is a description of the sorting task as described by Christian and Griffiths (2016) (pg. 65):

“Imagine you want to alphabetize your unsorted collection of books. A natural approach would be just to scan across the shelf looking for out-of-order pairs – Wallace followed by Pynchon, for instance – and flipping them around. Put Pynchon ahead of Wallace, then continue your scan, looping around to the beginning of the shelf each time you reach the end. When you make a complete pass without finding any more out-of-order pairs on the entire shelf, then you know the job is done.

This process is a Bubble Sort, and it lands us in quadratic time. There are n books out of order, and each scan through the shelf can move each one at most one position. (We spot a tiny problem, make a tiny fix.) So in the worst case, where the shelf is perfectly backward, at least one book will need to be moved n positions. Thus a maximum of n passes through n books, which gives us $O(n^2)$ in the worst case.⁷ For instance, it means that sorting five shelves of books will take not five times as long as sorting a single shelf, but twenty-five times as long.”

Converting the word description into **R** is fairly simple. We will have a vector of n numbers (we use a vector because its easy to step through the different positions), and we will scan through the vector once (and essentially find the smallest thing), and put it into the first position. Then we scan again from the second position and find the smallest thing remaining, and put it into the second position, and so on until the last scan which should have the remaining largest thing. If we desire a decreasing order, simply change the sense of the comparison.

Listing 7 is an **R** script that implements the algorithm – in the script the actual sort is treated as a function (we may actually want to use it again someday) which is loaded into the programming environment first, then an array is defined, and sorted. The program (outside of the sorting algorithm) is really quite simple.

- Load the sorting function.
- Load contents into an array to be sorted.
- Echo (print) the array (so we can verify the data are loaded as anticipated).
- Sort the array, put the results back into the array (an in-place sort).
- Report the results.

Listing 7. R code demonstrating the naive bubble sort.

```
#####
rm(list=ls()) # clear the object list (i.e. deallocate and clear memory)
### Bubble Sort Function -- Needs to be defined before sending array to sort ###
# Bubble Sort Function
```

⁷Actually, the average running time for Bubble Sort isn’t any better, as books will, on average, be $n/2$ positions away from where they’re supposed to end up. One would round the $n/2$ passes of n books up to $O(n^2)$.

```

# MyLocation: ~/Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/1-Lectures/Lecture03/ScriptsInLecture
# Bubble Sort with array indexing starting at [1]
# Compare to Python or C where arrays start at [0]
# by: Theodore G. Cleveland 2017-0317
#####
bubble <- function(array)
{
  # Prepare the sort, need to know how many things and need a temporary store
  swap <- numeric(0) # temporary store (aka swap location)
  howMany <- length(array) # how many things to be sorted
  # The actual sorting process
  for (irow in 1:(howMany-1))
  {
    for (jrow in 1:(howMany-irow))
    {
      if( array[jrow] > array[jrow+1])
      {
        swap <- array[jrow];
        array[jrow] <- array[jrow+1];
        array[jrow+1] <- swap;
      }
    }
  }
  # return result (sort in-place)
  return(array)
}
#####
#####
xarray <- c(1003,3.2,55.5,-0.0001,-6,666.6,102) # the array to sort
print(xarray)
xarray <- bubble(xarray)
print(xarray)
#####

```

Figure 38 is a screen capture of the script running. In the figure we see that the program (near the bottom of the file) assigns the values to the vector named array and the initial order of the array is [1003, 3.2, 55.5, -0.0001, -6, 666.6, 102]. The smallest value in the example is -6 and it appears in the 5-th position, not the 1-st as it should.

The first pass through the array will move the largest value, 1003, in sequence to the right until it occupies the last position. Repeated passes through the array move the remaining largest values to the right until the array is ordered. One can consider the values of the array at each scan of the array as a series of transformations (*irow*-th scan) – in practical cases we don't necessarily care about the intermediate values, but here because the size is manageable and we are trying to get our feet wet with algorithms, we can look at the values.

The sequence of results (transformations) after each pass through the array is shown in the following list:

1. Initial value: [1003, 3.2, 55.5, -0.0001, -6, 666.6, 102].
2. First pass: [3.2, 55.5, -0.0001, -6, 666.6, 102, 1003].
3. Second pass: [3.2, -0.0001, -6, 55.5, 102, 666.6, 1003].
4. Third pass: [-0.0001, -6, 3.2, 55.5, 102, 666.6, 1003].
5. Fourth pass: [-6, -0.0001, 3.2, 55.5, 102, 666.6, 1003].
6. Fifth pass: [-6, -0.0001, 3.2, 55.5, 102, 666.6, 1003]. Sorted, fast scan.
7. Sixth pass: [-6, -0.0001, 3.2, 55.5, 102, 666.6, 1003]. Sorted, fast scan.


```

R Console
~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/1-Lectures/Lecture03/
> # by: Theodore G. Cleveland 2017-0317
> #####
> bubble <- function(array)
+ {
+   # Prepare the sort, need to know how many things and need a temporary store
+   swap <- numeric(0) # temporary store (aka swap location)
+   howMany <- length(array) # how many things to be sorted
+   # The actual sorting process
+   for (irow in 1:(howMany-1))
+   {
+     for (jrow in 1:(howMany-irow))
+     {
+       if( array[jrow] > array[jrow+1])
+       {
+         swap <- array[jrow];
+         array[jrow] <- array[jrow+1];
+         array[jrow+1] <- swap;
+       }
+     }
+   }
+   # return result (sort in-place)
+   return(array)
+ }
> #####
> ##### Populate and sort an array #####
> xarray <- c(1003,3.2,55.5,-0.0001,-6,666.6,102) # the array to sort
> print(xarray)
[1] 1003.0000  3.2000  55.5000  -0.0001  -6.0000  666.6000  102.0000
> xarray <- bubble(xarray)
> print(xarray)
[1] -6.0000 -0.0001  3.2000  55.5000  102.0000  666.6000  1003.0000
>

```

Figure 38. Bubble Sort implemented in R.

We could probably add additional code to break from the scans when we have a single pass with no exchanges – while meaningless in this example, for larger collections of things, being able to break out when the sorting is complete is a nice feature.

2.5.2 Insertion Sort

The next type of sorting would be to select one item and locate it either left or right of an adjacent item based on its size – like sorting a deck of cards, or perhaps a better description – again using the bookshelf analog from Christian and Griffiths (2016) (pg. 65):

“..... You might take a different tack – pulling all the books off the shelf and putting them back in place one by one. You’d put the first book in the middle of the shelf, then take the second and compare it to the first, inserting it either to the right or to the left. Picking up the third book, you’d run through the books on the shelf from left to right until you found the right spot to tuck it in. Repeating this process, gradually all of the books would end up sorted on the shelf and you’d be done. Computer scientists call this, appropriately enough, Insertion Sort. The good news is that it’s arguably even more intuitive than Bubble Sort and doesn’t have quite the bad reputation. The bad news is that it’s not actually that much faster. You still have to do one insertion for each book. And each insertion still involves moving past about half the books on the shelf, on average, to find the correct place.

Although in practice Insertion Sort does run a bit faster than Bubble Sort, again we land squarely, if you will, in quadratic time. Sorting anything more than a single bookshelf is still an unwieldy prospect.”

Listing 8 is an **R** implementation of a straight insertion sort. The script is quite compact, and I used indentation and extra line spacing to keep track of the scoping delimiters. The sort works as follows, take the an element of the array (start with 2 and work to the right) and put it into a temporary location (called **swap** in my script). Then compare locations to the left of **swap**. If smaller, then break from the loop, exchange values, otherwise the values are currently ordered. Repeat (starting at the next element) , when all elements have been traversed the resulting vector is sorted. Here are the transformations for each pass through the outer loop:

1. Pass 0: [1003, 3.2, 55.5, -0.0001, -6, 666.6, 102], Initial array.
2. Pass 1: [3.2, 1003, 55.5, -0.0001, -6., 666.6, 102].
3. Pass 2: [3.2, 55.5, 1003, -0.0001, -6., 666.6, 102].
4. Pass 3: [-0.0001, 3.2, 55.5, 1003, -6., 666.6, 102].
5. Pass 4: [-6, -0.0001, 3.2, 55.5, 1003., 666.6, 102].
6. Pass 5: [-6, -0.0001, 3.2, 55.5, 666.6, 1003, 102].
7. Pass 6: [-6, -0.0001, 3.2, 55.5, 102, 666.6, 1003], Sorted array.

Figure 39 is a screen capture of the insertion sort in operation. Insertion sorting is

reasonably fast for small lists (about 50 or so elements) and forms the basis of the internal sorts in other routines that divide up the overall list into smaller lists, sort the smaller lists, then uses a merge to collate back to the overall list (now sorted).

Listing 8. R code demonstrating the insertion sort.

```
### Straight Insertion Sort Function by: Theodore G. Cleveland 2017-0317
rm(list=ls()) # clear the object list (i.e. deallocate and clear memory)
#####
insertSort <- function(array){
# Prepare the sort, need to know how many things and need a temporary store
swap <- numeric(0) # temporary store (aka swap location)
howMany <- length(array) # how many things to be sorted
for (j in 2:howMany) # select each position in turn
{
  test <- 0 # set a test value, used to insert later
  swap <- array[j] # current position to swap
  for (i in seq(j-1,1,by=-1)) #find place to insert by ...
  {
    if (array[i] <= swap) # test if current position is bigger
    {
      test <- 1 # if true set test to 1, break inner loop
      break
    }
    array[i+1] <- array[i] # otherwise exchange postions
  }
  if(test == 1) # if broke from loop, insert swap
  array[i+1] <- swap
  else
  i = 0
  array[i+1] <- swap # otherwise swap goes to first position
}
return(array) } # return result (sort in-place)
#####
xarray <- c(1003,3.2,55.5,-0.0001,-6,666.6,102) # the array to sort
print(xarray)
xarray <- insertSort(xarray)
print(xarray)
#####
```

```

RStudio
Project: (None)
MyInsertionSort.R
Source on Save
Run Source
24     test <- 1 # if true set test to 1, break inner loop
25     break
26   }
27   array[i+1] <- array[i] # otherwise exchange postions
28 }
29 if(test == 1) # if broke from loop, insert swap
30 array[i+1] <- swap
31 else
32 i = 0
33 array[i+1] <- swap # otherwise swap goes to first position
34
35 # return result (sort in-place)
36 return(array)
37 }
38 #####
39
40 #####
41 xarray <- c(1003,3.2,55.5,-0.0001,-6,666.6,102) # the array to sort
42 print(xarray)
43 xarray <- insertSort(xarray)
44 print(xarray)
45 #####
17:6 insertSort
R Script

Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/1-Lectures/Lecture03/ScriptsInLecture/
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/1-Lectures/Lecture03/ScriptsInLecture/MyInsertionSort.R')
[1] 1003.0000 3.2000 55.5000 -0.0001 -6.0000 666.6000 102.0000
[1] -6.0000 -0.0001 3.2000 55.5000 102.0000 666.6000 1003.0000
>

```

Figure 39. Insertion Sort implemented in R.

2.5.3 Merge Sort

A practical extension of these slow sorts is called the Merge Sort. It is an incredibly useful method. One simply breaks up the items into smaller arrays, sorts those arrays - then merges the sub-arrays into larger arrays (now already sorted), and finally merges the last two arrays into the final, single, sorted array.

Here is a better description, again from Christian and Griffiths (2016):

“..... information processing began in the US censuses of the nineteenth century, with the development, by Herman Hollerith and later by IBM, of physical punch-card sorting devices. In 1936, IBM began producing a line of machines called “collators” that could merge two separately **ordered** stacks of cards into one. *As long as the two stacks were themselves sorted*, the procedure of merging them into a single sorted stack was incredibly straightforward and took linear time: simply compare the two top cards to each other, move the smaller of them to the new stack you’re creating, and repeat until finished.

The program that John von Neumann wrote in 1945 to demonstrate the power of the stored-program computer took the idea of collating to its beautiful and ultimate conclusion. Sorting two cards is simple: just put the smaller one on top. And given a pair of two-card stacks, both of them sorted, you can easily collate them into an ordered stack of four. Repeating this trick a few times, you’d build bigger and bigger stacks, each one of them already sorted. Soon enough, you could collate yourself a perfectly sorted full deck – with a final climactic merge, like a riffle shuffle’s order-creating twin, producing the desired result. This approach is known today as Merge Sort, one of the legendary algorithms in computer science.”

There are several other variants related to Merge Sort; Quicksort and Heapsort being relatives. The creation of a Merge Sort is left to the reader if there is a need, and at this point we can just use the built-in `sort()` and/or `order()` functions in **R** – which implements either a Shellsort (useful if character strings are to be sorted) or Quicksort (used if numeric values are supplied). We also have to supply if we want increasing or decreasing sorts.

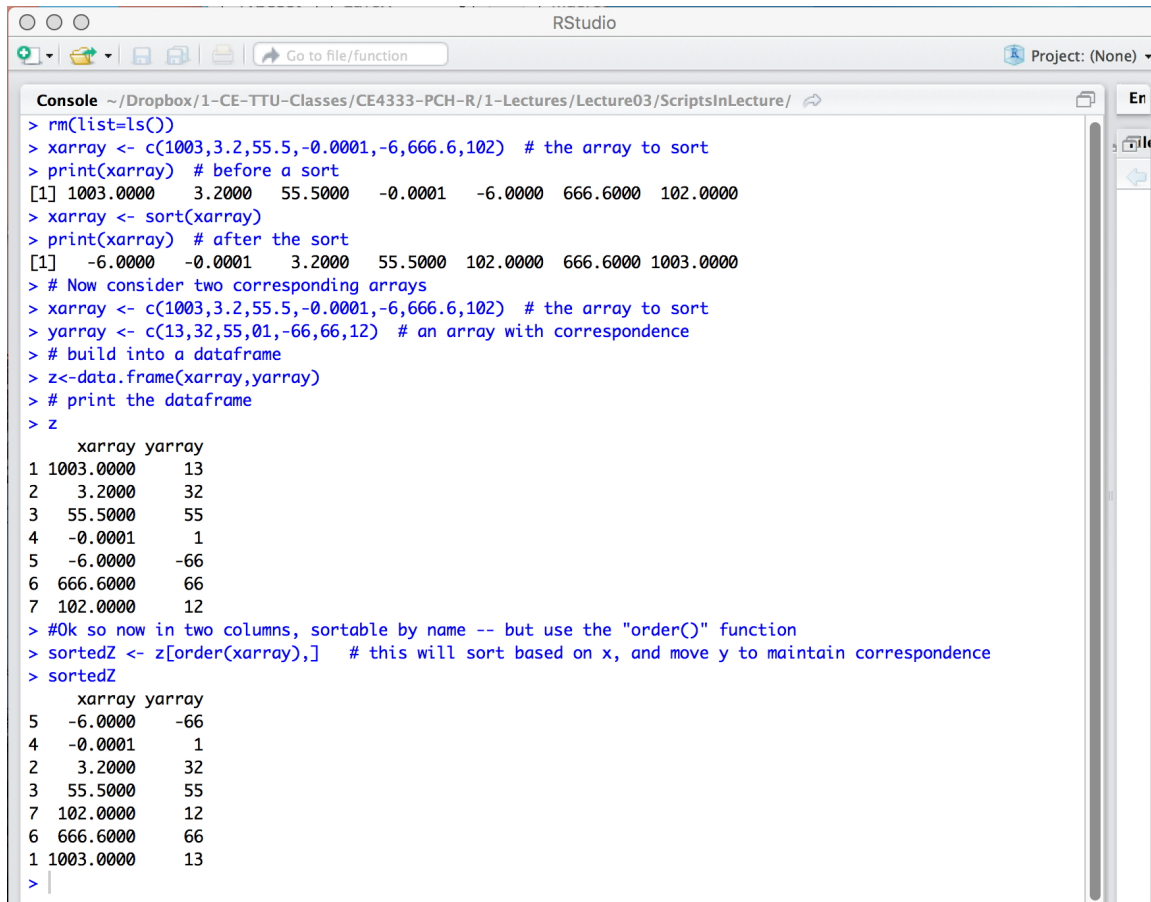
2.5.4 Built-In R Sorts

Figure 40 illustrates using the built-in functions. For an ordinary sort, we simply use the function name `sort()` and direct its output into an object (it can even be the same vector as shown in the figure).

If we wish to sort several related columns, based on values in one of the columns, it is easiest to construct a data frame (like a matrix), then order the contents based on one of the columns, and send the results to another data frame, or we can send

the result back to itself. Usually when we are manipulating multiple columns, we are operating in a “relational database” kind of mindset, and it is probably to our benefit to not destroy the original association structure. Be aware of the syntax of a dataframe function, you will notice there is a comma that appears at the end of the function that is important for the script to function.

For example, `z <- z[order(xarray),]` will function as shown, whereas `zztop <- z[order(xarray)]` will not.



```

RStudio
~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/1-Lectures/Lecture03/ScriptsInLecture/
Project: (None)
Console
> rm(list=ls())
> xarray <- c(1003,3.2,55.5,-0.0001,-6,666.6,102) # the array to sort
> print(xarray) # before a sort
[1] 1003.0000  3.2000  55.5000  -0.0001  -6.0000  666.6000  102.0000
> xarray <- sort(xarray)
> print(xarray) # after the sort
[1] -6.0000  -0.0001  3.2000  55.5000  102.0000  666.6000  1003.0000
> # Now consider two corresponding arrays
> xarray <- c(1003,3.2,55.5,-0.0001,-6,666.6,102) # the array to sort
> yarray <- c(13,32,55,01,-66,66,12) # an array with correspondence
> # build into a dataframe
> z<-data.frame(xarray,yarray)
> # print the dataframe
> z
  xarray yarray
1 1003.0000    13
2   3.2000    32
3  55.5000    55
4  -0.0001     1
5  -6.0000   -66
6  666.6000    66
7  102.0000    12
> #Ok so now in two columns, sortable by name -- but use the "order()" function
> sortedZ <- z[order(xarray),] # this will sort based on x, and move y to maintain correspondence
> sortedZ
  xarray yarray
5  -6.0000   -66
4  -0.0001     1
2   3.2000    32
3  55.5000    55
7  102.0000    12
6  666.6000    66
1  1003.0000    13
>

```

Figure 40. Sorting using built-in R functions.

Now if we return to the interpolation chapter just before this one, we can immediately see a need for sorting. The interpolation algorithm *assumes* that the explanatory structure (x-axis) is ordered, otherwise the interpolation equation will return garbage.

I conclude the section on sorting with one more quoted section from Christian and Griffiths (2016) about the value for sorting – which is already relevant to a lot of computational hydraulics:

“The poster child for the advantages of sorting would be an Internet search engine like Google. It seems staggering to think that Google can take the search phrase you typed in and scour the entire Internet for it in less than

half a second. Well, it can't – but it doesn't need to. If you're Google, you are almost certain that (a) your data will be searched, (b) it will be searched not just once but repeatedly, and (c) the time needed to sort is somehow “less valuable” than the time needed to search. (Here, sorting is done by machines ahead of time, before the results are needed, and searching is done by users for whom time is of the essence.) All of these factors point in favor of tremendous up-front sorting, which is indeed what Google and its fellow search engines do.”

Exercise Set 1

1. Build a function `getDensityUS()` that searches the table in Figure 36 and returns the density of water in US customary units for a value of temperature supplied in degrees Fahrenheit.

Submit your code and screen captures of the density for temperatures of $43^\circ F$, $146^\circ F$, and $210^\circ F$.

2. Later in the class we will need functions to return viscosity to compute head losses in pipe networks.

Build and test a function `getKinViscosityUS()` that searches the table in Figure ?? and returns the kinematic viscosity of water in US customary units for a value of temperature supplied in degrees Fahrenheit

Submit the code and screen captures of the kinematic viscosity for temperatures of $43^\circ F$, $146^\circ F$, and $210^\circ F$.

3. Build and test a function `getKinViscositySI()` that searches the table in Figure 36 and returns the kinematic viscosity of water in SI units for a value of temperature supplied in degrees Celsius

Submit the code and screen captures of the kinematic viscosity for temperatures of $13^\circ C$, $66^\circ C$, and $97^\circ C$.

4. Imagine you have two arrays, `array1` and `array2` that are linked in the sense that each element of `array1` is associated with the corresponding element of `array2`. You wish to sort based on contents of `array1` and maintain the correspondence of `array2`. In words we would simply modify the script to move and element of `array2` whenever you move an element of `array1`.

Modify the Insertion Sort script to accept two arrays, the sort is based on contents of the first array and you are to maintain correspondence with the second array.

Test your script on the following arrays:

```
array1 = [5, 6, 8, 2, 3, 4, 1]
```

```
array2 = [24, 35, 63, 3, 8, 15, 0]
```

when these are reordered, the result should be:

```
array1 = [1, 2, 3, 4, 5, 6, 8]
```

```
array2 = [0, 3, 8, 15, 35, 63]
```

Then modify your script to read input from an ASCII file that contains the two arrays as columns (like in a spreadsheet) where you will sort on the first column, and carry along the correspondence with the second column.

Apply the script on the file `es3-pr1.txt`.

5. Repeat the exercise above but use built-in **R** functions.⁸

This exercise set is also located on the class server as `ES-1`.

⁸My solution uses the `data.frame()` and `order()` functions. There are probably several other ways to accomplish the goal – corresponding sorts are hugely important in many practical situations.

3 Numerical Methods – Integrals, Derivatives, and Newton's Method

3.1 Numerical Integration of Functions

Numerical integration is the numerical approximation of

$$I = \int_a^b f(x)dx \quad (7)$$

Consider the problem of determining the shaded area under the curve $y = f(x)$ from $x = a$ to $x = b$, as depicted in Figure 41, and suppose that analytical integration is not feasible.




Figure 41. Schematic of Panels for Numerical Integration. .

The function may be known in tabular form from experimental measurements or it may be known in an analytical form. The function is taken to be continuous within the interval $a < x < b$. We may divide the area into n vertical panels, each of width $\Delta x = (b - a)/n$, and then add the areas of all strips to obtain $A \approx \int y dx$.

A representative panel of area A_i is shown with darker shading in the figure. Three useful numerical approximations are listed in the following sections. The approximations differ in how the function is represented by the panels — in all cases the function is approximated by known polynomial models between the panel end points.

In each case the greater the number of strips, and correspondingly smaller value of Δx , the more accurate the approximation. Typically, one can begin with a relatively small number of panels and increase the number until the resulting area approximation stops changing.

3.1.1 Rectangular Panels

Figure 42 is a schematic of a rectangular panels. The figure is assuming the function structure is known and can be evaluated at an arbitrary location in the Δx dimension. Each panels is treated as a rectangle, as shown by the representative panel whose




Figure 42. Rectangular Panel Schematic..

height y_m is chosen visually so that the small cross-hatched areas are as nearly equal as possible. Thus, we form the sum $\sum y_m$ of the effective heights and multiply by Δx . For a function known in analytical form, a value for y_m equal to that of the function at the midpoint $x_i + \Delta x/2$ may be calculated and used in the summation.

For tabulated functions, we have to choose to either take y_m as the value at the left endpoint or right endpoint. This limitation is often quite handy when we are trying to integrate a function that is integrable, but undefined on one endpoint.

Lets try some examples in **R**.

Problem: Find the area under the curve $y = x\sqrt{1+x^2}$ from $x = 0$ to $x = 2$.

Solution: One solution is shown in Figure 43, which is a screen capture of a rudimentary code that implements the rectangular panel method.⁹

```

RStudio
Project: (None)
RectangularPanelExample.R
Source on Save
Run Source
5 # source('PATH-TO-THE-FILE/RectangularPanelExample.R') #
6 # where PATH-TO-THE-FILE is replaced with the actual path on your machine #
7 #####
8 #           Function to be integrated (modify as needed) #
9 #####
10 y <- function(x){
11   y <- x * sqrt(1+x^2)
12   return(y)
13 }
14 #####
15 #           Get lower, upper and how many panels from user #
16 #####
17 xlow <- readline("What is the lower limit of integration? ")
18 xhigh <- readline("What is the upper limit of integration? ")
19 howMany <- readline("How many panels? ")
20 # Convert the strings into numeric values
21 xlow <- as.numeric(unlist(strsplit(xlow, ",")))
22 xhigh <- as.numeric(unlist(strsplit(xhigh, ",")))
23 howMany <- as.numeric(unlist(strsplit(howMany, ",")))
24 # Compute some constants
25 deltax <- (xhigh - xlow)/howMany
26 accumulated_area <- 0.0 # initialize the accumulator
27 xx <- xlow+deltax/2 # initial value for x at middle of left-most panel
28 #####
29 #           The actual numerical method #
30 #####
31 for (i in 1:howMany){
32   accumulated_area <- accumulated_area + deltax*y(xx) # y is the integrand function
33   xx <- xx+deltax
34 }
35 #####
36 #           Report Result #
37 #####
38 message("Approximate value of integral from ",xlow," to ",xhigh," is: ",accumulated_area)
19:41 (Untitled) R Script
Console
~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RectangularPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 4
Approximate value of integral from 0 to 2 is: 3.36170409533974
>

```

Figure 43. Rectangular panel example showing code and resulting computed area using just 4 panels.

The script does not implement any kind of error checking – we could enter text values for the lower and upper values of x as well as the number of panels to use, and the script would attempt to run. A better version would force us to enter numeric values, and check for undefined ranges and such; devotion to error trapping is typical for professional programs where you are going to distribute executable modules and not expect the end user to be a programmer.

⁹The exact solution is $A=3.393477$

For the time being, we will accept this approach (error trapping is left as an exercise), however in your own scripts you should implement error traps where possible – you may start without them but as you maintain your scripts, you will learn where data entry errors occur and trap them.¹⁰

The actual listing depicted in Figure 43 is shown in Listing 9

Listing 9. R code demonstrating Rectangular Panel Numerical Integration.

```
# R script to implement rectangular panel numerical integration
##### NOTE #####
## The interactive input requires the script to be sourced #
# In R console the command line would be #
# source('PATH-TO-THE-FILE/RectangularPanelExample.R') #
# where PATH-TO-THE-FILE is replaced with the actual path on your machine #
#####
# Function to be integrated (modify as needed) #
#####
y <- function(x){
  y <- x * sqrt(1+x^2)
  return(y)
}
#####
# Get lower, upper and how many panels from user #
#####
xlow <- readline("What is the lower limit of integration? ")
xhigh <- readline("What is the upper limit of integration? ")
howMany <- readline("How many panels? ")
# Convert the strings into numeric values
xlow <- as.numeric(unlist(strsplit(xlow, ",")))
xhigh <- as.numeric(unlist(strsplit(xhigh, ",")))
howMany <- as.numeric(unlist(strsplit(howMany, ",")))
# Compute some constants
deltax <- (xhigh - xlow)/howMany
accumulated_area <- 0.0 # initialize the accumulator
xx <- xlow+deltax/2 # initial value for x at middle of left-most panel
#####
# The actual numerical method #
#####
for (i in 1:howMany){
  accumulated_area <- accumulated_area + deltax*y(xx) # y is the integrand function
  xx <- xx+deltax
}
#####
# Report Result #
#####
message("Approximate value of integral from ",xlow," to ",xhigh," is: ",accumulated_area)
```

Figure 44 is the same program run using 4,400, and 4000 panels observe the difference in computed area as well as the results closeness to the exact solution.

3.1.2 Trapezoidal Panels

The trapezoidal panels are approximated as shown in Figure 45. The area A_i is the average height $(y_i + y_{i+1})/2$ times Δx . Adding the areas gives the area approximation as tabulated. For the example with the curvature shown, the approximation will be on the low side. For the reverse curvature, the approximation will be on the high side. The trapezoidal approximation is commonly used with tabulated values.

The same example as presented for rectangular panels is repeated, except using trapezoidal panels. The code is changed because we will evaluate at each end of the panel (so no fussing to find an intermediate estimate for where to evaluate the function).

¹⁰For example, traps are used to force the user to enter a value that actually is meaningful, or select a default value, or internally prevent division by zero.

```

1 # R script to implement rectangular panel numerical integration
2 ##### NOTE #####
3 ## The interactive input requires the script to be sourced #
4 # In R console the command line would be #
5 # source('PATH-TO-THE-FILE/RectangularPanelExample.R') #
6 # where PATH-TO-THE-FILE is replaced with the actual path on your machine #
7 #####
8 # Function to be integrated (modify as needed) #
9 #####
10 y <- function(x){
11   y <- x * sqrt(1+x^2)
12   return(y)
13 }
14 #####
15 # Get lower, upper and how many panels from user #
16 #####
17 xlow <- readline("What is the lower limit of integration? ")
18 xhigh <- readline("What is the upper limit of integration? ")
19 howMany <- readline("How many panels? ")
20 # Convert the strings into numeric values
21 xlow <- as.numeric(unlist(strsplit(xlow, ",")))
22 xhigh <- as.numeric(unlist(strsplit(xhigh, ",")))
23 howMany <- as.numeric(unlist(strsplit(howMany, ",")))
24 # Compute some constants
25 deltax <- (xhigh - xlow)/howMany
26 accumulated area <- 0 # initialize the accumulator
27
1:1 (Top Level)
    
```

```

> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RectangularPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 4
Approximate value of integral from 0 to 2 is: 3.36170409533974
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RectangularPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 400
Approximate value of integral from 0 to 2 is: 3.39344347820325
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RectangularPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 4000
Approximate value of integral from 0 to 2 is: 3.39344659765633
>
    
```

Figure 44. Rectangular panel example showing difference in computed area using 4, 400, and 4000 panels. The 4000 panel result is essentially equivalent to the exact solution. Such convergence to exact values is typical.




Figure 45. Trapezoidal Panel Schematic..

Figure 46 illustrates the trapezoidal method for approximating an integral. In the example, the left and right panel endpoints in x are set as separate variables x_{left} and x_{right} and incremented by Δx as we step through the count-controlled repetition to accumulate the area. The corresponding y values are computed within the loop and averaged, then multiplied by Δx and added to the accumulator. Finally the x values are incremented.

```

17 xlow <- readline("What is the lower limit of integration? ")
18 xhigh <- readline("What is the upper limit of integration? ")
19 howMany <- readline("How many panels? ")
20 # Convert the strings into numeric values
21 xlow <- as.numeric(unlist(strsplit(xlow, ",")))
22 xhigh <- as.numeric(unlist(strsplit(xhigh, ",")))
23 howMany <- as.numeric(unlist(strsplit(howMany, ",")))
24 # Compute some constants
25 deltax <- (xhigh - xlow)/howMany
26 accumulated_area <- 0.0 # initialize the accumulator
27 xleft <- xlow
28 xright <- xleft + deltax
29 #####
30 # The actual numerical method
31 #####
32 for (i in 1:howMany){
33   yleft <- y(xleft)
34   yright <- y(xright)
35   accumulated_area <- accumulated_area + (yleft+yright)*deltax/2
36   xleft <- xleft + deltax
37   xright <- xright + deltax
38 }
39 #####
40 # Report Result
41 #####
42 message("Approximate value of integral from ",xlow, " to ",xhigh, " is: ",accumulated_area)

```

```

> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/TrapezoidalPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 4
Approximate value of integral from 0 to 2 is: 3.45673099542291
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/TrapezoidalPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 400
Approximate value of integral from 0 to 2 is: 3.39345293109042
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/TrapezoidalPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 4000
Approximate value of integral from 0 to 2 is: 3.39344669218516
>

```

Figure 46. Trapezoidal panel example using 4, 400, and 4000 panels..

The actual listing depicted in Figure 46 is shown in Listing 10. Observe (at least for this example) the method appears more accurate than the rectangular method for the same number of panels, however also observe we are making twice as many function calls.

Listing 10. R code demonstrating Trapezoidal Panel Numerical Integration.

```

# R script to implement trapezoidal panel numerical integration
##### NOTE #####
## The interactive input requires the script to be sourced #
# In R console the command line would be #
# source('PATH-TO-THE-FILE/TrapezoidalPanelExample.R') #
# where PATH-TO-THE-FILE is replaced with the actual path on your machine #
#####
# Function to be integrated (modify as needed) #
#####
y <- function(x){
  y <- x * sqrt(1+x^2)
  return(y)
}
#####
# Get lower, upper and how many panels from user #
#####
xlow <- readline("What is the lower limit of integration? ")
xhigh <- readline("What is the upper limit of integration? ")
howMany <- readline("How many panels? ")
# Convert the strings into numeric values
xlow <- as.numeric(unlist(strsplit(xlow, ",")))
xhigh <- as.numeric(unlist(strsplit(xhigh, ",")))
howMany <- as.numeric(unlist(strsplit(howMany, ",")))
# Compute some constants
deltax <- (xhigh - xlow)/howMany
accumulated_area <- 0.0 # initialize the accumulator
xleft <- xlow
xright <- xleft + deltax
#####
# The actual numerical method #
#####
for (i in 1:howMany){
  yleft <- y(xleft)
  yright <- y(xright)
  accumulated_area <- accumulated_area + (yleft+yright)*deltax/2
  xleft <- xleft + deltax
  xright <- xright + deltax
}
#####
# Report Result #
#####
message("Approximate value of integral from ",xlow," to ",xhigh," is: ",accumulated_area)

```

3.1.3 Parabolic Panels

Parabolic panels approximate the shape of the panel with a parabola. The area between the chord and the curve (neglected in the trapezoidal solution) may be accounted for by approximating the function with a parabola passing through the points defined by three successive values of y .

This area may be calculated from the geometry of the parabola and added to the trapezoidal area of the pair of strips to give the area ΔA of the pair as illustrated. Adding all of the ΔA s produces the tabulation shown, which is known as Simpson's rule. To use Simpson's rule, the number n of strips must be even.

The same example as presented for rectangular panels is repeated, except using parabolic panels. The code is changed yet again because we will evaluate at each end of the panel as well as at an intermediate value.

Figure 48 is a screen capture of a parabolic panel integration. The actual script is also listed in Listing 11. In the script, I substituted $\frac{\Delta x}{2}$ for Δx from Figure 47, so the accumulation line has a 6 in the denominator (rather than the 3 in the figure).¹¹

Observe that the estimated integral for 400 and 4000 panels is nearly the same,

¹¹ ... $\frac{\Delta x}{2} \times \frac{1}{3} = \dots \frac{\Delta x}{6}$




Figure 47. Parabolic Panel Schematic..

```

21 xlow <- as.numeric(unlist(strsplit(xlow, ",")))
22 xhigh <- as.numeric(unlist(strsplit(xhigh, ",")))
23 howMany <- as.numeric(unlist(strsplit(howMany, ",")))
24 # Compute some constants
25 deltax <- (xhigh - xlow)/howMany
26 accumulated_area <- 0.0 # initialize the accumulator
27 xleft <- xlow
28 xmiddle <- xleft + deltax/2
29 xright <- xleft + deltax
30 #####
31 # The actual numerical method #
32 #####
33 for (i in 1:howMany){
34   yleft <- y(xleft)
35   ymiddle <- y(xmiddle)
36   yright <- y(xright)
37   accumulated_area <- accumulated_area + (yleft+4*ymiddle+yright)*deltax/6
38   xleft <- xright
39   xmiddle <- xleft + deltax/2
40   xright <- xleft + deltax
41 }
42 #####
43 # Report Result #
44 #####
45 message("Approximate value of integral from ",xlow," to ",xhigh," is: ",accumulated_area)

```

```

> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/ParabolicPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 4
Approximate value of integral from 0 to 2 is: 3.3933797287008
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/ParabolicPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 400
Approximate value of integral from 0 to 2 is: 3.39344662916564
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/ParabolicPanelExample.R')
What is the lower limit of integration? 0
What is the upper limit of integration? 2
How many panels? 4000
Approximate value of integral from 0 to 2 is: 3.39344662916593
>

```

Figure 48. Parabolic panel example using 4, 400, and 4000 panels.

suggesting no need to go beyond a certain number of panels. Algorithms that detect when to stop adding panels exist and would be implemented in many scientific and engineering programming applications.

Listing 11. R code demonstrating Parabolic Panel Numerical Integration.

```
# R script to implement trapezoidal panel numerical integration
##### NOTE #####
## The interactive input requires the script to be sourced #
# In R console the command line would be #
# source('/PATH-TO-THE-FILE/RectangularPanelExample.R') #
# where PATH-TO-THE-FILE is replaced with the actual path on your machine #
#####
# Function to be integrated (modify as needed) #
#####
y <- function(x){
  y <- x * sqrt(1+x^2)
  return(y)
}
#####
# Get lower, upper and how many panels from user #
#####
xlow <- readline("What is the lower limit of integration? ")
xhigh <- readline("What is the upper limit of integration? ")
howMany <- readline("How many panels? ")
# Convert the strings into numeric values
xlow <- as.numeric(unlist(strsplit(xlow, ",")))
xhigh <- as.numeric(unlist(strsplit(xhigh, ",")))
howMany <- as.numeric(unlist(strsplit(howMany, ",")))
# Compute some constants
deltax <- (xhigh - xlow)/howMany
accumulated_area <- 0.0 # initialize the accumulator
xleft <- xlow
xmmiddle <- xleft + deltax/2
xright <- xleft + deltax
#####
# The actual numerical method #
#####
for (i in 1:howMany){
  yleft <- y(xleft)
  ymiddle <- y(xmiddle)
  yright <- y(xright)
  accumulated_area <- accumulated_area + (yleft+4*ymiddle+yright)*deltax/6
  xleft <- xright
  xmmiddle <- xleft + deltax/2
  xright <- xleft + deltax
}
#####
# Report Result #
#####
message("Approximate value of integral from ",xlow," to ",xhigh," is: ",accumulated_area)
```

If we study all the forms of the numerical method we observe that the numerical integration method is really the sum of function values at specific locations in the interval of interest, with each value multiplied by a specific weight. In this development the weights were based on polynomials, but other methods use different weighting functions. An extremely important method is called Gaussian quadrature, which is outside the scope of the discussion herein — Gaussian quadrature routines are readily available within **R**. The method is valuable because one can approximate convolution integrals quite effectively using quadrature routines, while the number of function evaluations for a polynomial based approximation could become hopeless.

When the function values are tabular, we are going to have to accept the rectangular (with adaptations) and trapezoidal as our best tool to approximate an integral because we don't have any really effective way to evaluate the function between the tabulated values – if we were to use our interpolation routine from earlier, it's really going to be a kind of trapezoidal rule anyway.

3.2 Exercise Set 2

1. Write a script to approximate $\int_{1.8}^{3.4} e^x dx$ using rectangular panels.
Run your script using 6 and 600 panels.
 - (a) What is the analytical solution (e.g. do the calculus!)?
 - (b) What is the percent error between the analytical solution and the approximation using 6 panels?
 - (c) What is the percent error between the analytical solution and the approximation using 600 panels?
2. Write a script to approximate $\int_{1.8}^{3.4} e^x dx$ using trapezoidal panels.
Run your script using 6 and 600 panels.
 - (a) What is the analytical solution (e.g. do the calculus!)?
 - (b) What is the percent error between the analytical solution and the approximation using 6 panels?
 - (c) What is the percent error between the analytical solution and the approximation using 600 panels?
3. Based on the previous two exercises, which method do you think is more accurate for a given panel count? Why (do you think so)?
4. Write a script to approximate $\int_0^1 \ln(x) dx$ using rectangular panels.
Run your script using 6 and 600 panels.
 - (a) What is the analytical solution (e.g. do the calculus!)?
 - (b) What is the percent error between the analytical solution and the approximation using 6 panels?
 - (c) What is the percent error between the analytical solution and the approximation using 600 panels?
5. Write a script to approximate $\int_0^1 \ln(x) dx$ using trapezoidal panels.
Run your script using 6 and 600 panels.
 - (a) Did you get an error message — why?

This exercise set is also located on the class server as ES-2

3.3 Numerical Integration of Tabular Data

This subsection is going to work with tabular data — different from function evaluation, but similar. To be really useful, we need to learn how to read data from a file — manually entering tabular data is really time consuming, error prone, and just plain idiotic.

So in this subsection we will first learn how to read data from a file into a list, then we can process the list as if it were a function and integrate its contents.

3.3.1 Reading from a file – open, read, close files

R can read from an ASCII file (or even an Excel .csv file) using a multitude of methods. Common methods are `read.table(...)`, `read.csv(...)`, `read.csv2(...)`, `read.delim(...)`, and `read.xls(...)`.

One can also use primitives¹² to read individual rows in a file and process them.¹³

First, lets create a file named `MyFile.txt`. The extension is important so that we can examine the file with other tools (a text editor) and remember that it is an ASCII file. The contents of `MyFile.txt` are:

```
1 , 1
2 , 4
3 , 9
4 , 16
5 , 25
```

To read the contents into an **R** script we have to do the following:

1. Open a connection to the file — this is a concept common to all languages, it might be called something different, but the program needs to somehow know the location and name of the file.
2. Read the contents into an object — we have a lot of control on how this gets done, for the time being we won't exercise much control yet. When you do substantial programs, you will depend on the control of the reads (and writes).
3. Disconnect the file — this too is common to all languages. Its a really easy step to forget. Not a big deal if the program ends as planned but terrible if there is a error in the program and the connection is still open. Usually noting bad happens, but with an open connection it is possible for the file to get damaged. If that file represents millions of customers data, that's kind of a problem.

¹²Jargon to describe lower level tools within **R**

¹³We will use this approach later in the book – the interactive prompt and reads in the prior subsection are similar to this approach where input is read into a string, then the string is converted into the appropriate type of object (numeric or text).

The `read.table` class of functions handles all three of these steps for us, we do have to provide the filename and some information about the file structure. Later when we are doing network simulation and other hydraulic techniques, different parts of an input file will be read line-by-line and processed — for this task we will need to handle these three steps using primitives.

Figure 49 illustrates the process. The input file has 5 lines, these get read then echoed (printed) back to us. The actual script is pretty simple, notice how the `filepath` and `filename` character variables are defined, then pasted together to produce a full *absolute* file name.¹⁴

Listing 12 is a listing of the script used in Figure 49. The analyst should be able to deduce that the filenames could be read from user input using the prompting technique used in the earlier subsections, so if one is going to process a lot of similar files the explicit naming could be replaced with variable naming – it would probably be a good idea to confine the files to a reasonably memorable path.

Listing 12. R code demonstrating Reading from a File.

```
# R script to illustrate reading from a file using read.table
# The script is intentionally complicated to illustrate the
# three steps: open connection, read into object, close connection
filepath <- "~/Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-
Differentiation/RScripts"
filename <- "MyFile.txt"
fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
# Here we open the connection to the file (within read.table)
# Then the read.table attempts to read the entire file into an object named zz
# Upon either fail or success, read.table closes the connection
zz <- read.table(fileToRead,header=FALSE,sep=",") # comma seperated ASCII, No header
# Echo zz
print(zz)
```

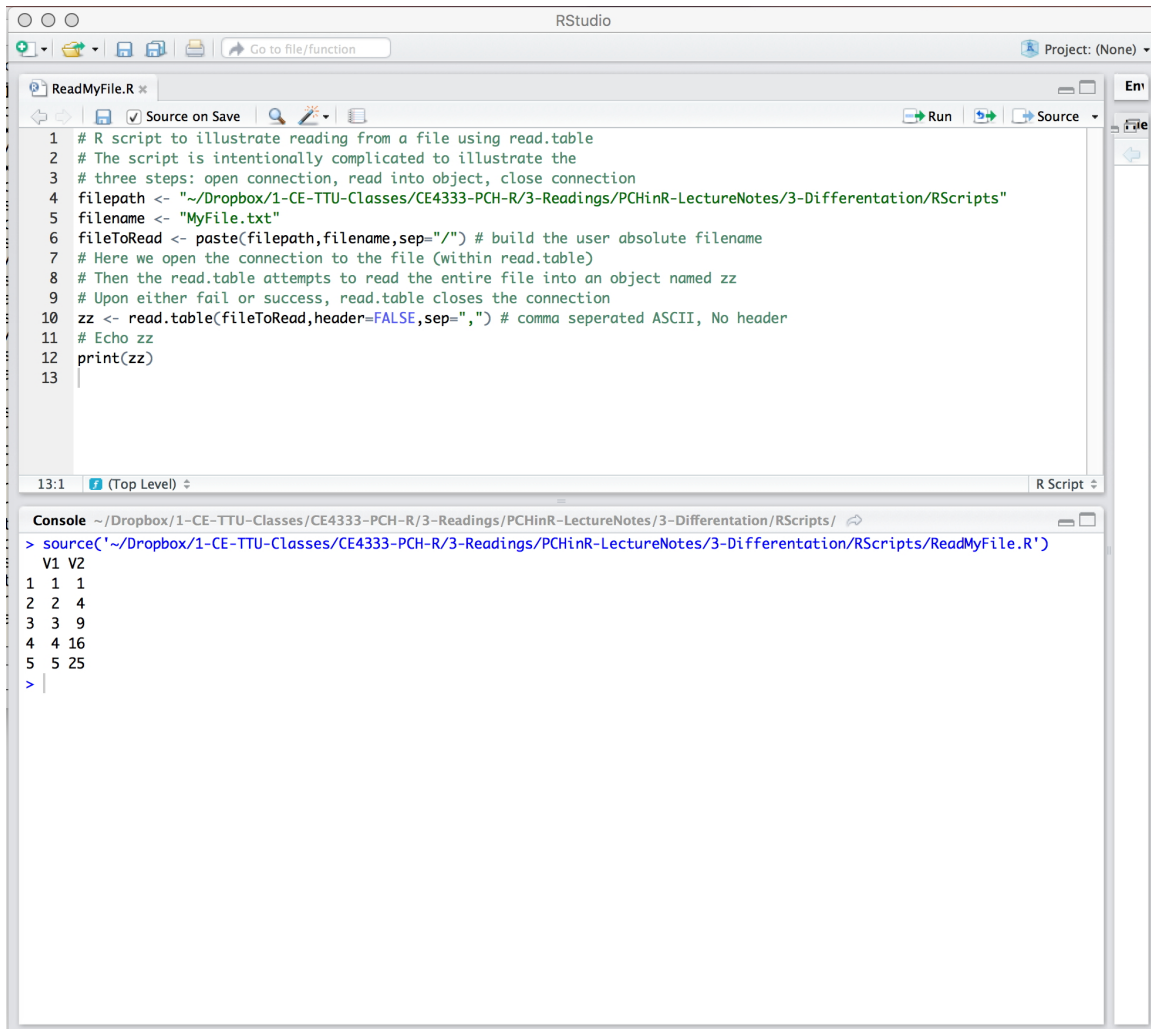
Now that we can read a file, we are now able to integrate tabular data.

3.4 Integrating tabular data

Suppose instead of a function we only have tabulations and wish to estimate the area under the curve represented by the tabular values. Then our integration rules from the prior sections still work more or less, except the rectangular panels will have to be shifted to either the left edge or right edge of a panel (where the tabulation exists).

Lets just examine an example. Suppose some measurement technology produced Table 1 a table of related values. The excitation variable is x and $f(x)$ is the response.

¹⁴A file name can specify all the directory names starting from the root of the tree; then it is called an absolute file name. Or it can specify the position of the file in the tree relative to a default directory; then it is called a relative file name. On the computer I used to write this workbook, the symbol `~`, is the root to my user account, then the remaining directories from that location are explicitly listed. The actual absolute name is `/Users/cleveland/Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts`



```

1 # R script to illustrate reading from a file using read.table
2 # The script is intentionally complicated to illustrate the
3 # three steps: open connection, read into object, close connection
4 filepath <- "~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts"
5 filename <- "MyFile.txt"
6 fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
7 # Here we open the connection to the file (within read.table)
8 # Then the read.table attempts to read the entire file into an object named zz
9 # Upon either fail or success, read.table closes the connection
10 zz <- read.table(fileToRead,header=FALSE,sep=",") # comma separated ASCII, No header
11 # Echo zz
12 print(zz)
13

```

```

> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts/ReadMyFile.R")
  V1 V2
1 1  1
2 2  4
3 3  9
4 4 16
5 5 25
>

```

Figure 49. Rudimentary file reading.

Table 1. Tabular values of an excitation–response relationship.

x	$f(x)$
1.0	1.543
1.1	1.668
1.2	1.811
1.3	1.971
1.4	2.151
1.5	2.352
1.6	2.577
1.7	2.828
1.8	3.107

To integrate this table using the trapezoidal method is straightforward. We will modify our earlier code to read the table (which we put into a file), and compute the integral.

Figure 50 is a screen capture of a script that implements the file read and the numerical integration. The conversion of the method from the functional form in the previous section is pretty straightforward. The main nuisance here is the syntax required to access the “x” values and the “y” values. In **R** the most generic approach is `object$name`, where *object* is the data frame name, and *name* is the variable (column) name. If you don’t use headers, **R** assigns names as V_1, V_2, \dots, V_{max} .

```

1 # R script to illustrate reading from a file using read.table
2 # The script is intentionally complicated to illustrate the
3 # three steps: open connection, read into object, close connection
4 filepath <- "~/Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts"
5 filename <- "MyTableOfData.txt"
6 fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
7 # Here we open the connection to the file (within read.table)
8 # Then the read.table attempts to read the entire file into an object named zz
9 # Upon either fail or success, read.table closes the connection
10 zz <- read.table(fileToRead,header=FALSE,sep=" ") # comma separated ASCII, No header
11 # Echo zz
12 print(zz)
13 # Trapezoidal Integration
14 # Compute some constants
15 howMany <- length(zz$V1) # get the row count from one of the columns
16 accumulated_area <- 0.0 # initialize the accumulator
17 #####
18 # The actual numerical method
19 #####
20 for (i in 1:(howMany-1)){#notice the we go to howMany minus 1
21   deltax <- zz$V1[i+1] - zz$V1[i] # compute the delta x this panel
22   yleft <- zz$V2[i]
23   yright <- zz$V2[i+1]
24   accumulated_area <- accumulated_area + (yleft+yright)*deltax/2
25 }
26 #####
27 # Report Result #
28 #####
29 message("Approximate value of integral from tabular data is: ",accumulated_area)

```

```

V1 V2
1 1.0 1.543
2 1.1 1.668
3 1.2 1.811
4 1.3 1.971
5 1.4 2.151
6 1.5 2.352
7 1.6 2.577
8 1.7 2.828
9 1.8 3.107
Approximate value of integral from tabular data is: 1.7683
>

```

Figure 50. Integrating tabular data..

Realistically the only other simple integration method for tabular data is the rectangular rule, either using the left edge of a panel or the right edge of a panel (and you could do both and average the result which would result in the same outcome as the trapezoidal method). For the sake of completeness lets do both and then compare the results from all four approaches (trapezoidal, rectangular-left, rectangular-right, average rectangular).

First, Figure 51 implements the file read and tabular integration using the rectangular panel method, evaluating the function at the left edge of each panel.

```

1 # R script to illustrate reading from a file using read.table
2 # The script is intentionally complicated to illustrate the
3 # three steps: open connection, read into object, close connection
4 filepath <- "~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts"
5 filename <- "MyTableOfData.txt"
6 fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
7 # Here we open the connection to the file (within read.table)
8 # Then the read.table attempts to read the entire file into an object named zz
9 # Upon either fail or success, read.table closes the connection
10 zz <- read.table(fileToRead,header=FALSE,sep=" ") # comma separated ASCII, No header
11 # Echo zz
12 print(zz)
13 # Rectangular Left
14 # Compute some constants
15 howMany <- length(zz$V1) # get the row count from one of the columns
16 accumulated_area <- 0.0 # initialize the accumulator
17 #####
18 # The actual numerical method #
19 #####
20 for (i in 1:(howMany-1)){ #notice the we go to howMany minus 1
21   deltax <- zz$V1[i+1] - zz$V1[i] # compute the delta x this panel
22   yleft <- zz$V2[i]
23   accumulated_area <- accumulated_area + yleft*deltax
24 }
25 #####
26 # Report Result #
27 #####
28 message("Approximate value of integral from tabular data is: ",accumulated_area)
29
5:32 (Top Level) R Script

```

```

Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts/
  V1  V2
1 1.0 1.543
2 1.1 1.668
3 1.2 1.811
4 1.3 1.971
5 1.4 2.151
6 1.5 2.352
7 1.6 2.577
8 1.7 2.828
9 1.8 3.107
Approximate value of integral from tabular data is: 1.6901
>

```

Figure 51. Integrating tabular data. Rectangular panel, evaluate at left edge..

Next, Figure 52 implements the file read and tabular integration using the rectangular panel method, evaluating the function at the left edge of each panel.

Now lets compare the results from using the three (four) approaches. Table 2 are the results by method.

Table 2. Comparison of tabular integration.

Method	Computed Area
Trapezoidal Panels	1.7683
Rectangular - Left Edge	1.6901
Rectangular - Right Edge	1.8465
Arithmetic Mean Rectangular	1.7683

What Table 2 illustrates is that the trapezoidal rule is simply the average of the rectangular rule evaluated at first the left-edge then the right-edge of a panel.

```

1 # R script to illustrate reading from a file using read.table
2 # The script is intentionally complicated to illustrate the
3 # three steps: open connection, read into object, close connection
4 filepath <- "~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts"
5 filename <- "MyTableOfData.txt"
6 fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
7 # Here we open the connection to the file (within read.table)
8 # Then the read.table attempts to read the entire file into an object named zz
9 # Upon either fail or success, read.table closes the connection
10 zz <- read.table(fileToRead,header=FALSE,sep=" ") # comma separated ASCII, No header
11 # Echo zz
12 print(zz)
13 # Trapezoidal Integration
14 # Compute some constants
15 howMany <- length(zz$V1) # get the row count from one of the columns
16 accumulated_area <- 0.0 # initialize the accumulator
17 - #####
18 #           The actual numerical method           #
19 - #####
20 for (i in 1:(howMany-1)){ #notice the we go to howMany minus 1
21   deltax <- zz$V1[i+1] - zz$V1[i] # compute the delta x this panel
22   yright <- zz$V2[i+1]
23   accumulated_area <- accumulated_area + yright*deltax
24 }
25 - #####
26 #           Report Result           #
27 - #####
28 message("Approximate value of integral from tabular data is: ",accumulated_area)
29

```

```

23:55 (Untitled) R Script
Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts/
  V1  V2
1 1.0 1.543
2 1.1 1.668
3 1.2 1.811
4 1.3 1.971
5 1.4 2.151
6 1.5 2.352
7 1.6 2.577
8 1.7 2.828
9 1.8 3.107
Approximate value of integral from tabular data is: 1.8465
>

```

Figure 52. Integrating tabular data. Rectangular panel, evaluate at right edge..

3.5 Numerical Differentiation

Similar in context to numerical integration is approximation of derivatives. If the functions are representable as functions, then differencing degenerates into the selection of an appropriate difference formula. If the function is tabular, the same decision is presented, but we have to pay additional attention to the quantity of observations available.

3.6 Difference Approximations for Tabulated Data

Here we will introduce differencing by an example. Suppose we want to convert a cumulative data series into an incremental data series. It is operationally related to numerical differentiation.

As an example (leading to an algorithm) consider the cumulative rainfall time series in Table 3.

We shall import this data into **R**, then plot the data, then construct a computational procedure to extract the incremental values from the cumulative values. To load the data into **R** we start the program and then read the contents of a data file that contains the data into **R**, then we will introduce the plotting tools in **R**. In addition to plotting, we will also learn about headers and attaching an object (which gives access to header names rather than the `object$name` structure).

Table 3. Cumulative Rainfall Time Series.

hours	cumulative_rain
0.0	0.00
0.5	1.06
1.0	2.99
1.5	4.80
2.0	4.80
2.5	4.80
3.0	4.80
3.5	4.80
4.0	4.80
4.5	4.80
5.0	4.80
5.5	4.80

The plot suggests that the values are accumulated at the end of the time interval, thus the value accumulated is some average “rate” multiplied by the time interval. The line segment between each point is called a “secant” line. The slope of each secant line, provides that average “rate”. So a fundamental computational step will be a function that computes the slope given any two points (assumed to be adjacent — so that’s why sorting can become important, although the program doesn’t actually care).

3.6.1 Slope of a Secant Line

`slopeOfSecant` is a prototype function that we write to that computes the slope of the secant line through two known points on a function $f(x_1)$ and $f(x_2)$. The function could be tabular or evaluated. The script assumes tabular in that the function is evaluates external to `slopeOfSecant`.

Listing 13. R code demonstrating the prototype function `slopeOfSecant`.

```
##### slope function prototype #####
slopeOfSecant<-function(f1,f2,x1,x2){
  slopeOfSecant <- (f2-f1)/(x2-x1);
  return(slopeOfSecant)
}
#####
```

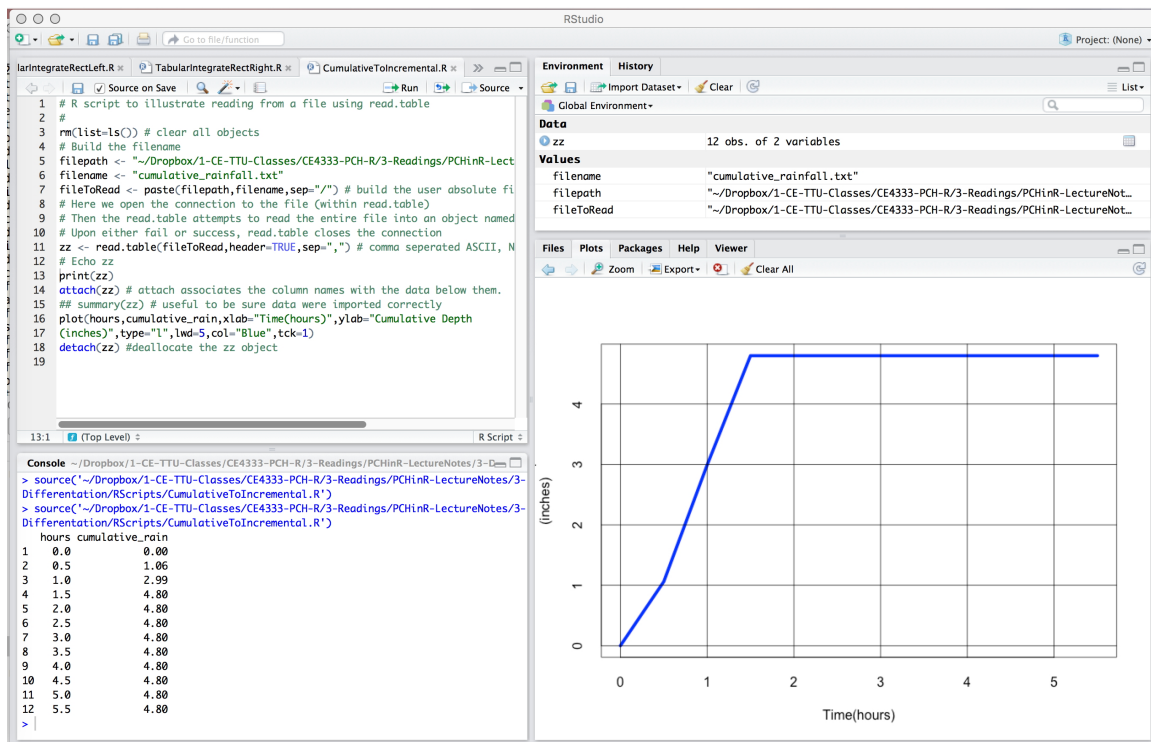



Figure 53. Plot of Cumulative Rainfall Time Series.

This slope is also a first order approximation of a derivative (forward, backward, and centered differences depending on values supplied). This function can then be used to compute “derivatives” of data series using a disaggregation function.

As an illustrative example, if we present parts of the cumulative rainfall data series we can recover the average rate between the inputs. Figure 54 is a screen capture of such a test.

3.6.2 Disaggregation

`disaggregate` is a prototype function that computes the slopes of the secant lines joining adjacent pairs of input data. Depending on the way the input arrays are presented to the `disaggregate` function, the function returns either the backward difference approximation to the function’s derivative or if an index is presented instead of actual t values, then the function returns the incremental values that when aggregated reconstruct the original input function.

Listing 14. R code demonstrating the prototype function `disaggregate()`.

```

#####      disaggregate function prototype      #####
# returns a vector of slopes computed by slopeOfSecant
disaggregate<-function(f,x,dfdx){
  n<-length(x) # length of vectors
  dfdx<-rep(0,n); # zero dfdx
  for (i in 2:n){dfdx[i]<-slopeOfSecant(f[i-1],f[i],x[i-1],x[i]);};
  dfdx[1]<-0;
  return(dfdx)}
#####

```




Figure 54. Script with slopeOfSecant prototype inserted and validation that we recover rate and can re-accumulate correctly.

3.6.3 Numerical Differentiation

A related concept is to determine the average rate for the time interval, the principal difference is that the rate occurs during the entire time interval and should be assigned to the beginning of the interval instead of the end of the interval. A subtle change in the `disaggregate` function can accomplish the task, we will name that new function `brbt`. The name is a mnemonic for “backward-rate, backward-time” differencing.

Listing 15. R code demonstrating the prototype function `brbt()`.

```

##### backward rate, backward time prototype #####
brbt<-function(f,x,dfdx){
  n<-length(x) # length of vectors
  dfdx<-rep(0,n); # zero dfdx
  for (i in 1:(n-1)){dfdx[i]<-slopeOfSecant(f[i],f[i+1],x[i],x[i+1]);};
  dfdx[n]<-0;
  return(dfdx)}
#####

```

Finally, putting everything together, we have the toolkit to determine the incremental rates (which is an approximation to the derivative of the cumulative rates) and incremental depths which are these individual rates multiplied by the length of the time interval. Figure 55 is a screen capture of the **R** script that implements these functions on the tabular data.

Listing 16 is a listing of the **R** script that produced Figure 55. The intermediate steps from Figure 54 is removed in this listing.

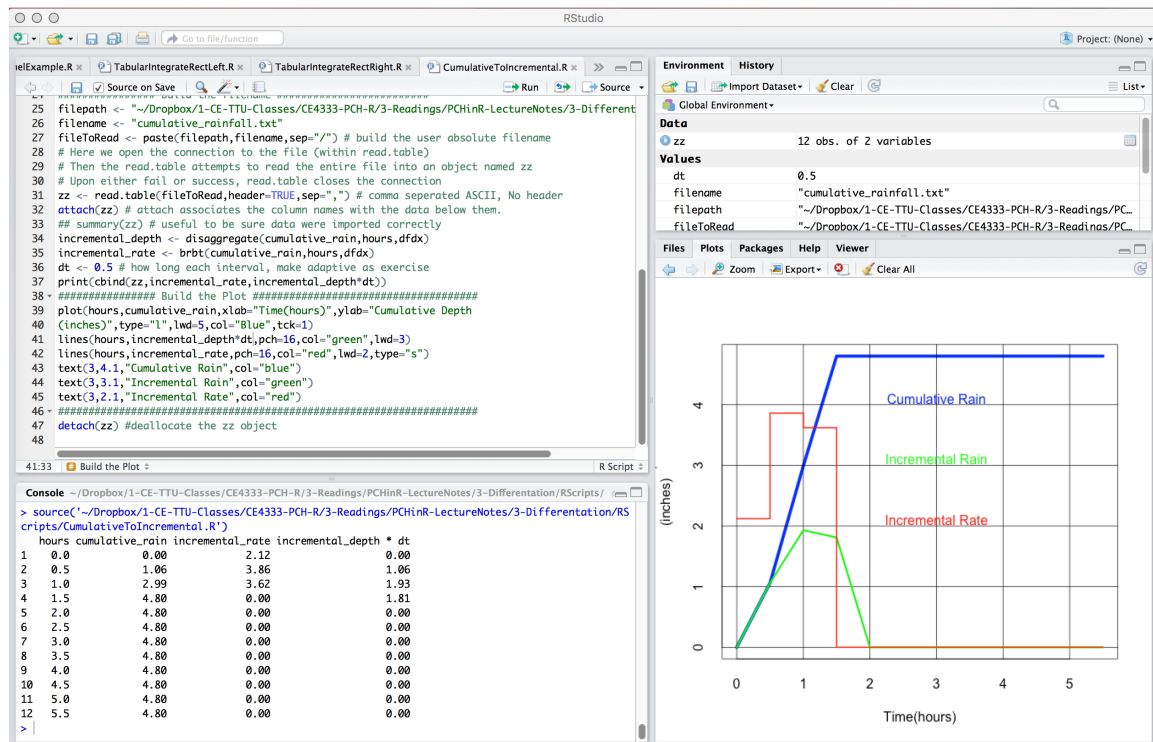


Figure 55. Plot of Cumulative Rainfall Time Series (BLUE), Incremental Depth Time Series (GREEN), and Average Rate Time Series (RED).

Listing 16. R code demonstrating Numerical Differencing.

```
# R script to illustrate numerical differencing
rm(list=ls()) # clear all objects
##### Prototype (forward define) Functions #####
##### slope function prototype #####
slopeOfSecant<-function(f1,f2,x1,x2){
  slopeOfSecant <- (f2-f1)/(x2-x1);
  return(slopeOfSecant)}
##### disaggregate function prototype #####
disaggregate<-function(f,x,dfdx){
  n<-length(x) # length of vectors
  dfdx<-rep(0,n); # zero dfdx
  for (i in 2:n){dfdx[i]<-slopeOfSecant(f[i-1],f[i],x[i-1],x[i]);};
  dfdx[1]<-0;
  return(dfdx)}
##### backward rate, backward time prototype #####
brbt<-function(f,x,dfdx){
  n<-length(x) # length of vectors
  dfdx<-rep(0,n); # zero dfdx
  for (i in 1:(n-1)){dfdx[i]<-slopeOfSecant(f[i],f[i+1],x[i],x[i+1]);};
  dfdx[n]<-0;
  return(dfdx)}
#####
##### Build the filename #####
filepath <- "~/Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/3-Differentiation/RScripts/"
filename <- "cumulative_rainfall.txt"
fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
# Here we open the connection to the file (within read.table)
# Then the read.table attempts to read the entire file into an object named zz
# Upon either fail or success, read.table closes the connection
zz <- read.table(fileToRead,header=TRUE,sep=",") # comma separated ASCII, No header
attach(zz) # attach associates the column names with the data below them.
## summary(zz) # useful to be sure data were imported correctly
incremental_depth <- disaggregate(cumulative_rain, hours, dfdx)
incremental_rate <- brbt(cumulative_rain, hours, dfdx)
dt <- 0.5 # how long each interval, make adaptive as exercise
incremental_depth <- incremental_depth*dt
print(cbind(zz, incremental_rate, incremental_depth))
##### Build the Plot #####
plot(hours, cumulative_rain, xlab="Time(hours)", ylab="Cumulative Depth
(inches)", type="l", lwd=5, col="Blue", tck=1)
```

```

lines(hours, incremental_depth*dt, pch=16, col="green", lwd=3)
lines(hours, incremental_rate, pch=16, col="red", lwd=2, type="s")
text(3, 4.1, "Cumulative Rain", col="blue")
text(3, 3.1, "Incremental Rain", col="green")
text(3, 2.1, "Incremental Rate", col="red")
#####
detach(zz) #deallocate the zz object

```

3.6.4 Aggregation

Aggregation is the compliment of disaggregation; instead of finding differences we are trying to produce cumulatives from incremental values or rates. Aggregation is to integration as disaggregation is to differentiation. Simple aggregation functions are straightforward to build. Numerical integration (already introduced) is a bit more challenging because there are many different ways to compute areas from tabular data – we will illustrated rectangular, trapezoidal, and parabolic panels.

We can insert a prototype `aggregate` function that simply adds elements in a series to prior elements and stores the value in another series. Another name for this kind of arithmetic is a running sum. Functionally, it is rectangular panel (evaluate from the left), numerical integration.

Listing 17. R code demonstrating the prototype function `aggregate()`.

```

##### aggregate function prototype #####
aggregate<-function(vector1,vector2){
n<-length(vector1)
# fill vector2 with zeros
vector2<-rep(0,n)
vector2[1]<-vector1[1]+0.0
for(i in 2:n)vector2[i]<-vector2[i-1]+vector1[i]
return(vector2)}
#####

```

We add this function to the prototype list at the top of the script and can run To illustrate the use of `aggregate` we will aggregate the incremental depths into the cumulative rainfall – we should recover the original cumulative rainfall series that was originally supplied.

3.7 Exercises

1. Add the `aggregate` prototype function to collection of prototype functions in the script. The add some code like:

```

new_cum_rain<-aggregate(incremental_depth,dummy)
plot(hours,cumulative_rain,xlab="Time(hours)",ylab="Cumulative Depth
(inches)",type="l",lwd=5,col="Blue",tck=1)
lines(hours,new_cum_rain,col="red",lwd=1.5)

```

Demonstrate that the two series are identical (e.g. plotting on top of one another).

2. (Advanced) Modify the `disaggregate()` prototype function to automatically determine the time spacing (x) and perform the correct multiplication within the function to return the correct increments. You only have to add one line of code to the prototype function at

```
for (i in 2:n){dfdx[i]<-slopeOfSecant(f[i-1],f[i],x[i-1],x[i]);
deltax <- # you need to define this in terms of x[i] and x[i-1]!
dfdx[i]<- deltax*dfdx[i];};
```

3.8 Finite-Difference Formulas

What we have just done is to explore the use of finite-difference approximations for derivatives. Some common formulas for difference formulas are listed below (without derivation – you should be able to find explanation in any numerical methods text). All the difference equations presented here are the result of truncated Taylor series expansions about x . The “order” refers to the magnitude of truncation error, and this magnitude is proportional to the step size (Δx) raised to a power (the order). Truncation error decreases as the step size is decreased, but one is approaching a divide-by-zero situation (because numerical methods don’t do limits just yet!).

3.8.1 First Derivatives

Equation 8 is a first-order backwards difference.

$$\frac{df}{dx} \approx \frac{f(x) - f(x - \Delta x)}{\Delta x} \quad (8)$$

Equation 9 is a first-order forwards difference.

$$\frac{df}{dx} \approx \frac{f(x + \Delta x) - f(x)}{\Delta x} \quad (9)$$

Equation 10 is a second-order central difference.

$$\frac{df}{dx} \approx \frac{f(x + \Delta x) - f(x - \Delta x)}{2\Delta x} \quad (10)$$

3.8.2 Second Derivatives

Equation 11 is a second-order central difference.

$$\frac{d^2 f}{dx^2} \approx \frac{f(x + \Delta x) - 2f(x) + f(x - \Delta x)}{\Delta x^2} \quad (11)$$

3.8.3 Third Derivatives

Equation 12 is a sixth-order central difference.

$$\frac{d^3 f}{dx^3} \approx \frac{f(x + 2\Delta x) - 2f(x + \Delta x) + 2f(x - \Delta x) + f(x - 2\Delta x)}{2\Delta x^3} \quad (12)$$

The procedure to generate such difference formulas is general and can supply estimates with approximations of any degree. The accuracy depends on the location and number of field variable values involved in the approximation. The selection of a formula is not at all trivial (especially with tabulations), but beyond the scope of this handbook.

Despite known complications, this is a general tool used in computational hydraulics and we will use it throughout the remainder of the handbook – in some examples it will not be obvious that it is finite differencing, and in others it will be explicitly obvious. The next section introduces Newton's method, and finite-differences will be used to approximate the derivative (Quasi-Newton) to implement the method. The procedure is really quite common and imbedded in a lot of the computational tools we use professionally.

3.9 Single Variable Quasi-Newton Methods

The application of fundamental principles of modeling and mechanics often leads to an algebraic or transcendental equation that cannot be easily solved and represented in a closed form. In these cases a numerical method is required to obtain an estimate of the root or roots of the expression.

Newton's method is an iterative technique that can produce good estimates of solutions to such equations. The method is employed by rewriting the equation in the form $f(x) = 0$, then successively manipulating guesses for x until the function evaluates to a value close enough to zero for the modeler to accept.

Figure 56 is a graph of some function whose intercept with the x -axis is unknown. The goal of Newton's method is to find this intersection (root) from a realistic first guess. Suppose the first guess is x_1 , shown on the figure as the right-most specific value of x . The value of the function at this location is $f(x_1)$. Because x_1 is supposed to be a root the difference from the value zero represents an error in the estimate. Newton's method simply provides a recipe for corrections to this error.

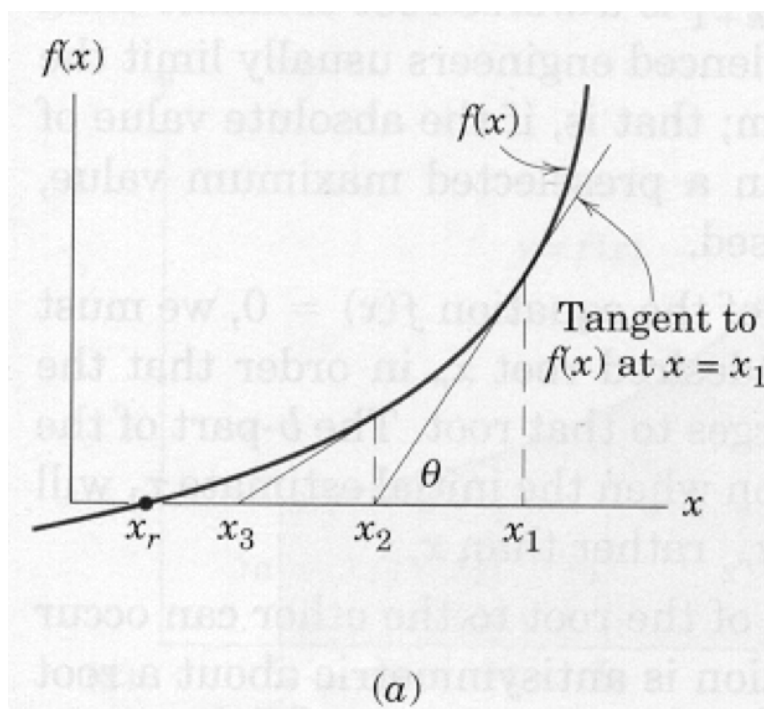


Figure 56. Graph of Arbitrary Function..

Provided x_1 is not near a minimum or maximum (slope of the function is not zero) then a better estimate of the root can be obtained by extending a tangent line from $x_1, f(x_1)$ to the x -axis. The intersection of this line with the axis represents a better estimate of the root.

This new estimate is x_2 . A formula for x_2 can be derived from the geometry of the triangle $x_2, f(x_1), x_1$. Recall from calculus that the tangent to a function at a

particular point is the first derivative of the function. Therefore, from the geometry of the triangle and the definition of tangent we can write,

$$\tan(\theta) = \left. \frac{df}{dx} \right|_{x_1} = \frac{f(x_1)}{x_1 - x_2} \quad (13)$$

Solving the equation for x_2 results in a formula that expresses x_2 in terms of the first guess plus a correction term.

$$x_2 = x_1 - \frac{f(x_1)}{\left. \frac{df}{dx} \right|_{x_1}} \quad (14)$$

The second term on the right hand side is the correction term to the estimate on the right hand side. Once x_2 is calculated we can repeat the formula substituting x_2 for x_1 and x_3 for x_2 in the formula. Repeated application usually leads to one of three outcomes:

1. a root;
2. divergence to $\pm\infty$; or
3. cycling.

These three outcomes are discussed below in various subsections along with some remedies.

The generalized formula is

$$x_{k+1} = x_k - \frac{f(x_k)}{\left. \frac{df}{dx} \right|_{x_k}} \quad (15)$$

If the derivative is evaluated using analytical derivatives the method is called Newton's method, if approximations to the derivative are used, it is called a quasi-Newton method.

3.9.1 Newton's Method — Using analytical derivatives

This subsection presents an example in **R** of implementing Newton's method with analytical derivatives. The algorithm itself is:

1. Write the function in proper form, and code it into a computer.
2. Write the derivative in proper form and code it into a computer.
3. Make an initial guess of the solution (0 and 1 are always convenient guesses).
4. Evaluate the function, evaluate the derivative, calculate their ratio.
5. Subtract the ratio from the current guess and save the result as the update.

6. Test for stopping:
 - (a) Did the update stay the same value? Yes, then stop, probably have a solution.
 - (b) Is the function nearly zero? Yes, then stop we probably have a solution.
 - (c) Have we tried too many updates? Yes, then stop the process is probably cycling, stop.
7. If stopping is indicated proceed to next step, otherwise proceed back to step 4.
8. Stopping indicated, report last update as the result (or report failure to find solution), and related information about the status of the numerical method.

The following example illustrates these step as well as a **R** implementation of Newton's method.

Suppose we wish to find a root (value of x) that satisfies Equation 16.

$$f(x) = e^x - 10\cos(x) - 100 \quad (16)$$

Then we will need to code it into a script. Here is a code fragment that will work:

Listing 18. R code fragment for the function calculation.

```
# Define Function Here
func <- function(x)
{
  func <- exp(x)-10*cos(x)-100;
  return(func);
}
```

The next step is to code the derivative. In this case, Equation 17 is the derivative of Equation 16.

$$\frac{df}{dx}(x) = e^x + 10\sin(x) \quad (17)$$

A code fragment to compute the value of the derivative at any value of x that will work is:

Listing 19. R code fragment for the derivative calculation.

```
# Define Derivative Here
dfdx <- function(x)
{
  dfdx <- exp(x) + 10*sin(x);
  return(dfdx);
}
```

Next we will need script to read in an initial guess, and ask us how many trials we will use to try to find a solution, as well as how close to zero we should be before we declare victory.

Listing 20. R code fragment for reading input data from the programmer.

```
# Read some values from the console
message('Enter an initial guess for X for Newton method : ')
xnow <- as.numeric(readline())
message('Enter iteration maximum : ')
HowMany <- as.numeric(readline())
message('Enter a tolerance value for stopping (e.g. 1e-06) : ')
HowSmall <- as.numeric(readline())
## There are several other ways to make these reads! The scan() function would probably
   also work.
```

The use of `HowSmall`; is called a zero tolerance. We will use the same numerical value for two tolerance tests. Also notice how we are using error traps to force numeric input. Probably overkill for this example, but we already wrote the code in an earlier chapter, so might as well use the code. Professional codes do a lot of error checking before launching into the actual processing — especially of the processing part is time consuming, its worth the time to check for obvious errors before running for a few hours then at some point failing because of an input value error that was predictable.

Now back to the tolerance tests. The first test is to determine if the update has changed or not. If it has not, we may not have a correct answer, but there is no point continuing because the update is unlikely to move further. The test is something like

IF $|x_{k+1} - x_k| < \text{Tol}$. THEN Exit and Report Results

The second test is if the function value is close to zero. The structure of the test is similar, just an different argument. The second test is something like

IF $|f(x_{k+1})| < \text{Tol}$. THEN Exit and Report Results

One can see from the nature of the two tests that a programmer might want to make the tolerance values different. This modification is left as a reader exercise.

Checking for maximum iterations is relatively easy, we just include code that checks for normal exit the loop.¹⁵

Now we simply connect the three fragments, and we have a working **R** script that implements Newton's method for Equation 16. Listing 21 is the entire code module that implements the method, makes the various tests, and reports results. Figure 57 is a screen capture of the program run in **R**.

The example is specific to the particular function provided, but the programmer could move the two functions `func` and `dfdx` into a user specified module, and then load that module in the program to make it even more generic. The next section will use such an approach to illustrate the ability to build a generalized Newton method and only have to program the function itself.

¹⁵Rather than breaking from the loop.

Listing 21. R code demonstrating Newton's Method calculations.

```

# Newtons Method in R
# Define Function Here
func <- function(x)
{
  func <- exp(x)-10*cos(x)-100;
  return(func);
}
# Define Derivative Here
dfdx <- function(x)
{
  dfdx <- exp(x) + 10*sin(x);
  return(dfdx);
}
# Newton's Method Here
# Read some values from the console
message('Enter an initial guess for X for Newton method : ')
xnow <- as.numeric(readline())
message('Enter iteration maximum : ')
HowMany <- as.numeric(readline())
message('Enter a tolerance value for stopping (e.g. 1e-06) : ')
HowSmall <- as.numeric(readline())
# Now start the iterations
for (i in 1:HowMany) {
  xnew <- xnow - func(xnow)/dfdx(xnow)
  # test for stopping
  if (abs(xnew-xnow) < HowSmall){
    message('Update not changing')
    xnow <- xnew
    print(cbind(xnow,xnew,func(xnew)))
    break
  }
  if (abs(func(xnew) < HowSmall)) {
    message('Function value close to zero')
    xnow <- xnew
    print(cbind(xnow,xnew,func(xnew)))
    break
  }
}
# next iteration
xnow <- xnew
}
if (i >= HowMany){
  message('Iteration limit reached')
  print(cbind(xnow,xnew,func(xnew)))
}

```

3.9.2 Newton's Method — Using Finite-Differences to estimate derivatives

A practical difficulty in using Newton's method is determining the value of the derivative in cases where differentiation is difficult. In these cases we can replace the derivative by a difference equation and then proceed as in Newton's method.

Recall from calculus that the derivative was defined as the limit of the difference quotient:

$$\left. \frac{df}{dx} \right|_x = \lim_{\Delta x \rightarrow 0} \frac{f(x + \Delta x) - f(x)}{\Delta x} \quad (18)$$

A good approximation to the derivative should be possible by using this formula with a small, but non-zero value for Δx .

$$\left. \frac{df}{dx} \right|_x \approx \frac{f(x + \Delta x) - f(x)}{\Delta x} \quad (19)$$

When one replaces the derivative with the difference formula the root finding method the resulting update formula is

```

NewtonExample.R x
Source on Save
1 # Newtons Method in R
1:1 (Top Level)

Console ~/Desktop/CE4333-PCH-R/
> source('~~/Desktop/CE4333-PCH-R/NewtonExample.R')
Enter an initial guess for X for Newton method :
1
Enter iteration maximum :
5
Enter a tolerance value for stopping (e.g. 1e-06) :
1e-06
Iteration limit reached
      xnow      xnew
[1,] 6.342875 6.342875 458.4462
> source('~~/Desktop/CE4333-PCH-R/NewtonExample.R')
Enter an initial guess for X for Newton method :
1
Enter iteration maximum :
10
Enter a tolerance value for stopping (e.g. 1e-06) :
1e-06
Iteration limit reached
      xnow      xnew
[1,] 4.593211 4.593211 0.0001774045
>
> source('~~/Desktop/CE4333-PCH-R/NewtonExample.R')
Enter an initial guess for X for Newton method :
1
Enter iteration maximum :
20
Enter a tolerance value for stopping (e.g. 1e-06) :
1e-06
Function value close to zero
      xnow      xnew
[1,] 4.593209 4.593209 1.944613e-10
>

```

Figure 57. Several runs of the program using the analytical derivative to illustrate different kinds of responses..

$$x_{k+1} = x_k - \frac{f(x_k)\Delta x}{f(x_k + \Delta x) - f(x_k)} \quad (20)$$

This root-finding method is called a quasi-Newton method.

Listing 22 is the code fragment that we change by commenting out the analytical

derivative and replacing it with a first-order finite difference approximation of the derivative. The numerical value $1e - 06$ is called the step size (Δx) and should be an input value (rather than built-in to the code as shown here) like the tolerance test values, and be passed to the function as another argument.

Listing 22. R code demonstrating Newton's Method calculations.

```
# Define Derivative Here
dfdx <- function(x)
{
  # dfdx <- exp(x) + 10*sin(x);
  dfdx <- (func(x + 1e-06) - func(x)) / (1e-06);
  # func must already exist before first call!
  return(dfdx);
}
```

Starting with the last example lets modify the analytical version of the code by inserting the above fragment in place of the analytical derivative. Listing 23 is the listing with the modification in place. Notice we have only changed a single line, and not have a more flexible tool. The next modification (left as an exercise) is to detach the creation of the function from the main algorithm, then we would have a general purpose Quasi-Newton's method.

Listing 23. R code demonstrating Newton's Method calculations using finite-difference approximation for the derivative.

```
# Newtons Method in R
# Define Function Here
func <- function(x)
{
  func <- exp(x)-10*cos(x)-100;
  return(func);
}
# Define Derivative Here
dfdx <- function(x)
{
  # dfdx <- exp(x) + 10*sin(x);
  dfdx <- (func(x + 1.0e-06) - func(x))/(1.0e-06)
  return(dfdx);
}
# Newton's Method Here
# Read some values from the console
message('Enter an initial guess for X for Newton method : ')
xnow <- as.numeric(readline())
message('Enter iteration maximum : ')
HowMany <- as.numeric(readline())
message('Enter a tolerance value for stopping (e.g. 1e-06) : ')
HowSmall <- as.numeric(readline())
# Now start the iterations
for (i in 1:HowMany) {
  xnew <- xnow - func(xnow)/dfdx(xnow)
  # test for stopping
  if (abs(xnew-xnow) < HowSmall){
    message('Update not changing')
    xnow <- xnew
    print(cbind(xnow,xnew,func(xnew)))
    break
  }
  if (abs(func(xnew) < HowSmall)) {
    message('Function value close to zero')
    xnow <- xnew
    print(cbind(xnow,xnew,func(xnew)))
    break
  }
}
# next iteration
xnow <- xnew
}
if (i >= HowMany){
  message('Iteration limit reached')
  print(cbind(xnow,xnew,func(xnew)))
}
```

Listing 23 is the main code. Notice how the function definitions are changed, in particular `dfdx`.

Figure 58 is a screen capture of the program run after the code modification above.

```

6   return(func);
7   }
8   # Define Derivative Here
9   dfdx <- function(x)
10  {
11  # dfdx <- exp(x) + 10*sin(x);
12  dfdx <- (func(x + 1.0e-06) - func(x))/(1.0e-06)
13  return(dfdx);
14  }
15  # Newton's Method Here

12:49 dfdx R Script
> source('~/Desktop/CE4333-PCH-R/NewtonExample.R')
Enter an initial guess for X for Newton method :
1
Enter iteration maximum :
20
Enter a tolerance value for stopping (e.g. 1e-06) :
1e-06
Function value close to zero
      xnow      xnew
[1,] 4.593209 4.593209 2.922036e-10
>

```

Figure 58. Program run after changing from analytical to finite-difference approximation for the derivative..

The advantage of the approximate derivative is that we don't have to do the calculus — just code in the function.

The obvious advantage of modular coding is to protect the parts of the code that are static, and just modify the function definitions. We can keep a working example around in case we break something and use that to find what we broke.

3.9.3 Method Fails

The three subsections below describe the ways that the method routinely fails, along with some suggestions for remedy. Generally we should plot the function before trying

to find a root, but sometimes the root finding is a component of a more complex program and we just want it to work. In that situation, the programmer would build in many more tests than the three above to try to force a result before giving up.

3.9.4 Multiple Roots

Figure 59 illustrates the behavior in the presence of multiple roots. When there are multiple roots the method will converge on the root that is defined by the initial guess.¹⁶ The initial estimate must be close enough to the desired root to converge to the root.¹⁷ Another challenge is what happens if the initial guess is at the divide




Figure 59. Multiple roots..

(the peak of the function in Figure 59); in such cases we may actually get a divergent solution because the slope of the function at that peak is nearly zero.

3.9.5 Cycling

Cycling can occur when the root is close to an inflection point of the function. Usual practice is to again limit the step size to prevent such behavior. Figure 60 is an illustration of cycling. A good remedy for cycling is to first detect the cycling, then provide a small “shove” to the guess. Examination of root finding codes often reveals a pseudo-random number generator within the code that will provide this shove when cycling is detected.

¹⁶This behavior is called “sensitive dependence on initial conditions”.

¹⁷Ironically, we need a good idea of the answer before we start the method.




Figure 60. Estimates cycling around a root..

3.9.6 Near-zero derivatives

The derivative in Equation 66 must not be zero, otherwise the guess corresponds to a maximum or minimum of the function and the tangent line will never intersect the x -axis. The derivative must not be too close to zero, otherwise the slope will be so small as to make the correction too large to produce a meaningful update. Usual practice is to limit the size of the correction term to some maximum and to use this maximum value whenever the formula prescribes a larger step. Divergence to $\pm\infty$ is usually explained by near-zero derivatives at the sign change. The bi-section method is a little more robust in this respect.

3.10 Related Concepts

A couple of other root finding methods are worth mentioning because they can sometimes serve as a fallback when Newton's method fails. Two robust methods are bisection and false-positioning. These are discussed in the suggested reading list in Chapra's textbook.

3.11 Exercise Set 3

1. Build a Newton's Method program (or use mine) and make the program request a tolerance value for "how close to zero" is the function, and "how small is the change in update values." Build your code using the modular approach (two files). Test your code using the same example in the notes.
2. Now modify the main and the function module to use approximate derivatives (the finite-difference formulation) and require the user supply a step size. Test the code using the same example in the notes.

For each of the exercises above, prepare documentation similar to the notes where you describe the salient points of your program.

3. Now use your program to find roots for the following equations:

(a) $\exp(x) - 3x^2 = 0$

(b) $\ln(x) - x + 2 = 0$

(c) $\tan(x) - x - 1 = 0$

For these three equations, document your search for roots. Identify if there are bad initial guesses that cause the program to fail to find a root. Also the equations may have multiple roots. If you discover multiple roots, identify the starting values one needs to use to converge to a particular root.

These exercises are also located on the class server in ES-3.

4 Tank Drain Simulation using Finite-Differences

This section examines how to simulate the time to drain a tank. This particular example is a common exercise in fluid mechanics and hydraulic engineering courses to motivate the concept of control volumes and the conservation of mass. The equations of motion are usually given (as they are here) based on an assumption of negligible energy loss as the tank drains¹⁸.

The purpose is to gain some practice with numerical modeling techniques before trying more complex methods in hydraulics.

4.1 Problem Description

Consider the tank depicted in Figure 61. The water depth in the tank at any instant is $z(t)$. The tank cross-section area is A_{tank} and the outlet cross section area is A_{out} . The product of A_{tank} and $z(t)$ is the volume in storage in the tank at some instant, and the outflow from the tank is $Q(t)$.

As a first model consider the computation of

1. Time to drain the tank from a known initial depth.
2. Storage in the tank at some given time.
3. Depth in tank at some given time.
4. Discharge rate at some instant in time.
5. How long to drain the tank to $\frac{1}{2}$ or $\frac{1}{4}$ full.

All these questions are reasonable kinds of questions that could arise in some engineering design situation or (more likely) in an operations scenario. In this example the tank has constant cross section, but it is not a far stretch to imagine that the tank could represent a reservoir of variable geometry. A practical reason to ask such questions arises in storm-water detention pond design (as well as reactor design in wastewater when dynamic flows are considered) — the time to drain impacts the remaining capacity in a detention pond, the remaining capacity is what provides some measure of control for back-to-back storms. If the tank drains too slowly, there will not be enough reserve for a back-to-back storm, whereas if the tank drains too fast it is hydraulically irrelevant.

While this problem seems detached from an open channel flow course, it [the problem] is not. Each node in the open channel equations where flow depth is computed is essentially such a tank. The tank area is not constant (instead the area is related to depth and reach distance), but like this tank the inflows and outflows are, in-part, governed by the depth already in the tank.

¹⁸Bernoulli's equation.




Figure 61. Storage tank with drain on bottom..

4.2 Computational Approach

This example will be handled first analytically (because it is simple to do so), then by a simple finite-difference scheme implemented in **R**

4.3 Problem Analysis — Development of the Analytical Solution

To write an equation or set of equations for this problem, the relationship of tank depth, areas, and discharges must be specified.

If we write Bernoulli's equation along a streamline in the tank from the free surface

to the outlet we would arrive at something like

$$\frac{p_{fs}}{\gamma} + \frac{V_{fs}^2}{2g} + z_{fs} = \frac{p_{out}}{\gamma} + \frac{V_{out}^2}{2g} + z_{out} \quad (21)$$

If we invoke the following reasonable assumptions:

1. $p_{fs} = 0$ Free surface at zero gage pressure.
2. $p_{out} = 0$ Pressure in jet is ≈ 0 .
3. $V_{fs} \ll V_{out}$ Free surface velocity is small relative to the outlet velocity.
4. $z_{out} = 0$ the outlet is the datum.

The resulting relationship between depth, and discharge is

$$Q(t) = A_{out} \times \sqrt{2gz(t)} \quad (22)$$

Now we have an equation of motion. A mass balance in the tank requires that the storage decrease as the tank drains. Thus relating the change in tank depth and discharge produces the ordinary differential equation (ODE) that governs (at least in our model world) the tank.

$$A_{out} \times \sqrt{2gz(t)} = -A_{tank} \times \frac{dz}{dt} \quad (23)$$

If one collects the constants $\frac{-A_{out}}{A_{tank}} \sqrt{2g} = \alpha$ the ODE is a little simpler to examine¹⁹

$$\alpha [z(t)]^{\frac{1}{2}} = \frac{dz}{dt} \quad (24)$$

Equation 24 can be separated and integrated

$$\int_0^T \alpha dt = \int_{z(0)}^0 \frac{dz}{[z(t)]^{\frac{1}{2}}} \quad (25)$$

The solution is

$$z(t) = (z(0)^{\frac{1}{2}} - \frac{\alpha}{2}t)^2 \quad (26)$$

¹⁹The constants can be collected in this case because the geometry stays constant in the tank — changing geometries would have areas as functions of depth and could not be collected in this fashion.

4.3.1 Application of Analytical Solution in R

This section presents the analytical solution coded in Excel, **R**, and **FORTRAN**.

Analytical solutions are usually straightforward to represent in environments like **R**. In this example, a single line of code builds the relationship and a couple more lines for a plot.

```
> depth<-function(alpha,time,initialDepth){(sqrt(initialDepth)-0.5*alpha*time)^2}
> tt<-seq(1,1000)
> plot(tt,depth(alpha,tt,5),xlab="Time (minutes)",ylab="Depth (meters)")
```

The actual “plot” is left as an exercise.

4.4 Problem Analysis — Development of the Finite-Difference Approximation

The finite difference approximation follows the same principles up to Equation 24 then approximates the derivative term by its difference quotient for small time steps²⁰

$$\alpha[z(t)]^{\frac{1}{2}} \approx \frac{\Delta z}{\Delta t} \quad (27)$$

Notice some features of 27. First the equal sign is changed to “approximately” equal; second the derivative is changed to a related rate. The remainder of the “equation” is unchanged. The trick here is to understand how to interpret the equation.

The approximation states that the rate of change of depth with time is approximately equal to the product of the geometric and gravitational constant and the square root of the current depth.

Using this approximation, and knowing the depth in the tank to start produces an algorithm to explore the tank behavior. First expand the difference equation.

$$\alpha[z(t)]^{\frac{1}{2}} = \frac{z(t + \Delta t) - z(t)}{\Delta t} \quad (28)$$

Then rearrange to isolate values at time $t + \Delta t$,

$$z(t + \Delta t) = z(t) + \Delta t \alpha[z(t)]^{\frac{1}{2}} \quad (29)$$

All the terms on the right hand side are known and the equation²¹ tells the modeler

²⁰Recall the fundamental theorem of calculus where the difference quotient is taken to the limit and this limit is called the derivative — here we don’t go to the limit, instead using small but finite steps.

²¹Now known as an update equation

how to approximate depth at a future time. Because the analysis assumed the difference quotient is close to the derivative, the time steps need to be kept pretty small (in fact in many problems the time step is constrained by the physics for stability and by the computation regime for precision).

The particular update here is an implementation of Euler's method²². There are other methods — the reader should observe that the time difference scheme could be backwards so that the discharge could be at an unknown time, or a weighted average of the two times, or a variety of other ways to approximate the derivative.

4.5 Application of the Finite-Difference Approximation in R

This section presents the same set of “solutions” using the finite-difference model. In these cases the underlying algorithm is that expressed by Equation 29.

To model the tank using **R** the modeler needs to write the necessary equation structure in the proper order. Listing 24 is a crude implementation — note how the update quotient is actually built with a test for near zero values to prevent attempted square root of a negative number. Also note how there is some flexibility to change the time step.

Listing 24. R code demonstrating time to drain calculations.

```
> # Time to Drain Model
> AreaTank<-987
> AreaPipe<-1
> alpha<-sqrt(2*9.8)*AreaPipe/AreaTank
> z<-numeric(0) # define the depth array
> z[1]<-5.0 # initial depth
> dt<-5.0 # time step size
> # program the update difference quotient
> dzdt<-function(alpha,depth){if(depth >= 0.001)alpha*sqrt(depth) else 0}
> # update many times
> for (i in 1:500){z[i+1]=z[i]-dt*dzdt(alpha,z[i])}
> length(z) # get length for plotting
[1] 501
> time<-seq(0,500)*dt
> plot(time,z,xlab="Time (minutes)",ylab="Depth (meters)")
```

The result of the plot call is displayed in Figure 62

The plot displays a lot of zero values, the student should explore tricks to plot only the interesting portion of the behavior. The student should also explore how to plot both the analytical and numerical solution on the same graph.

What this section presented is actually quite simple. The engineer in any case will have to conceptualize the physical system into a structure amenable to mathematical representation. The tools are the conservation of mass, momentum, and energy. Then relationships between components must be established — generally time, lengths, and forces are somehow related. These relationships, whether empirical or fundamental constitute the basis to build a model to reality.

²²One of many methods to approximate solutions to ordinary differential equations




Figure 62. Plot of relationship of elapsed time and depth of water in tank..

Next the modeler has to convert these relationships into a structure that can be solved using the tools at hand: in this case **R**. Finally the application is built and used for the problem of interest. While intellectually desirable to have a fairly general tool, the engineer should never be afraid to use a purpose built tool if a general tool is unavailable. Be sure to test the tool and know its limitations.

In the example a simple hydraulics computation was presented to motivate these modeling steps. Solutions were constructed in the **R** programming environments.

4.6 Exercise Set 4

1. Adapt the **R** script and build a simulator that approximates the depth of water in a cylindrical drum lying on its side as a function of time. Figure 63 is a sketch of the drum of interest. Also desired is the volume in storage in the drum as a function of time.




Figure 63. Sketch of drum dimensions.

The drum is drained by a 2-inch diameter short pipe at the bottom of the drum. The velocity of water in the pipe is $V_e = \sqrt{2gh}$ where g is gravitational acceleration and h is water depth in the tank above the outlet. The drum is 4-feet long and 2-feet in diameter. Simulate the time to drain from $\frac{1}{2}$ full to empty, then generalize to any starting depth (up to the tank diameter).

Produce a time versus depth and time versus storage plot for the drum using **R**.

Document your work in a short modeling report that includes conceptualization, problem analysis (development of requisite equations), coding, any testing, and finally the application of the model.

2. Adapt the **R** script and build a simulator for a trough of arbitrary dimensions as shown in Figure 64.

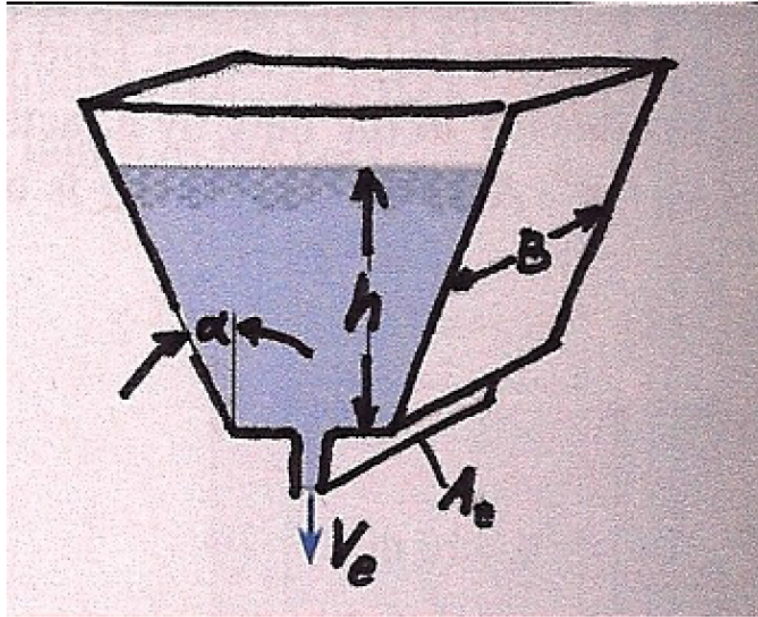


Figure 64. Sketch of trough dimensions.

The angle with the vertical of the sloping sides is α , and the distance between the parallel ends is B . The width of the trough is $W_0 + 2h \tan \alpha$, where h is the distance from the trough bottom. The velocity of water issuing from the opening in the bottom of the trough is equal to $V_e = \sqrt{2 g h}$. The area of the water stream at the bottom of the trough is A_e .

Produce a drainage curve (time versus depth) for the case where $h_0 = 5m$, $W_0 = 1m$, $\alpha = 30^\circ$, $B = 10m$, and $A_e = 1m^2$.

Document your work in a short modeling report that includes conceptualization, problem analysis (development of requisite equations), coding, any testing, and finally the application of the model.

These exercises are also located on the class server as ES-4.

5 Simultaneous Linear Systems of Equations

Many engineering simulations require the solution of simultaneous algebraic equations. These algebraic equation systems are either linear or non-linear in the unknown variables. Many computation schemes have been developed to solve the resulting systems, mostly depending on the structure of the systems (and the corresponding coefficient matrices).

The solution of linear (or non-linear for that matter) can be accomplished using either direct methods or iterative (successive approximation) methods. The method choice depends on:

1. The amount of computation required (size of the problem) and computer memory available.²³
2. The accuracy of the solution required.
3. The ability to control accuracy (i.e. find accurate enough solutions) to improve overall computation speed and throughput.

Direct solution methods lead to results by means of finite and predictable operations count, but at the expense of error amplification and difficulty to deal with near-singular systems. Iterative methods can converge to exact solutions, are robust in near-singular cases, but at the expense of a non-predictable number of operations.

In this chapter we will see how to solve systems using built-in method(s) in **R** and will also see the simplest of the iterative methods, Jacobi iteration. Jacobi iteration is presented for several reasons: it is simple to program, it shows the beauty of iteration when it works, and introduces a concept called pre-conditioning. For problems in this workbook, the built in `solve(...)` is recommended; we will use Jacobi iteration later on the the aquifer flow models, because the model equation structure is quite amenable to this kind of solution method.

For really large systems of equations iterative methods probably dominate because they are quite amenable to out-of-core solution — Jacobi iteration is ideal for parallel processing in a GPU²⁴

²³In the past, the memory was indeed an issue – its less so today; a really big problem of thousands of equations and thousands of variables might indeed be too big for any single computer array and would require out-of-core solver techniques, which I suspect are a slowly dying art.

²⁴Graphics Processing Unit — Nearly all our laptops have GPU; either an Intel, NVIDIA, or AMD. These are intended for rendering graphics, but can be directly accessed with the proper software tools and can perform floating point operations really quickly. For example on my laptop I have an NVIDIA GeForce GT750M which I can program using a CUDA toolkit. If I had a really large system to solve, I would try Jacobi iteration, make each equation a thread, the solution guess a thread, and the update a thread. Its relatively easy to multiply, add, and divide threads, so one could compute the update directly from parallel thread multiplication using the guess, then thread addition to update the guess, and repeat. GPU programming is beyond this handbook, but remember that one can trade efficiency for speed if the operations are simple vector arithmetic.

5.1 Numerical Linear Algebra – Matrix Manipulation

This section introduces use of matrices in **R** to learn how to address particular elements of a matrix – once that is understood, the remaining arithmetic is reasonably straightforward.

5.2 The Matrix — A data structure

Listing 25 is script fragment that reads in two different matrices **A** and **B**, and writes them back to the screen. While such an action alone is sort of meaningless, the code does illustrate how to read the two different files, and write back the result in a row wise fashion.

The two matrices are

$$\mathbf{A} = \begin{pmatrix} 12 & 7 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \quad (30)$$

and

$$\mathbf{B} = \begin{pmatrix} 5 & 8 & 1 & 2 \\ 6 & 7 & 3 & 0 \\ 4 & 5 & 9 & 1 \end{pmatrix} \quad (31)$$

Now that we have a way (albeit pretty arcane) for getting matrices into our program from a file²⁵ we can explore some elementary matrix arithmetic operations, and then will later move on to some more sophisticated operations, ultimately culminating in solutions to systems of linear equations (and non-linear systems in the next chapter).

²⁵The read from a file is a huge necessity — manually entering values will get old fast. I have written matrix generators whose purpose in life is to construct matrices and put them into files for subsequent processing — often these programs are pretty simple because of structure in a problem, at other times they rival the solution tool in complexity; once for a Linear Programming model (circa 1980's) I developed a code to write a 1200 X 1200 matrix to a file, which would be functionally impossible to enter by hand.

Listing 25. R code demonstrating reading in two matrices.

```

# R script for some matrix operations
##### READ IN DATA FROM A FILE #####
filepath <- "~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-
  LinearSystems/RScripts"
filename <- "MatrixA.txt"
fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
# Read the first file
yy <- read.table(fileToRead,header=FALSE,sep=",") # comma seperated ASCII, No header
filename <- "MatrixB.txt" # change the filename
fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
# Read the second file
zz <- read.table(fileToRead,header=FALSE,sep=",") # comma seperated ASCII, No header
##### Get Row and Column Counts #####
HowManyColumnsA <- length(yy)
HowManyRowsA <- length(yy$V1)
HowManyColumnsB <- length(zz)
HowManyRowsB <- length(zz$V1)
##### Build A and B Matrices #####
Amat <- matrix(0,nrow = HowManyRowsA, ncol = HowManyColumnsA)
Bmat <- matrix(0,nrow = HowManyRowsB, ncol = HowManyColumnsB)
for (i in 1:HowManyRowsA){
  for(j in 1:(HowManyColumnsA)){
    Amat[i,j] <- yy[i,j]
  }
}
rm(yy) # deallocate zz and just work with matrix and vectors
for (i in 1:HowManyRowsB){
  for(j in 1:(HowManyColumnsB)){
    Bmat[i,j] <- zz[i,j]
  }
}
rm(zz) # deallocate zz and just work with matrix and vectors
##### Echo Input #####
print(Amat)
print(Bmat)

```

5.3 Matrix Arithmetic

Analysis of many problems in engineering result in systems of simultaneous equations. We typically represent systems of equations with a matrix. For example the two-equation system,

$$\begin{aligned} 2x_1 + 3x_2 &= 8 \\ 4x_1 - 3x_2 &= -2 \end{aligned} \tag{32}$$

Could be represented by set of vectors and matrices²⁶

$$\mathbf{A} = \begin{pmatrix} 2 & 3 \\ 4 & -3 \end{pmatrix} \quad \mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \quad \mathbf{b} = \begin{pmatrix} 8 \\ -2 \end{pmatrix} \tag{33}$$

and the linear system then written as

$$\mathbf{A} \cdot \mathbf{x} = \mathbf{b} \tag{34}$$

²⁶Usually called “vector-matrix” form. Additionally, a vector is really just a matrix with column rank = 1 (a single column matrix).

So the “algebra” is considerably simplified, at least for writing things, however we now have to be able to do things like multiplication (indicated by \cdot) as well as the concept of addition and subtraction, and division (multiplication by an inverse). There are also several kinds of matrix multiplication – the inner product as required by the linear system, the vector (cross product), the exterior (wedge), and outer (tensor) product are a few of importance in both mathematics and engineering.

The remainder of this section will examine the more common matrix operations.

5.3.1 Matrix Definition

A matrix is a rectangular array of numbers.

$$\begin{pmatrix} 1 & 5 & 7 & 2 \\ 2 & 9 & 17 & 5 \\ 11 & 15 & 8 & 3 \end{pmatrix} \quad (35)$$

The size of a matrix is referred to in terms of the number of rows and the number of columns. The enclosing parenthesis are optional above, but become meaningful when writing multiple matrices next to each other. The above matrix is 3 by 4.

When we are discussing matrices we will often refer to specific numbers in the matrix. To refer to a specific element of a matrix we refer to the row number (i) and the column number (j). We will often call a specific element of the matrix, the $a_{i,j}$ -th element of the matrix. For example $a_{2,3}$ element in the above matrix is 17. In **R** we would refer to the element as `a_matrix[i][j]` or whatever the name of the matrix is in the program.

5.3.2 Multiply a matrix by a scalar

A scalar multiple of a matrix is simply each element of the matrix multiplied by the scalar value. Consider the matrix **A** below.

$$\mathbf{A} = \begin{pmatrix} 12 & 7 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \quad (36)$$

If the scalar is say 2, then $2 \times \mathbf{A}$ is computed by doubling each element of **A**, as

$$2\mathbf{A} = \begin{pmatrix} 24 & 14 & 6 \\ 8 & 10 & 12 \\ 17 & 16 & 18 \end{pmatrix} \quad (37)$$

In **R** we can simply perform the arithmetic as

Listing 26. R code demonstrating scalar multiplication.

```
#####
twoA <- 2 * Amat
print(twoA)
```

Figure 65 is an example using the earlier **A** matrix and multiplying it by the scalar value of 2.0.

```
#####
twoA <- 2 * Amat
print(twoA)
```

```

13 HowManyColumnsA <- length(yy)
14 HowManyRowsA <- length(yy$V1)
15 HowManyColumnsB <- length(zz)
16 HowManyRowsB <- length(zz$V1)
17 ##### Build A and B Matrices #####
18 Amat <- matrix(0,nrow = HowManyRowsA, ncol = HowManyColumnsA)
19 Bmat <- matrix(0,nrow = HowManyRowsB, ncol = HowManyColumnsB)
20 for (i in 1:HowManyRowsA){
21   for(j in 1:(HowManyColumnsA)){
22     Amat[i,j] <- yy[i,j]
23   }
24 }
25 rm(yy) # deallocate zz and just work with matrix and vectors
26 for (i in 1:HowManyRowsB){
27   for(j in 1:(HowManyColumnsB)){
28     Bmat[i,j] <- zz[i,j]
29   }
30 }
31 rm(zz) # deallocate zz and just work with matrix and vectors
32 ##### Echo Input #####
33 print(Amat)
34 #print(Bmat)
35 twoA <- 2 * Amat
36 print(twoA)
37
```

```

34:2 Echo Input
R Script

Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RSc
ripts/ReadMatrixNew.R')
  [,1] [,2] [,3]
[1,] 12  7  3
[2,]  4  5  6
[3,]  7  8  9
  [,1] [,2] [,3] [,4]
[1,]  5  8  1  2
[2,]  6  7  3  0
[3,]  4  5  9  1
  [,1] [,2] [,3]
[1,] 24 14  6
[2,]  8 10 12
[3,] 14 16 18
> |
```

Figure 65. Multiply each element in amatrix by a scalar .

5.3.3 Matrix addition (and subtraction)

Matrix addition and subtraction are also element-by-element operations. In order to add or subtract two matrices they must be the same size and shape. This requirement means that they must have the same number of rows and columns. To add or subtract a matrix we simply add or subtract the corresponding elements from each matrix.

For example consider the two matrices \mathbf{A} and $2\mathbf{A}$ below

$$\mathbf{A} = \begin{pmatrix} 12 & 7 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \quad 2\mathbf{A} = \begin{pmatrix} 24 & 14 & 6 \\ 8 & 10 & 12 \\ 17 & 16 & 18 \end{pmatrix} \quad (38)$$

For example the sum of these two matrices is the matrix named $3\mathbf{A}$, shown below:

$$\mathbf{A} + 2\mathbf{A} = \begin{pmatrix} 12 + 24 & 7 + 14 & 3 + 6 \\ 4 + 8 & 5 + 10 & 6 + 12 \\ 7 + 14 & 8 + 16 & 9 + 18 \end{pmatrix} = \begin{pmatrix} 36 & 21 & 9 \\ 12 & 15 & 18 \\ 21 & 24 & 27 \end{pmatrix} \quad (39)$$

Now to do the operation in \mathbf{R} , we need to read in the matrices, perform the addition, and write the result. In the code example in 66 I added a third matrix to store the result – generally we don't want to clobber existing matrices, so we will use the result instead.

Subtraction is performed in a similar fashion, except the subtraction operator is used.

The screenshot shows the RStudio interface with a script editor and a console. The script editor contains R code for building matrices A and B, and performing operations on them. The console shows the output of the script, including the dimensions of the matrices and the results of the operations.

```

14 HowManyRowsA <- length(yy$V1)
15 HowManyColumnsB <- length(zz)
16 HowManyRowsB <- length(zz$V1)
17 ##### Build A and B Matrices #####
18 Amat <- matrix(0,nrow = HowManyRowsA, ncol = HowManyColumnsA)
19 Bmat <- matrix(0,nrow = HowManyRowsB, ncol = HowManyColumnsB)
20 for (i in 1:HowManyRowsA){
21   for(j in 1:(HowManyColumnsA)){
22     Amat[i,j] <- yy[i,j]
23   }
24 }
25 rm(yy) # deallocate zz and just work with matrix and vectors
26 for (i in 1:HowManyRowsB){
27   for(j in 1:(HowManyColumnsB)){
28     Bmat[i,j] <- zz[i,j]
29   }
30 }
31 rm(zz) # deallocate zz and just work with matrix and vectors
32 ##### Echo Input #####
33 print(Amat)
34 #print(Bmat)
35 twoA <- 2 * Amat
36 print(twoA)
37 threeA <- Amat+twoA
38 print(threeA)
39

```

Console output:

```

> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/ReadMatrixNew.R')
      [,1] [,2] [,3]
[1,]  12   7   3
[2,]   4   5   6
[3,]   7   8   9
      [,1] [,2] [,3]
[1,]  24  14   6
[2,]   8  10  12
[3,]  14  16  18
      [,1] [,2] [,3]
[1,]  36  21   9
[2,]  12  15  18
[3,]  21  24  27
>

```

Figure 66. Add each element in A to each element in twoA, store the result in threeA..

5.3.4 Multiply a matrix

One kind of matrix multiplication is an inner product. Usually when matrix multiplication is mentioned without further qualification, the implied meaning is an inner product of the matrix and a vector (or another matrix).

Matrix multiplication is more complex than addition and subtraction. If two matrices such as a matrix \mathbf{A} (size $l \times m$) and a matrix \mathbf{B} (size $m \times n$) are multiplied together, the resulting matrix \mathbf{C} has a size of $l \times n$. The order of multiplication of matrices is extremely important²⁷.

To obtain $\mathbf{C} = \mathbf{A} \mathbf{B}$, the number of columns in \mathbf{A} must be the same as the number of rows in \mathbf{B} . In order to carry out the matrix operations for multiplication of matrices, the i, j -th element of \mathbf{C} is simply equal to the scalar (dot or inner) product of row i of \mathbf{A} and column j of \mathbf{B} .

Consider the example below

$$\mathbf{A} = \begin{pmatrix} 1 & 5 & 7 \\ 2 & 9 & 3 \end{pmatrix} \quad \mathbf{B} = \begin{pmatrix} 3 & -2 \\ -2 & 1 \\ 1 & 1 \end{pmatrix} \quad (40)$$

First, we would evaluate if the operation is even possible, \mathbf{A} has two rows and three columns. \mathbf{B} has three rows and two columns. By our implied multiplication “rules” for the multiplication to be defined the first matrix must have the same number of rows as the second matrix has columns (in this case it does), and the result matrix will have the same number of rows as the first matrix, and the same number of columns as the second matrix (in this case the result will be a 2X2 matrix).

$$\mathbf{C} = \mathbf{A} \mathbf{B} = \begin{pmatrix} c_{1,1} & c_{1,2} \\ c_{2,1} & c_{2,2} \end{pmatrix} \quad (41)$$

And each element of \mathbf{C} is the dot product of the row vector of \mathbf{A} and the column vector of \mathbf{B} .

²⁷Matrix multiplication is not transitive; $\mathbf{A} \mathbf{B} \neq \mathbf{B} \mathbf{A}$.

$$c_{1,1} = (1 \ 5 \ 7) \cdot \begin{pmatrix} 3 \\ -2 \\ 1 \end{pmatrix} = ((1)(3) + (5)(-2) + (7)(1)) = 0 \quad (42)$$

$$c_{1,2} = (1 \ 5 \ 7) \cdot \begin{pmatrix} -2 \\ 1 \\ 1 \end{pmatrix} = ((1)(-2) + (5)(1) + (7)(1)) = 10 \quad (43)$$

$$c_{2,1} = (2 \ 9 \ 3) \cdot \begin{pmatrix} 3 \\ -2 \\ 1 \end{pmatrix} = ((2)(3) + (9)(-2) + (3)(1)) = -9 \quad (44)$$

$$c_{2,2} = (2 \ 9 \ 3) \cdot \begin{pmatrix} -2 \\ 1 \\ 1 \end{pmatrix} = ((2)(-2) + (9)(1) + (3)(1)) = 8 \quad (45)$$

Making the substitutions results in :

$$\mathbf{C} = \mathbf{AB} = \begin{pmatrix} 0 & 10 \\ -9 & 8 \end{pmatrix} \quad (46)$$

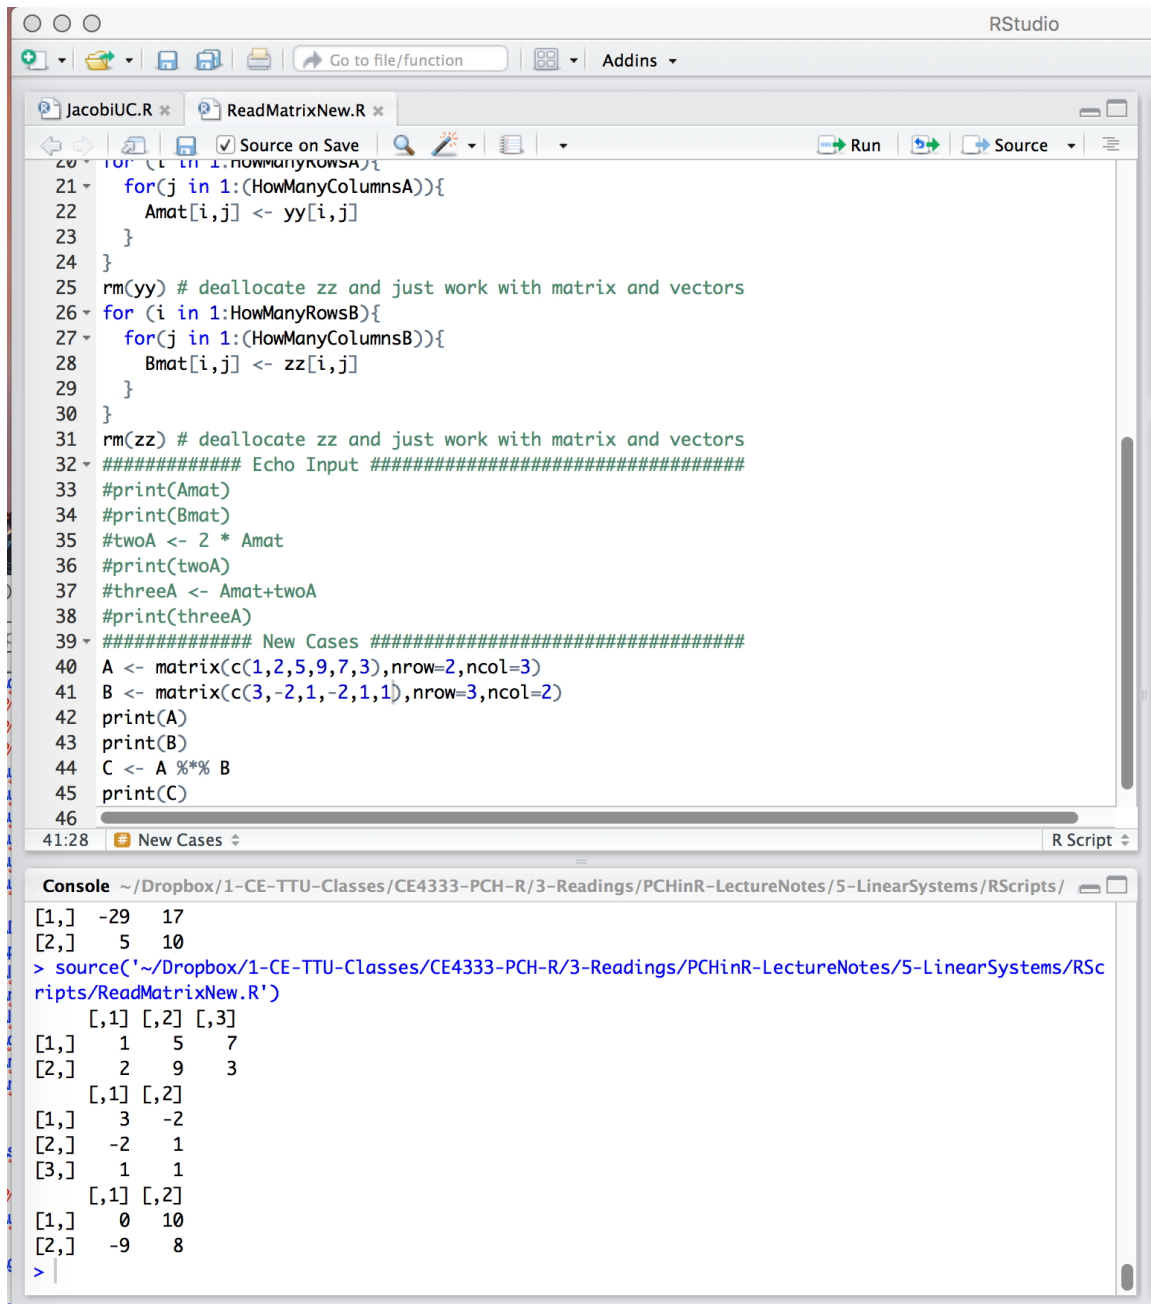
So in an algorithmic sense we will have to deal with three matrices, the two source matrices and the destination matrix. We will also have to manage element-by-element multiplication and be able to correctly store through rows and columns. In **R** this manipulation is handled for us by the matrix multiply operator `% * %`.

Figure 67 is a script that multiplies the two matrices above and prints the result.²⁸

5.3.5 Identity matrix

In computational linear algebra we often need to make use of a special matrix called the “Identity Matrix”. The Identity Matrix is a square matrix with all zeros except the $i, i0$ -th element (diagonal) which is equal to 1:

²⁸Internal to **R** the actual code for the multiplication is three nested for-loops. The outer loop counts based rows of the first matrix, the middle loop counts based on columns of the second matrix, and the inner most loop counts based on columns of the first matrix (or rows of the second matrix). In many practical cases we may actually have to manipulate at the element level — similar to how the **zz** object was put into a matrix explicitly above.



```

RStudio
Go to file/function
Addins
JacobiUC.R * ReadMatrixNew.R *
Source on Save
Run Source
20 for (i in 1:HowManyRowsA){
21   for(j in 1:(HowManyColumnsA)){
22     Amat[i,j] <- yy[i,j]
23   }
24 }
25 rm(yy) # deallocate zz and just work with matrix and vectors
26 for (i in 1:HowManyRowsB){
27   for(j in 1:(HowManyColumnsB)){
28     Bmat[i,j] <- zz[i,j]
29   }
30 }
31 rm(zz) # deallocate zz and just work with matrix and vectors
32 ##### Echo Input #####
33 #print(Amat)
34 #print(Bmat)
35 #twoA <- 2 * Amat
36 #print(twoA)
37 #threeA <- Amat+twoA
38 #print(threeA)
39 ##### New Cases #####
40 A <- matrix(c(1,2,5,9,7,3),nrow=2,ncol=3)
41 B <- matrix(c(3,-2,1,-2,1,1),nrow=3,ncol=2)
42 print(A)
43 print(B)
44 C <- A %*% B
45 print(C)
46
41:28 New Cases R Script
Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/
[1,] -29 17
[2,] 5 10
> source('~Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RSc
ripts/ReadMatrixNew.R')
      [,1] [,2] [,3]
[1,]  1   5   7
[2,]  2   9   3
      [,1] [,2]
[1,]  3  -2
[2,] -2   1
[3,]  1   1
      [,1] [,2]
[1,]  0  10
[2,] -9   8
>

```

Figure 67. Matrix multiplication example.

$$\mathbf{I}_{3 \times 3} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (47)$$

Usually we don't bother with the size subscript i used above and just stipulate that the matrix is sized as appropriate. Multiplying any matrix by (a correctly sized) identity matrix results in no change in the matrix. $\mathbf{IA} = \mathbf{A}$

In **R** the identity matrix is easily created using `<matrix_name> <- diag(dimension)`.

5.3.6 Matrix Inverse

In many practical computational and theoretical operations we employ the concept of the inverse of a matrix. The inverse is somewhat analogous to “dividing” by the matrix. Consider our linear system

$$\mathbf{A} \cdot \mathbf{x} = \mathbf{b} \quad (48)$$

If we wished to solve for \mathbf{x} we would “divide” both sides of the equation by \mathbf{A} . Instead of division (which is essentially left undefined for matrices) we instead multiply by the inverse of the matrix²⁹. The inverse of a matrix \mathbf{A} is denoted by \mathbf{A}^{-1} and by definition is a matrix such that when \mathbf{A}^{-1} and \mathbf{A} are multiplied together, the identity matrix \mathbf{I} results. e.g. $\mathbf{A}^{-1}\mathbf{A} = \mathbf{I}$

Lets consider the matrixes below

$$\mathbf{A} = \begin{pmatrix} 2 & 3 \\ 4 & -3 \end{pmatrix} \quad (49)$$

$$\mathbf{A}^{-1} = \begin{pmatrix} \frac{1}{6} & \frac{1}{6} \\ \frac{2}{9} & -\frac{1}{9} \end{pmatrix} \quad (50)$$

We can check that the matrices are indeed inverses of each other using **R** and matrix multiplication — it should return an identity matrix.

Figure 68 is our multiplication script modified where $\mathbf{A} = \mathbf{A}$ and $\mathbf{B} = \mathbf{A}^{-1}$ perform the multiplication and then report the result. The result is the identity matrix regardless of the order of operation.³⁰

Now that we have some background on what an inverse is, it would be nice to know how to find them — that is a remarkably challenging problem. Here we examine a classical algorithm for finding an inverse if we really need to — computationally we only invert if necessary, there are other ways to “divide” that are faster.

5.3.7 Gauss-Jordan method of finding \mathbf{A}^{-1}

There are a number of methods that can be used to find the inverse of a matrix using elementary row operations. An elementary row operation is any one of the three operations listed below:

²⁹The matrix inverse is the multiplicative inverse of the matrix – we are defining the equivalent of a division operation, just calling it something else. This issue will be huge later on in our workbook, especially when we are dealing with non-linear systems

³⁰Why do you think this is so, when above we stated that multiplication is intransitive?

```

RStudio
Go to file/function
Addins
JacobiUC.R * ReadMatrixNew.R *
Source on Save
Run Source
20 for (i in 1:nrow(A)){
21   for(j in 1:(HowManyColumnsA)){
22     Amat[i,j] <- yy[i,j]
23   }
24 }
25 rm(yy) # deallocate yy and just work with matrix and vectors
26 for (i in 1:HowManyRowsB){
27   for(j in 1:(HowManyColumnsB)){
28     Bmat[i,j] <- zz[i,j]
29   }
30 }
31 rm(zz) # deallocate zz and just work with matrix and vectors
32 ##### Echo Input #####
33 #print(Amat)
34 #print(Bmat)
35 #twoA <- 2 * Amat
36 #print(twoA)
37 #threeA <- Amat+twoA
38 #print(threeA)
39 ##### New Cases #####
40 A <- matrix(c(2,4,3,-3),nrow=2,ncol=2)
41 B <- matrix(c(1/6,2/9,1/6,-1/9),nrow=2,ncol=2)
42 print(A)
43 print(B)
44 C <- A %*% B
45 print(C)
46
41:31 New Cases R Script
Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/
  [,1] [,2]
[1,]  0  10
[2,] -9   8
> source('~/.Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/ReadMatrixNew.R')
  [,1] [,2]
[1,]  2   3
[2,]  4  -3
  [,1] [,2]
[1,] 0.1666667 0.1666667
[2,] 0.2222222 -0.1111111
  [,1] [,2]
[1,]  1   0
[2,]  0   1
>

```

Figure 68. Matrix multiplication used to check an inverse..

1. Multiply or divide an entire row by a constant.
2. Add or subtract a multiple of one row to/from another.
3. Exchange the position of any 2 rows.

The Gauss-Jordan method of inverting a matrix can be divided into 4 main steps. In order to find the inverse we will be working with the original matrix, augmented with the identity matrix – this new matrix is called the augmented matrix (because

no-one has tried to think of a cooler name yet).

$$\mathbf{A}|\mathbf{I} = \left(\begin{array}{cc|cc} 2 & 3 & 1 & 0 \\ 4 & -3 & 0 & 1 \end{array} \right) \quad (51)$$

We will perform elementary row operations based on the left matrix to convert it to an identity matrix – we perform the same operations on the right matrix and the result when we are done is the inverse of the original matrix.

So here goes – in the theory here, we also get to do infinite-precision arithmetic, no rounding/truncation errors.

1. Divide row one by the $a_{1,1}$ value to force a 1 in the $a_{1,1}$ position. This is elementary row operation 1 in our list above.

$$\mathbf{A}|\mathbf{I} = \left(\begin{array}{cc|cc} 2/2 & 3/2 & 1/2 & 0 \\ 4 & -3 & 0 & 1 \end{array} \right) = \left(\begin{array}{cc|cc} 1 & 3/2 & 1/2 & 0 \\ 4 & -3 & 0 & 1 \end{array} \right) \quad (52)$$

2. For all rows below the first row, replace row_j with $row_j - a_{j,1} * row_1$. This happens to be elementary row operation 2 in our list above.

$$\mathbf{A}|\mathbf{I} = \left(\begin{array}{cc|cc} 1 & 3/2 & 1/2 & 0 \\ 4 - 4(1) & -3 - 4(3/2) & 0 - 4(1/2) & 1 - 4(0) \end{array} \right) = \left(\begin{array}{cc|cc} 1 & 3/2 & 1/2 & 0 \\ 0 & -9 & -2 & 1 \end{array} \right) \quad (53)$$

3. Now multiply row_2 by $\frac{1}{a_{2,2}}$. This is again elementary row operation 1 in our list above.

$$\mathbf{A}|\mathbf{I} = \left(\begin{array}{cc|cc} 1 & 3/2 & 1/2 & 0 \\ 0 & -9/-9 & -2/-9 & 1/-9 \end{array} \right) = \left(\begin{array}{cc|cc} 1 & 3/2 & 1/2 & 0 \\ 0 & 1 & 2/9 & -1/9 \end{array} \right) \quad (54)$$

4. For all rows above and below this current row, replace row_j with $row_j - a_{2,2} * row_2$. This happens to again be elementary row operation 2 in our list above. What we are doing is systematically converting the left matrix into an identity matrix by multiplication of constants and addition to eliminate off-diagonal values and force 1 on the diagonal.

$$\mathbf{A}|\mathbf{I} = \quad (55)$$

$$\left(\begin{array}{cc|cc} 1 & 3/2 - (3/2)(1) & 1/2 - (3/2)(2/9) & 0 - (3/2)(-1/9) \\ 0 & 1 & 2/9 & -1/9 \end{array} \right) = \quad (56)$$

$$\left(\begin{array}{cc|cc} 1 & 0 & 1/6 & 1/6 \\ 0 & 1 & 2/9 & -1/9 \end{array} \right) \quad (57)$$

5. As far as this example is concerned we are done and have found the inverse. With more than a 2X2 system there will be many operations moving up and down the matrix to eliminate the off-diagonal terms.

So the next logical step is to build an algorithm to perform these operations for us.

In **R** inversion is simply performed using the `solve(...)` function where the only argument passed to the function is the matrix.³¹

Figure 69 is a screen capture of using `solve(...)` to find the inverse of **A**. The result is identical to the input matrix \mathbf{A}^{-1} above. While we now have the ability to solve linear systems by rearrangement into

$$\mathbf{x} = \mathbf{A}^{-1} \cdot \mathbf{b} \quad (58)$$

this is generally not a good approach (we are solving n linear systems to obtain the inverse, instead of only the one we seek!).

Instead to solve a linear system, we would supply the coefficient matrix **A** and the right hand side **b**, and then supply these two matrices to the solve routine (e.g. `x <- solve(A,b)`).

³¹If we have to write code ourselves, its not terribly hard, but is lengthy and consequently error-prone. Sometimes we have no choice, but in this workbook, we will use the built-in tool as much as possible. **R** does not use Gaussian reduction unless we tell it to do so, it implements a factorization called LU (or Cholesky) decomposition, then computes the inverse by repeated solution of a linear system with the right hand side being selected from one of the identify matrix columns (as was done above).

The screenshot shows the RStudio interface with a script editor and a console. The script editor contains the following R code:

```

20 for (i in 1:HowManyRowsA){
21   for(j in 1:(HowManyColumnsA)){
22     Amat[i,j] <- yy[i,j]
23   }
24 }
25 rm(yy) # deallocate zz and just work with matrix and vectors
26 for (i in 1:HowManyRowsB){
27   for(j in 1:(HowManyColumnsB)){
28     Bmat[i,j] <- zz[i,j]
29   }
30 }
31 rm(zz) # deallocate zz and just work with matrix and vectors
32 ##### Echo Input #####
33 #print(Amat)
34 #print(Bmat)
35 #twoA <- 2 * Amat
36 #print(twoA)
37 #threeA <- Amat+twoA
38 #print(threeA)
39 ##### New Cases #####
40 A <- matrix(c(2,4,3,-3),nrow=2,ncol=2)
41 B <- matrix(c(1/6,2/9,1/6,-1/9),nrow=2,ncol=2)
42 print(A)
43 print(B)
44 C <- A %*% B
45 #print(C)
46 A_inv <- solve(A)
47 print(A_inv)
48

```

The console output shows the execution of the script:

```

> source('~/.Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/
Rscripts/ReadMatrixNew.R')
  [,1] [,2]
[1,]  2   3
[2,]  4  -3
      [,1]      [,2]
[1,] 0.1666667 0.1666667
[2,] 0.2222222 -0.1111111
      [,1]      [,2]
[1,] 0.1666667 0.1666667
[2,] 0.2222222 -0.1111111
>

```

Figure 69. The matrix inversion script showing results of a run and various input and output..

5.4 Jacobi Iteration – An iterative method to find solutions

Iterative methods are often more rapid and economical in storage requirements than the direct methods in `solve(...)`.³² The methods are useful (necessary) for non-linear systems of equations — we will use this feature later when we find solutions to networks of pipelines.

Lets consider a simple example:

$$\begin{aligned} 8x_1 + 1x_2 - 1x_3 &= 8 \\ 1x_1 - 7x_2 + 2x_3 &= -4 \\ 2x_1 + 1x_2 + 9x_3 &= 12 \end{aligned} \quad (59)$$

The solution is $x_1 = 1$, $x_2 = 1$, $x_3 = 1$. We begin the iterative scheme by refactoring each equation in terms of a single variable (there is a secret pivot step to try to make the system diagonally dominant – the example above has already been pivoted, or “pre-conditioned” for the solution method):

$$\begin{aligned} x_1 &= 1.000 && -0.125x_2 & 0.125x_3 \\ x_2 &= 0.571 & 0.143x_1 && 0.286x_3 \\ x_3 &= 1.333 & -0.222x_1 & -0.111x_2 \end{aligned} \quad (60)$$

Then supply an initial guess of the solution (e.g. $(0, 0, 0)$) and put these values into the right-hand side, the resulting left-hand side is an improved (hopefully) solution. Repeat the process until the solution stops changing, or goes obviously haywire.

This sequence of operation for the example above produces the results listed in Table 4.

Table 4. Jacobi Iteration Solution Sequence.

Iteration:	1-st	2-nd	3-rd	4-th	5th	6-th	7-th	8-th
x_1	0	1.000	1.095	0.995	0.993	1.002	1.001	1.000
x_2	0	0.571	1.095	1.026	0.990	0.998	1.001	1.000
x_3	0	1.333	1.048	0.969	1.000	1.004	1.001	1.000

As a practical matter, refactoring the equations can instead be accomplished by computing the inverse of each diagonal coefficient – and matrix multiplication, scalar division, and vector addition are all that is required to find a solution (if the method will actually work).

In linear algebra terms the Jacobi iteration method (without refactoring) performs the following steps:

³²The **R** `solve` routine is pretty robust, if you tell it `sparse=TRUE` it has a lot of internal methods to pre-condition the problem for fast solution. But for really big systems we may wish to program our own solver — especially if these systems have some special and predictable structure.

1. Read in \mathbf{A} , \mathbf{b} , and $\mathbf{x}_{\text{guess}}$.
2. Construct a vector from the diagonal elements of \mathbf{A} . This vector, \mathbf{W} , will have one column, and same number of rows as \mathbf{A} .
3. Perform matrix arithmetic to compute an error vector, $\mathbf{residual} = \mathbf{A} \cdot \mathbf{x}_{\text{guess}} - \mathbf{b}$.
4. Divide this error vector by the diagonal weights $\mathbf{update} = \mathbf{residual}/\mathbf{W}$
5. Update the solution vector $\mathbf{x}_{\text{new}} = \mathbf{x}_{\text{guess}} - \mathbf{update}$
6. Test for stopping, if not indicated, move the new solution into the guess and return to step 3.
7. If time to stop, then report result and stop.

Listing 27 implements in **R** the algorithm described above to find solutions by the Jacobi iteration method. The script does not pre-condition the linear system (so we have to do that ourselves).

Listing 27. R code demonstrating Jacobi Iteration.

```
# R script to implement Jacobi Iteration Method to
# find solution to simultaneous linear equations
# assumes matrix is pre-conditioned to diagonal dominant
# assumes matrix is non-singular
##### READ IN DATA FROM A FILE #####
filepath <- "~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-
LinearSystems/RScripts"
filename <- "LinearSystem000.txt"
fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
# Here we open the connection to the file (within read.table)
# Then the read.table attempts to read the entire file into an object named zz
# Upon either fail or success, read.table closes the connection
zz <- read.table(fileToRead,header=FALSE,sep=",") # comma separated ASCII, No header
##### Row and Column Counts #####
HowManyColumns <- length(zz)
HowManyRows <- length(zz$V1)
tolerance <- 1e-12 #stop when error vector is small
itermax <- 200 # maximum number of iterations
##### Build A, x, and B #####
Amat <- matrix(0,nrow = HowManyRows, ncol = (HowManyColumns-2) )
xguess <- numeric(0)
Bvec <- numeric(0)
Wvec <- numeric(0)
#####
for (i in 1:HowManyRows){
  for(j in 1:(HowManyColumns-2)){
    Amat[i,j] <- zz[i,j]
  }
  Bvec[i] <- zz[i,HowManyColumns-1]
  xguess[i] <- zz[i,HowManyColumns]
  Wvec[i] <- Amat[i,i]
}
rm(zz) # deallocate zz and just work with matrix and vectors
##### Implement Jacobi Iteration #####
for(iter in 1:itermax){
  Bguess <- Amat %*% xguess
  residue <- Bguess - Bvec
  xnew <- xguess - residue/Wvec
  xguess <- xnew
  testval <- t(residue) %*% residue
  if (testval < tolerance) {
    message("sum squared error vector small : ",testval);
    break
  }
}
if( iter == itermax) message("Method Fail")
message(" Number Iterations : ", iter)
message(" Coefficient Matrix : ")
print(cbind(Amat))
message(" Solution Vector : ")
print(cbind(xguess))
message(" Right-Hand Side Vector : ")
print(cbind(Bvec))
```

Figure 70 is a screen capture of the script in Listing 27 applied to the example problem.

```

1 # R script to implement Jacobi Iteration Method to
2 # find solution to simultaneous linear equations
3 # assumes matrix is pre-conditioned to diagonal dominant
4 # assumes matrix is non-singular
5 ##### READ IN DATA FROM A FILE #####
6 filepath <- "~/Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSys
7 filename <- "LinearSystem000.txt"
8 fileToRead <- paste(filepath,filename,sep="/") # build the user absolute filename
9 # Here we open the connection to the file (within read.table)
10 # Then the read.table attempts to read the entire file into an object named zz
11 # Upon either fail or success, read.table closes the connection
12 zz <- read.table(fileToRead,header=FALSE,sep=",") # comma separated ASCII, No header
13 ##### Row and Column Counts #####
14 HowManyColumns <- length(zz)
15 HowManyRows <- length(zz$V1)
16 tolerance <- 1e-12 #stop when error vector is small
17 itermax <- 200 # maximum number of iterations
18 ##### Build A, x, and B #####
19 Amat <- matrix(0,nrow = HowManyRows, ncol = (HowManyColumns-2) )
20 xguess <- numeric(0)
21
17:28 Row and Column Counts R Script

```

```

Console ~/Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/
> source('~/.Dropbox/1-CE-TTU-Classes/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/5-LinearSystems/RScripts/JacobiUC.R')
sum squared error vector small : 1.7061092645239e-13
Number Iterations : 15
Coefficient Matrix :
  [,1] [,2] [,3]
[1,]  8   1  -1
[2,]  1  -7   2
[3,]  2   1   9
Solution Vector :
  [,1]
[1,]  1
[2,]  1
[3,]  1
Right-Hand Side Vector :
  Bvec
[1,]  8
[2,] -4
[3,] 12
>

```

Figure 70. Jacobi Iteration applied to Example Problem.

6 Simultaneous Non-Linear Systems of Equations

Non-linear systems are extensions of the linear systems cases except the systems involve products and powers of the unknown variables. Non-linear problems are often quite difficult to manage, especially when the systems are large (many rows and many variables).

The solution to non-linear systems, if non-trivial yet alone even possible, are iterative.

Within the iterative steps is a linearization component – these linear systems which are intermediate computations within the overall solution process are treated by an appropriate linear system method (direct or iterative).

In **R** it is sometimes successful to solve by the nonlinear minimization tool built-in, but neither efficient, nor particularly useful when the system gets large. On the **CRAN** there are a couple of packages devoted to non-linear systems, and these would be reasonable places to consider.

In this chapter we will illustrate an iterative technique called Quasi-Linearization, and the next chapter we will formally extend Newton's method to multi-dimensional cases.

$$\begin{aligned} x^2 + y^2 &= 4 \\ e^x + y &= 1 \end{aligned} \tag{61}$$

Suppose we have a solution guess x_k, y_k , which of course could be wrong, but we could linearize about that guess as

$$\mathbf{A} = \begin{pmatrix} x_k & + y_k \\ 0 & + 1 \end{pmatrix} \mathbf{x} = \begin{pmatrix} x_{k+1} \\ y_{k+1} \end{pmatrix} \mathbf{b} = \begin{pmatrix} 4 \\ 1 - e^{x_k} \end{pmatrix} \tag{62}$$

Now the system is linear, and we can solve for \mathbf{x}_{k+1} much like the Jacobi iteration of the previous chapter. If the system is convergent (not all are) then we update in the same fashion, and repeat until complete.

Listing 28 is a script that implements the quasi-linearization method. The starting vector is crucial, and the next several screen captures illustrate good starting vectors (resulting in a solution) and poor ones.

Listing 28. R code demonstrating Non-Linear by quasi-linearization.

```

# R script to solve non-linear example by quasi-linearization
Amat <- matrix(0,nrow=2,ncol=2)
Brhs <- numeric(0)
x_guess <- c(-1.9, 0.8)
maxiter <- 20
message("Initial Guess"); print(x_guess); message("Original Equations - x_guess ");
message( x_guess[1]^2 + x_guess[2]^2, " : should be 4 ");
message( exp(x_guess[1]) + x_guess[2], " : should be 1 ")
# Construct the current quasi-linear model
for (iter in 1:maxiter){
  Amat[1,1] <- x_guess[1]; Amat[1,2] <- x_guess[2];
  Amat[2,1] <- 0 ; Amat[2,2] <- 1;
  Brhs[1] <- 4
  Brhs[2] <- 1-exp(x_guess[1])
  # Solve for the new guess
  x_new <- solve(Amat,Brhs)
  # Update
  x_guess <- x_new
}
print(Amat); print(Brhs);
message("Current Guess"); print(x_new)
message("Original Equations - x_new")
message( x_new[1]^2 + x_new[2]^2, " : should be 4 ")
message( exp(x_new[1]) + x_new[2], " : should be 1 ")

```

Figure 71 is a screen capture of the algorithm started near a solution, that sort-of converges to the solution. Not really satisfying, but at least not divergent.

```

RStudio
Updates Not Installed
Some updates could not be
automatically.

QuasiLinear.R
Source on Save
Run Source

1 # R script to solve non-linear example by quasi-linearization
2 Amat <- matrix(0,nrow=2,ncol=2)
3 Brhs <- numeric(0)
4 x_guess <- c(-1.83, 0.8)
5 maxiter <- 20
6 message("Initial Guess"); print(x_guess); message("Original Equations - x_guess ");
7 message( x_guess[1]^2 + x_guess[2]^2, " : should be 4 ")
8 message( exp(x_guess[1]) + x_guess[2], " : should be 1 ")
9 # Construct the current quasi-linear model
10 for (iter in 1:maxiter){
11   Amat[1,1] <- x_guess[1]; Amat[1,2] <- x_guess[2];
12   Amat[2,1] <- 0 ; Amat[2,2] <- 1;
13   Brhs[1] <- 4
14   Brhs[2] <- 1-exp(x_guess[1])
15   # Solve for the new guess
16   x_new <- solve(Amat,Brhs)
17   # Update
18   x_guess <- x_new
19 }
20 print(Amat); print(Brhs);
21 message("Current Guess"); print(x_new)
22 message("Original Equations - x_new")
23 message( x_new[1]^2 + x_new[2]^2, " : should be 4 ")
24 message( exp(x_new[1]) + x_new[2], " : should be 1 ")

20:26 (Top Level) R Script

Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/
> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/QuasiLinear.R")
Initial Guess
[1] -1.83 0.80
Original Equations - x_guess
3.9889 : should be 4
0.96041356775173 : should be 1
      [,1]      [,2]
[1,] -1.820072 0.8367488
[2,] 0.000000 1.0000000
[1] 4.0000000 0.8379858
Current Guess
[1] -1.8124652 0.8379858
Original Equations - x_new
3.98725021975519 : should be 4
1.00123705001415 : should be 1
>

```

Figure 71. Quasi-linear, started near a solution, converges (sort-of) to the solution.

Figure 72 is a screen capture of the algorithm started near a solution, that fails to converge — it actually diverged.

```

1 # R script to solve non-linear example by quasi-linearization
2 Amat <- matrix(0,nrow=2,ncol=2)
3 Brhs <- numeric(0)
4 x_guess <- c(1, -1.7323)
5 maxiter <- 20
6 message("Initial Guess"); print(x_guess); message("Original Equations - x_guess ");
7 message( x_guess[1]^2 + x_guess[2]^2, " : should be 4 ")
8 message( exp(x_guess[1]) + x_guess[2], " : should be 1 ")
9 # Construct the current quasi-linear model
10 for (iter in 1:maxiter){
11 Amat[1,1] <- x_guess[1]; Amat[1,2] <- x_guess[2];
12 Amat[2,1] <- 0 ; Amat[2,2] <- 1;
13 Brhs[1] <- 4
14 Brhs[2] <- 1-exp(x_guess[1])
15 # Solve for the new guess
16 x_new <- solve(Amat,Brhs)
17 # Update
18 x_guess <- x_new
19 }
20 print(Amat); print(Brhs);
21 message("Current Guess"); print(x_new)
22 message("Original Equations - x_new")
23 message( x_new[1]^2 + x_new[2]^2, " : should be 4 ")
24 message( exp(x_new[1]) + x_new[2], " : should be 1 ")

```

```

Console ~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/
> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/QuasiLinear.R")
Initial Guess
[1] 1.0000 -1.7323
Original Equations - x_guess
4.00086329 : should be 4
0.985981828459045 : should be 1
      [,1] [,2]
[1,] -9.189108 0.3293203
[2,] 0.000000 1.0000000
[1] 4.0000000 0.9998979
Current Guess
[1] -0.3994635 0.9998979
Original Equations - x_new
1.1593668337585 : should be 4
1.67057760019411 : should be 1
>

```

Figure 72. Quasi-linear, started near a solution, fails to converge.

What is really needed is a much more reliable algorithm. Sometimes the non-linear minimization tools can successfully be used. We will try that next.

Lets restructure our equation system a bit into

$$\mathbf{f}(\mathbf{x}) = \begin{cases} f_1(x, y) = x^2 + y^2 - 4 \\ f_2(x, y) = e^x + y - 1 \end{cases} \quad (63)$$

At the solution \mathbf{x} , the result should be $\mathbf{f}(\mathbf{x}) = \mathbf{0}$. But if we are not at a solution, then the result will be non-zero (and represents the error) — one tool we have is a non-linear minimization tool in \mathbf{R} that can minimize functions. So now we need the sum-of-squared errors, which with vectors is simply the inner product of the vector with itself:

$$\mathbf{F}(\mathbf{x}) = \mathbf{f}(\mathbf{x})^T \cdot \mathbf{f}(\mathbf{x}) \quad (64)$$

So lets rewrite the script to construct $\mathbf{f}(\mathbf{x})$ and $\mathbf{F}(\mathbf{x})$, then implement the non-linear minimizer, `nlm(...)` in **R**. Listing 29 is a listing that implements these changes. Notice the two prototype functions, the first takes vector input and returns vector output — internal (to the function) definition of a vector using `func <- numeric(0)` provides the memory space in the program.³³

Listing 29. R code demonstrating Non-Linear by Minimization.

```
# R script for system of non-linear equations using minimization
# WARNING -- This is not recommended for large systems
# forward define the functions
##### f(x) #####
func <- function(x_vector){
  func <- numeric(0)
  func[1] <- x_vector[1]^2 + x_vector[2]^2 - 4
  func[2] <- exp(x_vector[1]) + x_vector[2] - 1
  return(func)
}
##### F(x) #####
bigF <- function(x_vector){
  vector <- numeric(0)
  vector <- func(x_vector)
  bigF <- t(vector) \%*\% vector
  return(bigF)
}
#####
# forward define some variables
# starting guess
x_guess <- c(1,-1.7)
result <- nlm(bigF,x_guess)
message(" Estimated bigF Value : ",result$minimum)
message(" Estimated x_vector Value : ")
print(result$estimate)
message(" Estimated func Value : ")
print(func(result$estimate))
```

Figure 73 is a screen capture of the script for the first solution to the system of equations, we have started quite close to a solution and the method converges to the correct solution. The object named `result` contains several items of which we have only accessed two. Notice how we have addressed these items using the `name$attribute` method.

Figure 74 is a screen capture of the script for the second solution to the system of equations, we have started quite close to a solution and the method converges to the correct solution.

Naturally, to be really useful we should test the method for starting values relatively far from the solution; Figure 75 is a screen capture of such testing for a few different start vectors. Observe that the solution at (-1.8,0.8) is the preferred solution in most cases unless we start very close to the second solution at (1,-1.7). This kind of preference to one solution over another is quite common in non-linear systems (sometimes these particular solutions are called attractors). The related observation is that we can find starting vectors that simply fail — this phenomenon is also quite common (sometimes called sensitive dependence on initial conditions).

³³If you get an error message with the words ... `Atomic ...`, it means that something in a function is trying to address a variable for which there is no space, or trying to address a global (external to the function) variable directly. These are pretty hard errors to debug (fix), so I have gotten into the habit of building and testing the prototype functions before I even try to get the rest of the program to run.

```

1 # R script for system of non-linear equations using minimization
2 # WARNING -- This is not recommended for large systems
3 # forward define the functions
4 ##### f(x) #####
5 func <- function(x_vector){
6   func <- numeric(0)
7   func[1] <- x_vector[1]^2 + x_vector[2]^2 - 4
8   func[2] <- exp(x_vector[1]) + x_vector[2] - 1
9   return(func)
10 }
11 ##### F(x) #####
12 bigF <- function(x_vector){
13   vector <- numeric(0)
14   vector <- func(x_vector)
15   bigF <- t(vector) %*% vector
16   return(bigF)
17 }
18 #####
19 # forward define some variables
20 # starting guess
21 x_guess <- c(-1.8162, 0.8374)
22 result <- nlm(bigF, x_guess)
23 message(" Estimated bigF Value : ", result$minimum)
24 message(" Estimated x_vector Value : ")
25 print(result$estimate)
26 message(" Estimated func Value : ")
27 print(func(result$estimate))

```

```

> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/NLM-Method.R")
Estimated bigF Value : 5.6969192604848e-11
Estimated x_vector Value :
[1] -1.8162678  0.8373615
Estimated func Value :
[1] 2.999973e-06 -6.925991e-06
>

```

Figure 73. Solution using `nlm(...)`. Start vector (-1.8,0.8).

```

1 # R script for system of non-linear equations using minimization
2 # WARNING -- This is not recommended for large systems
3 # forward define the functions
4 ##### f(x) #####
5 func <- function(x_vector){
6   func <- numeric(0)
7   func[1] <- x_vector[1]^2 + x_vector[2]^2 - 4
8   func[2] <- exp(x_vector[1]) + x_vector[2] - 1
9   return(func)
10 }
11 ##### F(x) #####
12 bigF <- function(x_vector){
13   vector <- numeric(0)
14   vector <- func(x_vector)
15   bigF <- t(vector) %*% vector
16   return(bigF)
17 }
18 #####
19 # forward define some variables
20 # starting guess
21 x_guess <- c(1, -1.7)
22 result <- nlm(bigF, x_guess)
23 message(" Estimated bigF Value : ", result$minimum)
24 message(" Estimated x_vector Value : ")
25 print(result$estimate)
26 message(" Estimated func Value : ")
27 print(func(result$estimate))

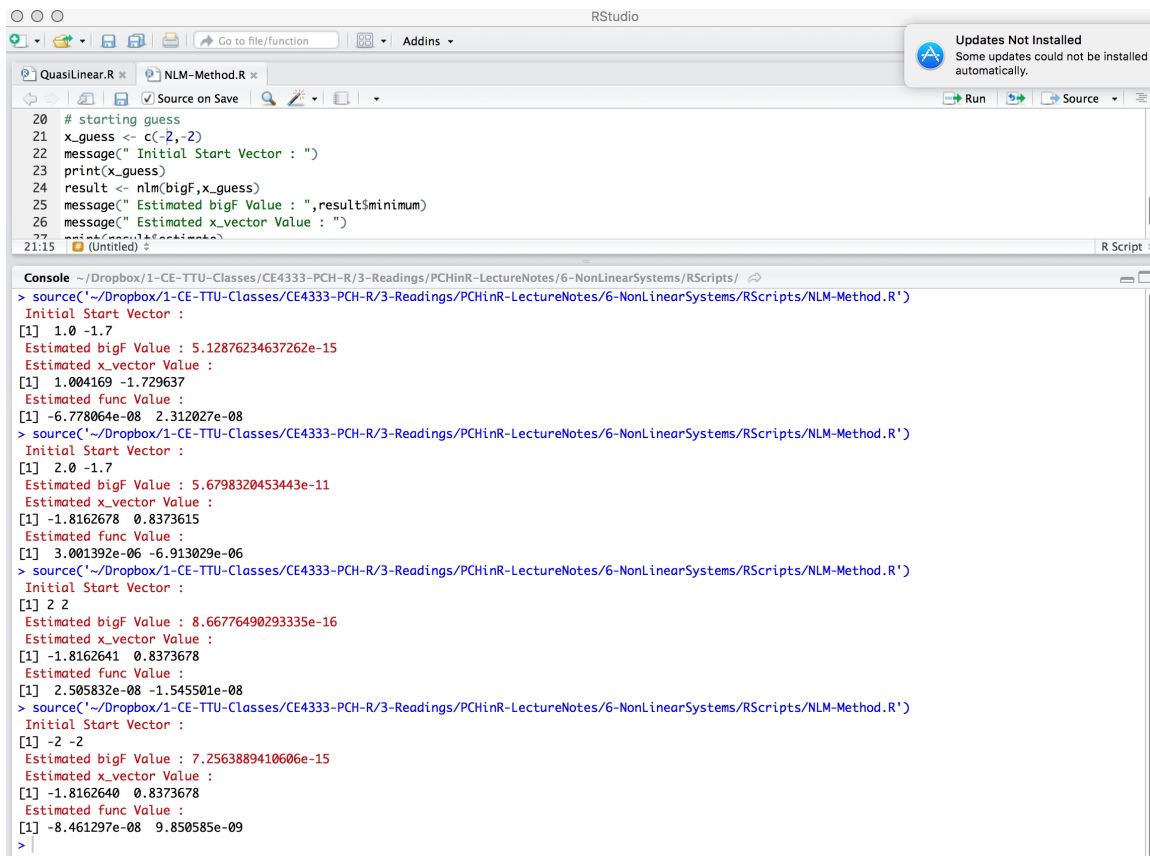
```

```

> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/NLM-Method.R")
Estimated bigF Value : 5.12876234637262e-15
Estimated x_vector Value :
[1] 1.004169 -1.729637
Estimated func Value :
[1] -6.778064e-08  2.312027e-08
>

```

Figure 74. Solution using `nlm(...)`. Start vector (1.0,-1.7).



```

20 # starting guess
21 x_guess <- c(-2,-2)
22 message(" Initial Start Vector : ")
23 print(x_guess)
24 result <- nlm(bigF,x_guess)
25 message(" Estimated bigF Value : ",result$minimum)
26 message(" Estimated x_vector Value : ")
27 print(result$x)

```

```

> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/NLM-Method.R")
Initial Start Vector :
[1] 1.0 -1.7
Estimated bigF Value : 5.12876234637262e-15
Estimated x_vector Value :
[1] 1.004169 -1.729637
Estimated func Value :
[1] -6.778064e-08 2.312027e-08
> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/NLM-Method.R")
Initial Start Vector :
[1] 2.0 -1.7
Estimated bigF Value : 5.6798320453443e-11
Estimated x_vector Value :
[1] -1.8162678 0.8373615
Estimated func Value :
[1] 3.001392e-06 -6.913029e-06
> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/NLM-Method.R")
Initial Start Vector :
[1] 2 2
Estimated bigF Value : 8.66776490293335e-16
Estimated x_vector Value :
[1] -1.8162641 0.8373678
Estimated func Value :
[1] 2.505832e-08 -1.545501e-08
> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHinR-LectureNotes/6-NonLinearSystems/RScripts/NLM-Method.R")
Initial Start Vector :
[1] -2 -2
Estimated bigF Value : 7.2563889410606e-15
Estimated x_vector Value :
[1] -1.8162640 0.8373678
Estimated func Value :
[1] -8.461297e-08 9.850585e-09
>

```

Figure 75. Solution using `nlm(...)`. Varying start vectors.

Using a non-linear minimization technique to solve systems of non-linear equations is not recommended for anything bigger than a few equations (maybe as many as 8 or 9). Quasi-linearization is a good technique — the example here is intentionally pathological. The next chapter presents a better technique than quasi-linearization that can be used for large systems (assuming they will converge at all), and it is the method that will be used for pipeline networks.

7 Numerical Methods – Multiple Variable Quasi-Newton Method

This chapter formally presents the Newton-Raphson method as the preferred alternative to using an optimizer routine to solve systems of non-linear equations. The method is used later in the document to solve for flows and heads in a pipeline network.

Lets return to our previous example where the function \mathbf{f} is a vector-valued function of a vector argument.

$$\mathbf{f}(\mathbf{x}) = \begin{matrix} f_1 = & x^2 & + & y^2 & -4 \\ f_2 = & e^x & + & y & -1 \end{matrix} \quad (65)$$

Lets also recall Newtons method for scalar valued function of a single variable.

$$x_{k+1} = x_k - \frac{f(x_k)}{\frac{df}{dx}|_{x_k}} \quad (66)$$

Extending to higher dimensions, the value x become the vector \mathbf{x} and the function $f()$ becomes the vector function $\mathbf{f}()$. What remains is an analog for the first derivative in the denominator (and the concept of division of a matrix).

The analog to the first derivative is a matrix called the Jacobian which is comprised of the first derivatives of the function \mathbf{f} with respect to the arguments \mathbf{x} . For example for a 2-value function of 2 arguments (as our example above)

$$\frac{df}{dx}|_{x_k} \Rightarrow \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} \end{pmatrix} \quad (67)$$

Next recall that division is replaced by matrix multiplication with the multiplicative inverse, so the analogy continues as

$$\frac{1}{\frac{df}{dx}|_{x_k}} \Rightarrow \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} \end{pmatrix}^{-1} \quad (68)$$

Lets name the Jacobian $\mathbf{J}(\mathbf{x})$.

So the multi-variate Newton's method can be written as

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{J}(\mathbf{x})^{-1}|_{x_k} \cdot \mathbf{f}(\mathbf{x})|_{x_k} \quad (69)$$

In the linear systems chapter we did find a way to solve for an inverse, but its not necessary – a series of rearrangement of the system above yields a nice scheme tthat does not require inversion of a matrix.

First, move the \mathbf{x}_k to the left-hand side.

$$\mathbf{x}_{k+1} - \mathbf{x}_k = -\mathbf{J}(\mathbf{x})^{-1}|_{x_k} \cdot \mathbf{f}(\mathbf{x})|_{x_k} \quad (70)$$

Next multiply both sides by the Jacobian.

$$\mathbf{J}(\mathbf{x})|_{x_k} \cdot (\mathbf{x}_{k+1} - \mathbf{x}_k) = -\mathbf{J}(\mathbf{x})|_{x_k} \cdot \mathbf{J}(\mathbf{x})^{-1}|_{x_k} \cdot \mathbf{f}(\mathbf{x})|_{x_k} \quad (71)$$

Recall a matrix multiplied by its inverse returns the identity matrix (the matrix equivalent of unity)

$$-\mathbf{J}(\mathbf{x})|_{x_k} \cdot (\mathbf{x}_{k+1} - \mathbf{x}_k) = \mathbf{f}(\mathbf{x})|_{x_k} \quad (72)$$

So we now have an algorithm:

1. Start with an initial guess \mathbf{x}_k , compute $\mathbf{f}(\mathbf{x})|_{x_k}$, and $\mathbf{J}(\mathbf{x})|_{x_k}$.
2. Test for stopping. Is $\mathbf{f}(\mathbf{x})|_{x_k}$ close to zero? If yes, exit and report results, otherwise continue.
3. Solve the linear system $\mathbf{J}(\mathbf{x})|_{x_k} \cdot (\mathbf{x}_{k+1} - \mathbf{x}_k) = \mathbf{f}(\mathbf{x})|_{x_k}$.
4. Test for stopping. Is $(\mathbf{x}_{k+1} - \mathbf{x}_k)$ close to zero? If yes, exit and report results, otherwise continue.
5. Compute the update $\mathbf{x}_{k+1} = \mathbf{x}_k - (\mathbf{x}_{k+1} - \mathbf{x}_k)$, then
6. Move the update into the guess vector $\mathbf{x}_k \leftarrow \mathbf{x}_{k+1}$ and repeat step 1. Stop after too many steps.

Now to repeat the example from the previous chapter, except we will employ this algorithm.

The function (repeated)

$$\mathbf{f}(\mathbf{x}) = \begin{matrix} f_1 = x^2 + y^2 - 4 \\ f_2 = e^x + y - 1 \end{matrix} \quad (73)$$

Then the Jacobian, here we will compute it analytically because we can

$$\mathbf{J}(\mathbf{x}) \Rightarrow \begin{pmatrix} 2x & 2y \\ e^x & 1 \end{pmatrix} \quad (74)$$

Listing 30 is a listing that implements the Newton-Raphson method with analytical derivatives.

Listing 30. R code demonstrating Newton's Method calculations.

```
# R script for system of non-linear equations using Newton-Raphson with analytical
# derivatives
# forward define the functions
##### f(x) #####
func <- function(x_vector){
  func <- numeric(0)
  func[1] <- x_vector[1]^2 + x_vector[2]^2 - 4
  func[2] <- exp(x_vector[1]) + x_vector[2] - 1
  return(func)
}
##### J(x) #####
jacob <- function(x_vector){
  jacob <- matrix(0,nrow=2,ncol=2)
  jacob[1,1] <- 2*x_vector[1] ; jacob[1,2] <- 2*x_vector[2];
  jacob[2,1] <- exp(x_vector[1]); jacob[2,2] <- 1 ;
  return(jacob)
}
##### Solver Parameters #####
x_guess <- c(2.,-0.8)
tolerancef <- 1e-9 # stop if function gets to zero
tolerancex <- 1e-9 # stop if solution not changing
maxiter <- 20 # stop if too many iterations
x_now <- x_guess
##### Newton-Raphson Algorithm #####
for (iter in 1:maxiter){
  funcNow <- func(x_now)
  testf <- t(funcNow) %*% funcNow
  if(testf < tolerancef){
    message("f(x) is close to zero : ", testf);
    break
  }
  dx <- solve(jacob(x_now),funcNow)
  testx <- t(dx) %*% dx
  if(testx < tolerancex){
    message("solution change small : ", testx);
    break
  }
  x_now <- x_now - dx
}
#####
if( iter == maxiter) {message("Maximum iterations -- check if solution is converging : ")}
message("Initial Guess"); print(x_guess);
message("Initial Function Value: "); print(func(x_guess));
message("Exit Function Value : ");print(func(x_now));
message("Exit Vector : "); print(x_now)
```

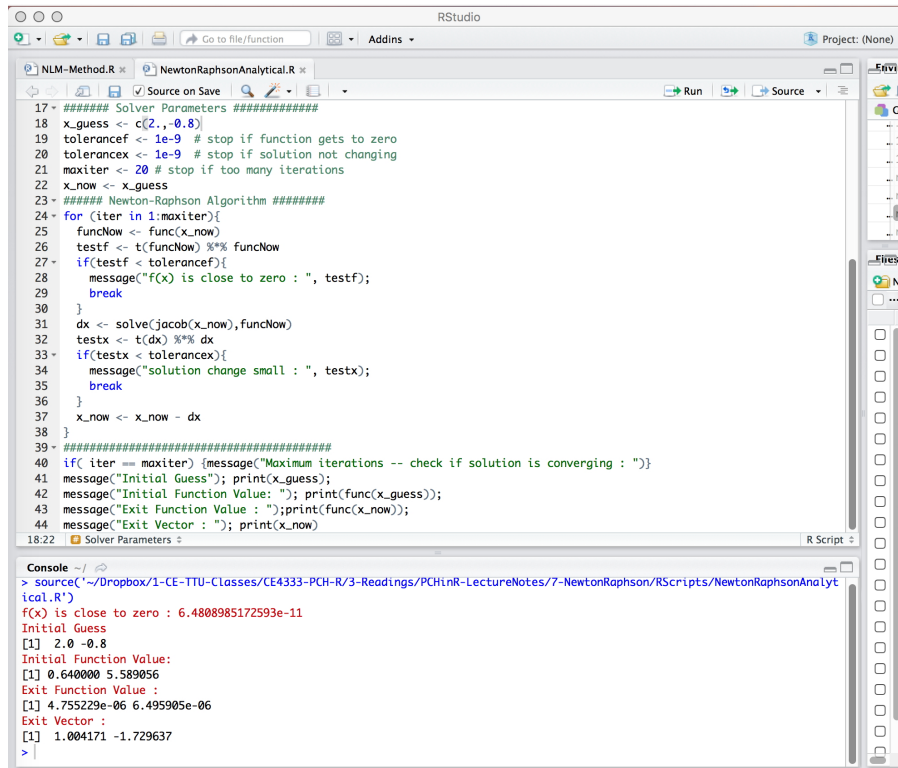
Figure 76 implements the script in Listing 30 for the example problem.

The next variant is to approximate the derivatives – usually a Finite-Difference approximation is used, either forward, backward, or centered differences – generally determined based on the actual behavior of the functions themselves or by trial and error. For really huge systems, we usually make the program itself make the adaptations as it proceeds.

The coding for a finite-difference representation of a Jacobian is shown in Listing 31. In constructing the Jacobian, we observe that each column of the Jacobian is simply the directional derivative of the function with respect to the variable associated with the column. For instance, the first column of the Jacobian in the example is first derivative of the first function (all rows) with respect to the first variable, in this case x . The second column is the first derivative of the second function with respect to the second variable, y . This structure is useful to generalize the Jacobian construction method because we can write (yet another) prototype function that can take the directional derivatives for us, and just insert the returns as columns. The example listing is specific to the 2X2 function in the example, but the extension to more general cases is evident.

Listing 31. R code demonstrating Newton's Method calculations using finite-difference approximations to the partial derivatives.

```
# R script for system of non-linear equations using Newton-Raphson with
# finite-difference approximated derivatives
# forward define the functions
##### f(x) #####
func <- function(x_vector){
  func <- numeric(0)
  func[1] <- x_vector[1]^2 + x_vector[2]^2 - 4
  func[2] <- exp(x_vector[1]) + x_vector[2] - 1
  return(func)
}
##### J(x) #####
jacob <- function(x_vector,func){ #supply a vector and the function name
# the columns of the jacobian are just directional derivatives
dv <- 1e-06 #perturbation value for finite difference
df1 <- numeric(0);
df2 <- numeric(0);
dxv <- x_vector;
dyv <- x_vector;
# perturb the vectors
dxv[1] <- dxv[1]+dv;
dyv[2] <- dyv[2]+dv;
df1 <- (func(dxv) - func(x_vector))/dv;
df2 <- (func(dyv) - func(x_vector))/dv;
jacob <- matrix(0,nrow=2,ncol=2)
# for a more general case should put this into a loop
jacob[1,1] <- df1[1] ; jacob[1,2] <- df2[1] ;
jacob[2,1] <- df1[2] ; jacob[2,2] <- df2[2] ;
return(jacob)
}
##### Solver Parameters #####
x_guess <- c(2.,-0.8)
tolerancef <- 1e-9 # stop if function gets to zero
tolerancex <- 1e-9 # stop if solution not changing
maxiter <- 20 # stop if too many iterations
x_now <- x_guess
##### Newton-Raphson Algorithm #####
for (iter in 1:maxiter){
  funcNow <- func(x_now)
  testf <- t(funcNow) %*% funcNow
  if(testf < tolerancef){
    message("f(x) is close to zero : ", testf);
    break
  }
  dx <- solve(jacob(x_now,func),funcNow)
  testx <- t(dx) %*% dx
  if(testx < tolerancex){
    message("solution change small : ", testx);
    break
  }
  x_now <- x_now - dx
}
#####
if( iter == maxiter) {message("Maximum iterations -- check if solution is converging : ")}
message("Initial Guess"); print(x_guess);
message("Initial Function Value: "); print(func(x_guess));
message("Exit Function Value : ");print(func(x_now));
message("Exit Vector : "); print(x_now)
```



```

17- ##### Solver Parameters #####
18 x_guess <- c(2., -0.8)
19 tolerancef <- 1e-9 # stop if function gets to zero
20 tolerancecx <- 1e-9 # stop if solution not changing
21 maxiter <- 20 # stop if too many iterations
22 x_now <- x_guess
23 ##### Newton-Raphson Algorithm #####
24 for (iter in 1:maxiter){
25   funcNow <- func(x_now)
26   testf <- t(funcNow) %>% funcNow
27   if(testf < tolerancef){
28     message("f(x) is close to zero : ", testf);
29     break
30   }
31   dx <- solve(jacob(x_now), funcNow)
32   testx <- t(dx) %>% dx
33   if(testx < tolerancecx){
34     message("solution change small : ", testx);
35     break
36   }
37   x_now <- x_now - dx
38 }
39 #####
40 if( iter == maxiter) {message("Maximum iterations -- check if solution is converging : ")}
41 message("Initial Guess"); print(x_guess);
42 message("Initial Function Value: "); print(func(x_guess));
43 message("Exit Function Value : "); print(func(x_now));
44 message("Exit Vector : "); print(x_now)
18:22 Solver Parameters

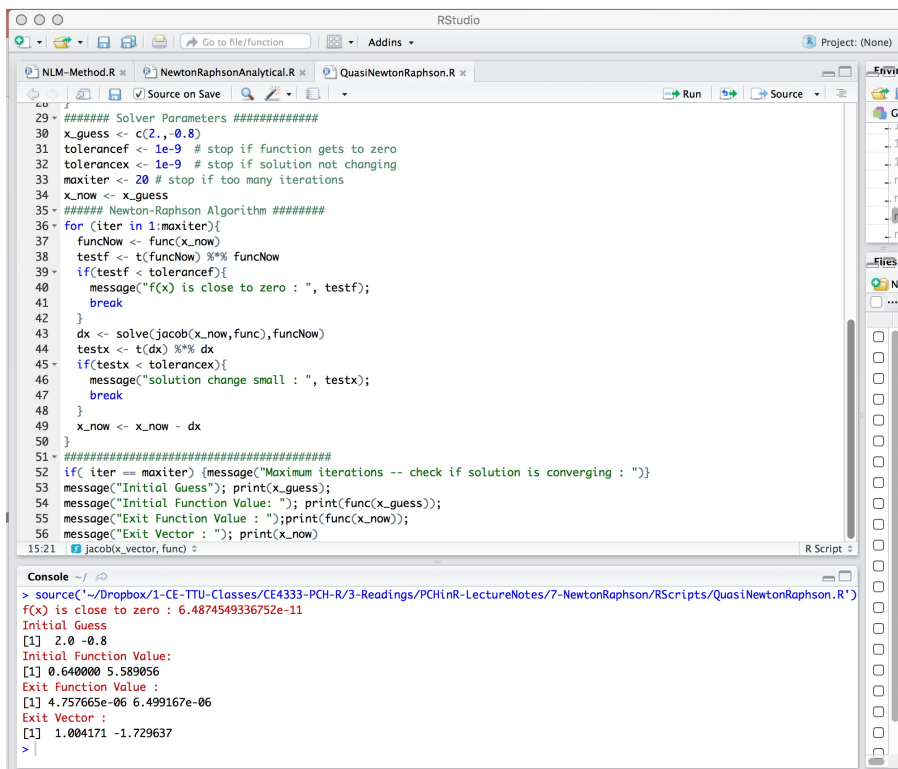
```

```

Console ~/
> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHiR-LectureNotes/7-NewtonRaphson/RScripts/NewtonRaphsonAnalytical.R")
f(x) is close to zero : 6.4808985172593e-11
Initial Guess
[1] 2.0 -0.8
Initial Function Value:
[1] 0.640000 5.589056
Exit Function Value :
[1] 4.755229e-06 6.495905e-06
Exit Vector :
[1] 1.004171 -1.729637
>

```

Figure 76. Newton-Raphson using Analytical Derivatives.



```

29- ##### Solver Parameters #####
30 x_guess <- c(2., -0.8)
31 tolerancef <- 1e-9 # stop if function gets to zero
32 tolerancecx <- 1e-9 # stop if solution not changing
33 maxiter <- 20 # stop if too many iterations
34 x_now <- x_guess
35 ##### Newton-Raphson Algorithm #####
36 for (iter in 1:maxiter){
37   funcNow <- func(x_now)
38   testf <- t(funcNow) %>% funcNow
39   if(testf < tolerancef){
40     message("f(x) is close to zero : ", testf);
41     break
42   }
43   dx <- solve(jacob(x_now, func), funcNow)
44   testx <- t(dx) %>% dx
45   if(testx < tolerancecx){
46     message("solution change small : ", testx);
47     break
48   }
49   x_now <- x_now - dx
50 }
51 #####
52 if( iter == maxiter) {message("Maximum iterations -- check if solution is converging : ")}
53 message("Initial Guess"); print(x_guess);
54 message("Initial Function Value: "); print(func(x_guess));
55 message("Exit Function Value : "); print(func(x_now));
56 message("Exit Vector : "); print(x_now)
15:21 jacob(x_vector, func)

```

```

Console ~/
> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/3-Readings/PCHiR-LectureNotes/7-NewtonRaphson/RScripts/QuasiNewtonRaphson.R")
f(x) is close to zero : 6.4874549336752e-11
Initial Guess
[1] 2.0 -0.8
Initial Function Value:
[1] 0.640000 5.589056
Exit Function Value :
[1] 4.757665e-06 6.499167e-06
Exit Vector :
[1] 1.004171 -1.729637
>

```

Figure 77. Newton-Raphson using Finite-Difference Approximated Derivatives.

8 Pipelines and Networks

Pipe networks, like single path pipelines, are analyzed for head losses in order to size pumps, determine demand management strategies, and ensure minimum pressures in the system. Conceptually the same principles are used for steady flow systems: conservation of mass and energy; with momentum used to determine head losses.

8.1 Pipe Networks – Topology

Network topology refers to the layout and connections. Networks are built of nodes (junctions) and arcs (links).

8.1.1 Continuity (at a node)

Water is considered incompressible in steady flow in pipelines and pipe networks, and the conservation of mass reduces to the volumetric flow rate, Q ,

$$Q = AV \quad (75)$$

where A is the cross sectional of the pipe, and V is the mean section velocity. Typical units for discharge is liters per second (lps), gallons per minute (gpm), cubic meters per second (cms), cubic feet per second (cfs), and million gallons per day (mgd). The continuity equation in two cross-sections of a pipe as depicted in Figure 78 is

$$A_1V_1 = A_2V_2 \quad (76)$$

Junctions (nodes) are where two or more pipes join together. A three-pipe junction node with constant external demand is shown in Figure 10. The continuity equation for the junction node is

$$Q_1 - Q_2 - Q_3 - D = 0 \quad (77)$$




Figure 78. Continuity of mass (discharge) across a change in cross section.

In design analysis, all demands on the system are located at junctions (nodes), and the flow connecting junctions is assumed to be uniform across the cross sections (so




Figure 79. Continuity of mass (discharge) across a node (junction).

that mean velocities apply). If a substantial demand is located between nodes, then an additional node is established at the demand location.

8.1.2 Energy Loss (along a link)

Equation 101 is the one-dimensional steady flow form of the energy equation typically applied for pressurized conduit hydraulics.

$$\frac{p_1}{\rho g} + \alpha_1 \frac{V_1^2}{2g} + z_1 + h_p = \frac{p_2}{\rho g} + \alpha_2 \frac{V_2^2}{2g} + z_2 + h_t + h_l \quad (78)$$

where $\frac{p}{\rho g}$ is the pressure head at a location, $\alpha \frac{V^2}{2g}$ is the velocity head at a location, z is the elevation, h_p is the added head from a pump, h_t is the added head extracted by a turbine, and h_l is the head loss between sections 1 and 2. Figure 87 is a sketch that illustrates the various components in Equation 101.

In network analysis this energy equation is applied to a link that joins two nodes. Pumps and turbines would be treated as separate components (links) and their hy-




Figure 5-1 Definition sketch for terms in the energy equation

Figure 80. Definition sketch for energy equation.

draulic behavior must be supplied using their respective pump/turbine curves.

8.1.3 Velocity Head

The velocity in $\alpha \frac{V^2}{2g}$ is the mean section velocity and is the ratio of discharge to flow area. The kinetic energy correction coefficient is

$$\alpha = \frac{\int_A u^3 dA}{V^3 A} \quad (79)$$

where u is the point velocity in the cross section (usually measured relative to the centerline or the pipe wall; axial symmetry is assumed). Generally values of α are 2.0 if the flow is laminar, and approach unity (1.0) for turbulent flow. In most water distribution systems the flow is usually turbulent so α is assumed to be unity and the velocity head is simply $\frac{V^2}{2g}$.

8.1.4 Added Head — Pumps

The head supplied by a pump is related to the mechanical power supplied to the flow. Equation 102 is the relationship of mechanical power to added pump head.

$$\eta P = Q \rho g h_p \quad (80)$$

where the power supplied to the motor is P and the “wire-to-water” efficiency is η .

If the relationship is re-written in terms of added head³⁴ the pump curve is

$$h_p = \frac{\eta P}{Q \rho g} \quad (81)$$

This relationship illustrates that as discharge increases (for a fixed power) the added head decreases. Power scales at about the cube of discharge, so pump curves for computational application typically have a mathematical structure like

$$h_p = H_{\text{shutoff}} - K_{\text{pump}} Q^{\text{exponent}} \quad (82)$$

8.1.5 Extracted Head — Turbines

The head recovered by a turbine is also an “added head” but appears on the loss side of the equation. Equation 109 is the power that can be recovered by a turbine (again using the concept of “water-to-wire” efficiency is

$$P = \eta Q \rho g h_t \quad (83)$$

8.2 Pipe Head Loss Models

The Darcy-Weisbach, Chezy, Manning, and Hazen-Williams formulas are relationships between physical pipe characteristics, flow parameters, and head loss. The Darcy-Weisbach formula is the most consistent with the energy equation formulation being derivable (in structural form) from elementary principles.

$$h_{L_f} = f \frac{L V^2}{D 2g} \quad (84)$$

where h_{L_f} is the head loss from pipe friction, f is a dimensionless friction factor, L is the pipe length, D is the pipe characteristic diameter, V is the mean section velocity, and g is the gravitational acceleration.

The friction factor, f , is a function of Reynolds number Re_D and the roughness ratio $\frac{k_s}{D}$.

$$f = \sigma(Re_D, \frac{k_s}{D}) \quad (85)$$

The structure of σ is determined experimentally. Over the last century the structure is generally accepted to be one of the following depending on flow conditions and pipe properties

1. Laminar flow (Eqn 2.36, pg. 17 Chin (2006)) :

$$f = \frac{64}{Re_D} \quad (86)$$

³⁴A negative head loss!

2. Hydraulically Smooth Pipes(Eqn 2.34 pg. 16 Chin (2006)):

$$\frac{1}{\sqrt{f}} = -2\log_{10}\left(\frac{2.51}{Re_d\sqrt{f}}\right) \quad (87)$$

3. Hydraulically Rough Pipes(Eqn 2.34 pg. 16 Chin (2006)):

$$\frac{1}{\sqrt{f}} = -2\log_{10}\left(\frac{k_e}{3.7D}\right) \quad (88)$$

4. Transitional Pipes (Colebrook-White Formula)(Eqn 2.35 pg. 17 Chin (2006)):

$$\frac{1}{\sqrt{f}} = -2\log_{10}\left(\frac{k_e}{3.7D} + \frac{2.51}{Re_d\sqrt{f}}\right) \quad (89)$$

5. Transitional Pipes (Jain Formula)(Eqn 2.39 pg. 19 Chin (2006)):

$$f = \frac{0.25}{\left[\log_{10}\left(\frac{k_e}{3.7D} + \frac{5.74}{Re_d^{0.9}}\right)\right]^2} \quad (90)$$

8.3 Pipe Networks Solution Methods

Several methods are used to produce solutions (estimates of discharge, head loss, and pressure) in a network. An early one, that only involves analysis of loops is the Hardy-Cross method. A later one, more efficient, is a Newton-Raphson method that uses node equations to balance discharges and demands, and loop equations to balance head losses. However, a rather ingenious method exists developed by Haman and Brameller (1971), where the flow distribution and head values are determined simultaneously. The task here is to outline the Haman and Brameller (1971) method on the problem below – first some necessary definitions and analysis.

The fundamental procedure is:

1. Continuity is written at nodes (node equations).
2. Energy loss (gain) is written along links (pipe equations).
3. The entire set of equations is solved simultaneously.

8.4 Network Analysis

Figure 81 is a sketch of the problem that will be used. The network supply is the fixed-grade node in the upper left hand corner of the drawing. The remaining nodes (N1 – N4) have demands specified as the purple outflow arrows. The pipes are labeled




Figure 81. Pipe network for illustrative example with supply and demands identified. Pipe dimensions and diameters are also depicted..

(P1 – P6), and the red arrows indicate a positive flow direction, that is, if the flow is in the indicated direction, the numerical value of flow (or velocity) in that link would be a positive number.

Define the flows in each pipe and the total head at each node as Q_i and H_i where the subscript indicates the particular component identification. Expressed as a vector, these unknowns are:

$$[Q_1, Q_2, Q_3, Q_4, Q_5, Q_6, H_1, H_2, H_3, H_4] = \mathbf{x}$$

If we analyze continuity for each node we will have 4 equations (corresponding to each node) for continuity, for instance for Node N2 the equation is

$$Q_2 - Q_3 - Q_6 = 4$$

Similarly if we define head loss in any pipe as $\Delta H_i = f \frac{8L_i}{\pi^2 g D_i^5} |Q_i| Q_i$ or $\Delta H_i = L_i Q_i$, where $L_i = f \frac{8L_i}{\pi^2 g D_i^5} |Q_i|$, then we have 6 equations (corresponding to each pipe) for energy, for instance for Pipe (P2) the equation is³⁵

$$-L_2 Q_2 - H_1 + H_2 = 0$$

³⁵The seemingly awkward way of writing the equations will become apparent shortly!

If we now write all the node equations then all the pipe equations we could construct the following coefficient matrix below.³⁶

$$\begin{array}{cccccccccc}
 \hline
 1 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 1 & -1 & -1 & 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
 \hline
 -L_1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\
 0 & -L_2 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\
 0 & 0 & -L_3 & 0 & 0 & 0 & 0 & 1 & 0 & -1 \\
 0 & 0 & 0 & -L_4 & 0 & 0 & 1 & 0 & -1 & 0 \\
 0 & 0 & 0 & 0 & -L_5 & 0 & 0 & 0 & 1 & -1 \\
 0 & 0 & 0 & 0 & 0 & -L_6 & 0 & -1 & 1 & 0 \\
 \hline
 \end{array}$$

Declare the name of this matrix $\mathbf{A}(\mathbf{x})$, where \mathbf{x} denotes the unknown vector of Q augmented by H as above. Next consider the right-hand-side at the correct solution (as of yet still unknown!) as

$$[0, 4, 3, 1, -100, 0, 0, 0, 0, 0] = \mathbf{b}$$

So if the coefficient matrix is correct then the following system would result:

$$\mathbf{A}(\mathbf{x}) \cdot \mathbf{x} = \mathbf{b}$$

which would look like

$$\begin{pmatrix}
 1 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 1 & -1 & -1 & 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
 \hline
 -L_1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\
 0 & -L_2 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\
 0 & 0 & -L_3 & 0 & 0 & 0 & 0 & 1 & 0 & -1 \\
 0 & 0 & 0 & -L_4 & 0 & 0 & 1 & 0 & -1 & 0 \\
 0 & 0 & 0 & 0 & -L_5 & 0 & 0 & 0 & 1 & -1 \\
 0 & 0 & 0 & 0 & 0 & -L_6 & 0 & -1 & 1 & 0 \\
 \hline
 \end{pmatrix}
 \begin{pmatrix}
 Q_1 \\
 Q_2 \\
 Q_3 \\
 Q_4 \\
 Q_5 \\
 Q_6 \\
 H_1 \\
 H_2 \\
 H_3 \\
 H_4
 \end{pmatrix}
 =
 \begin{pmatrix}
 0 \\
 4 \\
 3 \\
 1 \\
 -100 \\
 0 \\
 0 \\
 0 \\
 0 \\
 0
 \end{pmatrix}
 \quad (91)$$

Observe, the system is non-linear because the coefficient matrix depends on the current values of Q_i for the L_i terms. However, the system is full-rank (rows == columns) so it is a candidate for Newton-Raphson.

³⁶The horizontal lines divide the node and the pipe equations. The upper partition are the node equations in Q and H, the lower partition are the pipe equations in Q and H

Further observe that the upper partition from column 6 and smaller is simply the node-arc incidence matrix, and the lower partition for the same columns only contains L_i terms on its diagonal, the remainder is zero. Next observe that the partition associated with heads in the node equations is the zero-matrix.

Lastly (and this is important!) the lower right partition is the transpose of the node-arc incidence matrix subjected to scalar multiplication of -1 . The importance is that all the information needed to find a solution is contained in the node-arc incidence matrix and the right-hand-side – the engineer does not need to identify closed loops (nor does the computer need to find closed loops).

The trade-off is a much larger system of equations, however solving large systems is far easier than searching a directed graph to identify closed loops, furthermore we obtain the heads as part of the solution process.

9 Pipelines Network Analysis

The prior chapter introduced the non-linear system that results from the analysis of the pipeline network. This chapter continues the effort and produces a workable **R** script that can compute flows and heads given just the node-arc incidence matrix, and pipe properties.

Recall from the prior chapter the non-linear system to be solved is

$$\mathbf{A}(\mathbf{x}) \cdot \mathbf{x} = \mathbf{b}$$

which would look like

$$\begin{pmatrix} 1 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \hline -L_1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & -L_2 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & -L_3 & 0 & 0 & 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & -L_4 & 0 & 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & -L_5 & 0 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 & 0 & -L_6 & 0 & -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} Q_1 \\ Q_2 \\ Q_3 \\ Q_4 \\ Q_5 \\ Q_6 \\ H_1 \\ H_2 \\ H_3 \\ H_4 \end{pmatrix} = \begin{pmatrix} 0 \\ 4 \\ 3 \\ 1 \\ \hline -100 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad (92)$$

The system is non-linear because the coefficient matrix depends on the current values of Q_i for the L_i terms. The upper partition from column 6 and smaller is simply the node-arc incidence matrix, and the lower partition for the same columns only contains L_i terms on its diagonal, the remainder is zero. Next observe that the partition associated with heads in the node equations is the zero-matrix. The lower right partition is the transpose of the node-arc incidence matrix subjected to scalar multiplication of -1 . So using the Newton-Raphson approach discussed earlier we develop a script in **R** that produces estimates of discharge and total head in the system depicted in Figure 81.

9.1 Script Structure

The script will need to accomplish several tasks including reading the node-arc incidence matrix supplied as the file in Figure 82 and convert the strings into numeric values. The script will also need some support functions defined before constructing the matrix.

The rows of the input file are:

```

4
6
1.00 0.67 0.67 0.67 0.67 0.5
800 800 700 700 800 600
0.00001 0.00001 0.00001 0.00001 0.00001 0.00001
0.000011
1 1 1 1 1 1
1 -1 0 -1 0 0
0 1 -1 0 0 1
0 0 0 1 -1 -1
0 0 1 0 1 0
0 4 3 1 -100 0 0 0 0 0

```

Figure 82. Input file for the problem.

1. The node count.
2. The pipe count.
3. Pipe diameters, in feet.
4. Pipe lengths, in feet.
5. Pipe roughness heights, in feet.
6. Kinematic viscosity in feet²/second.
7. Initial guess of flow rates (unbalanced OK, non-zero vital!)
8. The next four rows are the node-arc incidence matrix.
9. The last row is the demand (and fixed-grade node total head) vector.

9.1.1 Support Functions

The Reynolds number will need to be calculated for each pipe at each iteration of the solution, so a Reynolds number function will be useful. For circular pipes, the following equation should work,

$$Re(Q) = \frac{8L}{\mu\pi D}|Q| \quad (93)$$

The Jain equation (Jain, 1976) that directly computes friction factor from Reynolds number, diameter, and roughness is

$$f(k_s, D, Re) = \frac{0.25}{[\log(\frac{k_s}{3.7D} + \frac{5.74}{Re^{0.9}})]^2} \quad (94)$$

Once you have the Reynolds number for a pipe, and the friction factor, then the head loss factor that will be used in the coefficient matrix (and the Jacobian) is

$$L_i = f \frac{8L_i}{\pi^2 g D_i^5} |Q_i| \quad (95)$$

These three support functions are coded in **R** as shown in Listing 32.

Listing 32. R Code to compute Reynolds numbers and friction factors

```
#####
##### Forward Define Support Functions #####
#####
# Jain Friction Factor Function -- Tested OK 23SEP16
friction_factor <- function(roughness,diameter,reynolds){
  temp1 <- roughness/(3.7*diameter);
  temp2 <- 5.74/(reynolds^(0.9));
  temp3 <- log10(temp1+temp2);
  temp3 <- temp3^2;
  friction_factor <- 0.25/temp3;
  return(friction_factor)
}
# Velocity Function
velocity <- function(diameter,discharge){
  velocity <- discharge/(0.25*pi*diameter^2)
  return(velocity)
}
# Reynolds Number Function
reynolds_number <- function(velocity,diameter,mu){
  reynolds_number <- abs(velocity)*diameter/mu
  return(reynolds_number)
}
# Geometric factor function
k_factor <- function(howlong,diameter,gravity){
  k_factor <- (16*howlong)/(2.0*gravity*pi^2*diameter^5)
  return(k_factor)
}
```

9.1.2 Augmented and Jacobian Matrices

The $\mathbf{A}(\mathbf{x})$ is built using the node-arc incidence matrix (which does not change), and the current values of L_i . You will also need to build the Jacobian of $\mathbf{A}(\mathbf{x})$ to implement the update as-per Newton-Raphson.

A brief review; at the solution we can write

$$[\mathbf{A}(\mathbf{x})] \cdot \mathbf{x} - \mathbf{b} = \mathbf{f}(\mathbf{x}) = \mathbf{0} \quad (96)$$

Lets assume we are not at the solution, so we need a way to update the current value of \mathbf{x} . Recall from Newton's method (for univariate cases) that the update formula is

$$x_{k+1} = x_k - \left(\frac{df}{dx} \Big|_{x_k} \right)^{-1} f(x_k) \quad (97)$$

The Jacobian will play the role of the derivative, and \mathbf{x} is now a vector (instead of a single variable). Division is not defined for matrices, but the multiplicative inverse is (the inverse matrix), and plays the role of division. Hence, the extension to the pipeline case is

$$\mathbf{x}_{k+1} = \mathbf{x}_k - [\mathbf{J}(\mathbf{x}_k)]^{-1} \mathbf{f}(\mathbf{x}_k) \quad (98)$$

where $\mathbf{J}(\mathbf{x}_k)$ is the Jacobian of the coefficient matrix \mathbf{A} evaluated at \mathbf{x}_k . Although a bit cluttered, here is the formula for a single update step, with the matrix, demand vector, and the solution vector in their proper places.

$$\mathbf{x}_{k+1} = \mathbf{x}_k - [\mathbf{J}(\mathbf{x}_k)]^{-1} \{[\mathbf{A}(\mathbf{x}_k)] \cdot \mathbf{x}_k - \mathbf{b}\} \quad (99)$$

As a practical matter we actually never invert the Jacobian³⁷, instead we solve the related Linear system of

$$[\mathbf{J}(\mathbf{x}_k)] \cdot \Delta \mathbf{x} = \{[\mathbf{A}(\mathbf{x}_k)] \cdot \mathbf{x}_k - \mathbf{b}\} \quad (100)$$

for $\Delta \mathbf{x}$, then perform the update as $\mathbf{x}_{k+1} = \mathbf{x}_k - \Delta \mathbf{x}$

The Jacobian of the pipeline model is a matrix with the following properties:

1. The partition of the matrix that corresponds to the node formulas (upper left partition) is identical to the original coefficient matrix — it will be comprised of 0 or ± 1 in the same pattern at the equivalent partition of the \mathbf{A} matrix.
2. The partition of the matrix that corresponds to the pipe head loss terms (lower left partition), will consist of values that are twice the values of the coefficients in the original coefficient matrix (at any supplied value of \mathbf{x}_k).
3. The partition of the matrix that corresponds to the head terms (lower right partition), will consist of values that are identical to the original matrix.
4. The partition of the matrix that corresponds to the head coefficients in the node equations (upper right partition) will also remain unchanged.

You will want to take advantage of problem structure to build the Jacobian (you could just finite-difference the coefficient matrix to approximate the partial derivatives, but that is terribly inefficient if you already know the structure).

9.1.3 Stopping Criteria, and Solution Report

You will need some way to stop the process – the three most obvious (borrowed from Newton’s method) are:

1. Approaching the correct solution (e.g. $[\mathbf{A}(\mathbf{x})] \cdot \mathbf{x} - \mathbf{b} = \mathbf{f}(\mathbf{x}) = \mathbf{0}$).

³⁷Inverting the matrix every step is computationally inefficient, and unnecessary. As an example, solving the system in this case would at worst take 10 row operations each step, but nearly 100 row operations to invert at each step – to accomplish the same result, generate an update. Now imagine when there are hundreds of nodes and pipes!

2. Update vector is not changing (e.g. $\mathbf{x}_{k+1} = \mathbf{x}_k$), so either have an answer, or the algorithm is stuck.
3. You have done a lot of iterations (say 100).

Listing 33 is a code fragment to find the flow distribution and heads for the example problem. Not listed is the forward defined functions already listed above – these should be placed into the script in the location shown (or directly sourced into the code in **R**).

Listing 33. R Code to Implement Pipe Network Solution

This fragment reads the data file and converts it into numeric values and reports back the values.

```
# Steady Flow in a Pipe Network Using Hybrid Method (and Newton-Raphson) based on
# Haman YM, Brameller A. Hybrid method for the solution of piping networks. Proc IEEE
# 1971;118(11):1607?12.
#
# Clear all existing objects
rm(list=ls())
#####
#####Forward Define Support Functions Go Here #####
#####
# Read Input Data Stream from File
zz <- file("PipeNetwork.txt", "r") # Open a connection named zz to file named PipeNetwork.
txt
nodeCount <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
pipeCount <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
diameter <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
distance <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
roughness <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
viscosity <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
flowguess <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
nodearcs <- (readLines(zz, n = nodeCount, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
rhs_true <- (readLines(zz, n = pipeCount+nodeCount, ok = TRUE, warn = TRUE,encoding = "
unknown", skipNul = FALSE))
close(zz) # Close connection zz
#
# Convert Input Stream into Numeric Structures
diameter <-as.numeric(unlist(strsplit(diameter,split=" ")))
distance <-as.numeric(unlist(strsplit(distance,split=" ")))
roughness <-as.numeric(unlist(strsplit(roughness,split=" ")))
viscosity <-as.numeric(unlist(strsplit(viscosity,split=" ")))
flowguess <-as.numeric(unlist(strsplit(flowguess,split=" ")))
nodearcs <-as.numeric(unlist(strsplit(nodearcs,split=" ")))
rhs_true <-as.numeric(unlist(strsplit(rhs_true,split=" ")))
# convert nodearcs a matrix
# We will need to augment this matrix for the actual solution -- so after augmentation will
deallocate the memory
nodearcs <-matrix(nodearcs,nrow=nodeCount,ncol=pipeCount,byrow = TRUE)
# echo input
message("Node Count = ",nodeCount)
message("Pipe Count = ",pipeCount)
message("Pipe Lengths = "); distance
message("Pipe Diameters = "); diameter
message("Pipe Roughness = "); roughness
message("Fluid Viscosity = ",viscosity)
message("Initial Guess = "); flowguess
message("Node-Arc-Incidence Matrix = "); nodearcs
#
```

Listing 34 is a code fragment to construct the coefficient matrix structure for the non-changing part and allocate variables for the Newton-Raphson method.

Listing 34. R Code to Implement Pipe Network Solution

This fragment constructs the initial $\mathbf{A}(\mathbf{x})$ matrix and allocates variables used in the iteration loop.

```
# create the augmented matrix
headCount <- nodeCount
flowCount <- pipeCount
augmentedRowCount <- nodeCount+pipeCount
augmentedColCount <- flowCount+headCount
augmentedMat <- matrix(0,nrow=augmentedRowCount,ncol=augmentedColCount,byrow = TRUE)
#
augmentedMat
# build upper left partition of matrix -- this partition is constants from node-arc matrix
for (i in 1:nodeCount){
  for (j in 1:flowCount){
    augmentedMat[i,j] <- nodearcs[i,j]
  }
}
augmentedMat
# build lower right partition of matrix -- this partition is -1*transpose(node-arc) matrix
istart <- nodeCount+1
iend <- nodeCount+pipeCount
jstart <- flowCount+1
jend <- flowCount+headCount
for (i in istart:iend ){
  for(j in jstart:jend ){
    augmentedMat[i,j] <- -1*nodearcs [j-jstart+1,i-istart+1]
  }
}
augmentedMat
# here it should be safe to delete the nodearc matrix
rm(nodearcs)
# Need some vorking vectors
HowMany <- 50
tolerance1 <- 1e-24
tolerance2 <- 1e-24
velocity_pipe <-numeric(0)
reynolds <- numeric(0)
friction <- numeric(0)
geometry <- numeric(0)
lossfactor <- numeric(0)
jacbMatrix <- matrix(0,nrow=augmentedRowCount,ncol=augmentedColCount,byrow = TRUE)
gq <- numeric(0)
solvecguess <- numeric(length=augmentedRowCount)
solvecnew <- numeric(length=augmentedRowCount)
solvecguess[1:flowCount] <- flowguess [1:flowCount]

# compute geometry factors (only need once, goes outside iteration loop)
for (i in 1:pipeCount)
{
  geometry[i] <- k_factor(distance[i],diameter[i],32.2)
}
geometry
```

Listing 35 is the code fragment that implements the iteration loop of the Newton-Raphson method. Within each iteration, the support functions are repeatedly used to construct the changing part of the coefficient and Jacobian matrices, solving the resulting linear system, performing the vector update, and testing for stopping.

Listing 35. R Code to Implement Pipe Network Solution

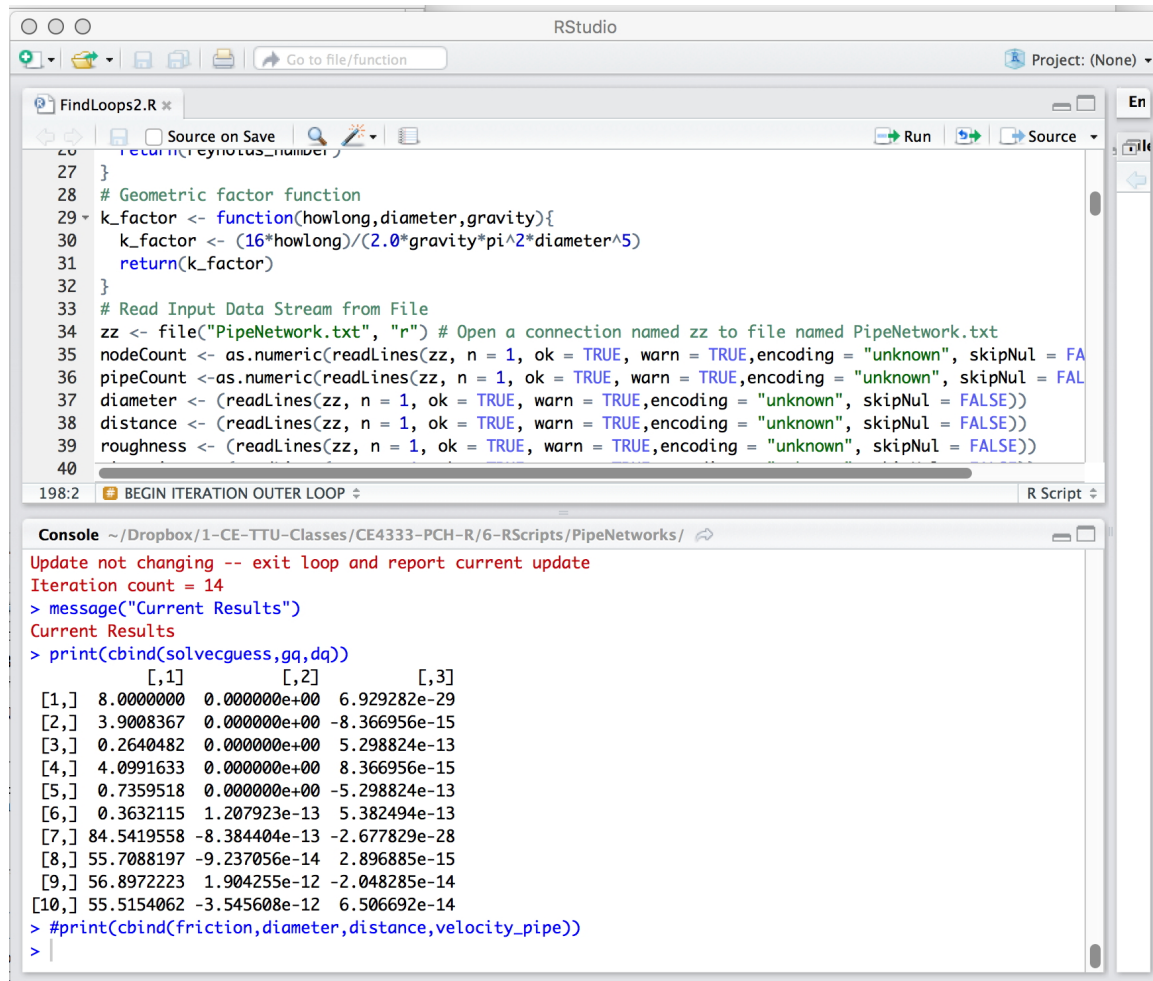
This fragment executes the iteration loop where the Newton-Raphson method and updates are implemented.

```

# going to wrap below into an iteration loop -- first a single instance
for (iteration in 1:HowMany){
##### BEGIN ITERATION OUTER LOOP #####
# compute current velocity
for (i in 1:pipeCount)
{
  velocity_pipe[i]<-velocity(diameter[i],flowguess[i])
}
# compute current reynolds
for (i in 1:pipeCount)
{
  reynolds[i]<-reynolds_number(velocity_pipe[i],diameter[i],viscosity)
}
# compute current friction factors
for (i in 1:pipeCount)
{
  friction[i]<-friction_factor(roughness[i],diameter[i],reynolds[i])
}
# compute current loss factor
for (i in 1:pipeCount)
{
  lossfactor[i] <- friction[i]*geometry[i]*abs(flowguess[i])
}
# build the function matrix
# operate on the lower left partition of the matrix
istart <- nodeCount+1
iend <- nodeCount+pipeCount
jstart <- 1
jend <- flowCount
for (i in istart:iend){
  for(j in jstart:jend){
    if ((i-istart+1) == j)  augmentedMat[i,j] <- -1*lossfactor[j]
  }
}
# now build the current jacobian
# slick trick -- we will copy the current function matrix, then modify the lower left
  partition
  jacobMatrix <- augmentedMat
# build the function matrix
# operate on the lower left partition of the matrix
istart <- nodeCount+1
iend <- nodeCount+pipeCount
jstart <- 1
jend <- flowCount
for (i in istart:iend){
  for(j in jstart:jend){
    if ((i-istart+1) == j)  jacobMatrix[i,j] <- 2*jacobMatrix[i,j]
  }
}
# now build the gq() vector
gq <- augmentedMat %*% solvecguess - rhs_true
gq
dq <- solve(jacobMatrix,gq)
# update the solution vector
solvecnew <- solvecguess - dq
solvecnew
# # now test for stopping
test <- abs(solvecnew - solvecguess)
if( t(test) %*% test < tolerance1){
  message("Update not changing -- exit loop and report current update")
  message("Iteration count = ",iteration)
  solvecguess <- solvecnew
  flowguess[1:flowCount] <- solvecguess[1:flowCount]
  break
}
test <- abs(gq)
if( t(test) %*% test < tolerance2 ){
  message("G(Q) close to zero -- exit loop and report current update")
  message("Iteration count = ",iteration)
  solvecguess <- solvecnew
  flowguess[1:flowCount] <- solvecguess[1:flowCount]
  break
}
solvecguess <- solvecnew
flowguess[1:flowCount] <- solvecguess[1:flowCount]
##### END OF ITERATION OUTER LOOP #####
}
message("Current Results")
print(cbind(solvecguess,gq,dq))
print(cbind(friction,diameter,distance,velocity_pipe))

```

Figure 83 is a screen capture of the script running the example problem. The first column in the output is the solution vector. The first 6 rows are the flows in pipes P1-P6. The remaining 4 rows are the heads at nodes N1-N4.



```

20 }
21 return(reynolds_number)
22 }
23 }
24 # Geometric factor function
25 k_factor <- function(howlong,diameter,gravity){
26   k_factor <- (16*howlong)/(2.0*gravity*pi^2*diameter^5)
27   return(k_factor)
28 }
29 # Read Input Data Stream from File
30 zz <- file("PipeNetwork.txt", "r") # Open a connection named zz to file named PipeNetwork.txt
31 nodeCount <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul = FA
32 pipeCount <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul = FAL
33 diameter <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul = FALSE))
34 distance <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul = FALSE))
35 roughness <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul = FALSE))
36 }
37 }
38 }
39 }
40 }
198:2 BEGIN ITERATION OUTER LOOP
Update not changing -- exit loop and report current update
Iteration count = 14
> message("Current Results")
Current Results
> print(cbind(solvecguess,gq,dq))
      [,1]      [,2]      [,3]
[1,]  8.0000000  0.000000e+00  6.929282e-29
[2,]  3.9008367  0.000000e+00 -8.366956e-15
[3,]  0.2640482  0.000000e+00  5.298824e-13
[4,]  4.0991633  0.000000e+00  8.366956e-15
[5,]  0.7359518  0.000000e+00 -5.298824e-13
[6,]  0.3632115  1.207923e-13  5.382494e-13
[7,]  84.5419558 -8.384404e-13 -2.677829e-28
[8,]  55.7088197 -9.237056e-14  2.896885e-15
[9,]  56.8972223  1.904255e-12 -2.048285e-14
[10,] 55.5154062 -3.545608e-12  6.506692e-14
> #print(cbind(friction,diameter,distance,velocity_pipe))
>

```

Figure 83. Screen capture of R script for pipe network analysis.

9.2 Exercises

- Figure 84 is a five-pipe network with a water supply source at Node 1, and demands at Nodes 1-5. Table 5 is a listing of the node and pipe data.




Figure 84. Layout of Simple Network.

Table 5. Node and Pipe Data.

Pipe ID	Diameter (inches)	Length (feet)	Roughness (feet)
P1	8	800	0.00001
P2	8	700	0.00001
P3	8	700	0.00001
P4	8	800	0.00001
P5	6	600	0.00001
Node ID	Demand (CFS)	Elevation (feet)	Head (feet)
N1	2.0	0.0	100
N2	4.0	0.0	?
N3	3.0	0.0	?
N4	1.0	0.0	?

Code the script, build an input file, and determine the flow distribution. In your solution you are to supply

- An analysis showing the development of the node-arc incidence matrix based on the flow directions in Figure 84,
 - The input file you constructed to provide the simulation values to your script, and
 - A screen capture (or output file) showing the results.
- Code the script and determine the flow distribution in Figures 85 and 86. Assume Node N1 has a total head of 300 feet.

In your solution you are to supply




Figure 85. Pipe network for illustrative example with supply and demands identified. Pipe lengths (in feet) and diameters (in feet) are also depicted..




Figure 86. Pipe network for illustrative example with pipes and nodes labeled..

- An analysis showing the development of the node-arc incidence matrix based on the flow directions in Figure 86,
- The input file you constructed to provide the simulation values to your script, and
- A screen capture (or output file) showing the results.

3. Modify the script to include node elevation information to compute pressures.
Assume all nodes are at elevation 200 feet.

10 Pumps and Valves

The addition of pumps, turbines, and valves increases some of the complexity for a network simulator. Valves and other fittings like elbows and such, that have a fixed setting are modeled as links and the resulting equations look much like pipe loss equations.

Pumps while also logically categorized as links are more complex because their head loss behavior is firstly negative – that is they add head to a flow system, and their ability to actually function is governed by their own performance curve. First we will revieve the modified Bernoulli equation again and then construct a prototype pump function to add to the program and simulate pump performance.

10.1 Energy Loss (along a link)

Equation 101 is the one-dimensional steady flow form of the energy equation typically applied for pressurized conduit hydraulics.

$$\frac{p_1}{\rho g} + \alpha_1 \frac{V_1^2}{2g} + z_1 + h_p = \frac{p_2}{\rho g} + \alpha_2 \frac{V_2^2}{2g} + z_2 + h_t + h_l \quad (101)$$

where $\frac{p}{\rho g}$ is the pressure head at a location, $\alpha \frac{V^2}{2g}$ is the velocity head at a location, z is the elevation, h_p is the added head from a pump, h_t is the added head extracted by a turbine, and h_l is the head loss between sections 1 and 2. Figure 87 is a sketch that illustrates the various components in Equation 101.

In network analysis this energy equation is applied to a link that joins two nodes. Pumps and turbines would be treated as separate components (links) and their hydraulic behavior must be supplied using their respective pump/turbine curves.

10.1.1 Added Head — Pumps

The head supplied by a pump is related to the mechanical power supplied to the flow. Equation 102 is the relationship of mechanical power to added pump head.

$$\eta P = Q \rho g h_p \quad (102)$$

where the power supplied to the motor is P and the “wire-to-water” efficiency is η .

If the relationship is re-written in terms of added head³⁸ the pump curve is

$$h_p = \frac{\eta P}{Q \rho g} \quad (103)$$

³⁸A negative head loss!




Figure 5-1 Definition sketch for terms in the energy equation

Figure 87. Definition sketch for energy equation.

Figure 88 is a typical pump curve depicting the kind of information available from a manufacturer of a pump.




Figure 88. Pump Curve.

In introductory fluid mechanics we spend effort to match the pump curve to the system curve (head losses in our distribution system) and that match tells us how the pump-system combination should function. The pump curve relationship, as well as Equation 103, illustrates that as discharge increases (for a fixed power) the added head decreases. Power scales at about the cube of discharge, so pump curves for

computational application typically have a mathematical structure like

$$h_p = H_{shutoff} - K_{pump} Q^{\text{exponent}} \tag{104}$$

In computational hydraulics we will need to represent the added as a head loss term (with opposite sign), and the functional form represented by Equation 104 is a good starting point. Practical (professional) programs will allow the curve to be represented in a tabular form and will use interpolation (just like our examples earlier) to specify the added head at a particular flow rate.

The next example will illustrate how to add pumps into the model.

Example 1: Pipe network with pumps

Figure 89 is a sketch of the problem that will be used. The network supply is the fixed-grade node in the upper left hand corner of the drawing – in this example its head is set at zero. The remaining nodes (N1 – N4) have demands specified as the purple outflow arrows. The pipes are labeled (P2 – P6), and the red arrows indicate a positive flow direction, that is, if the flow is in the indicated direction, the numerical value of flow (or velocity) in that link would be a positive number. The pump replaces




Figure 89. Pipe Network with a Pump.

pipe (P1) from the previous version of this example. We will use the observation that we really only need to identify which links are pumps, substitute in the correct added head component and then solve the system as in the earlier example.

We have to specify how the pump curve will be represented. In this example we will use a functional form.

$$h_p(Q) = H_{shutoff} - K_{pump} \times Q^n \tag{105}$$

For this example we will use the following numerical values for the pump function: $H_{shutoff} = 104.54$ feet, $K_{pump} = 0.25$ feet/cfs², and $n = 2$.

The $h_p(Q)$ is actually written as an added head factor, just like the friction factor, so we will use absolute values of flow so the term at each computational step can be placed in the augmented matrix as if it were a head loss term; the solver will not know the difference.

The actual functional form employed is

$$h_p(Q) = [H_{shutoff}/|Q| - K_{pump} \times |Q|]Q \quad (106)$$

As before the sign of Q at the solution conveys flow direction. The program example does not trap the potential divide by zero error $H_{shutoff}/|Q|$, but one could test for zero flow, and just apply the shutoff head. Listing 36 implements the prototype function described above.

Listing 36. R Code to pump prototype function

```
.....
# Pump Curve factor function
p_factor <- function(shutoff,constant,exponent,flow){
  p_factor <- shutoff/abs(flow) - constant*abs(flow^(exponent-1))
  return(p_factor)
}
```

Next we have to read in the pump characteristics, I decided to just have pumps replace links (so I won't have to rebuild a node-arc-incidence matrix), so the pump characteristics are

1. Link ID – the index of the pipe that is replaced by a pump.
2. Shutoff head.
3. K_{pump} .
4. Exponent on the pump curve, n . Typically it will be larger than 1.0.

Listing 37 implements the reads from the input file, and builds the pump matrix.

Listing 37. R Code to include pumps in a pipeline network

```
# Read Input Data Stream from File
zz <- file("PipeNetwork.txt", "r") # Open a connection named zz to file named PipeNetwork.
txt
pumpCount <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
nodeCount <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
.....
rhs_true <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
pumps <- (readLines(zz, n = pumpCount, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
  = FALSE))
close(zz) # Close connection zz
.....
pumps <-as.numeric(unlist(strsplit(pumps,split=" ")))
# convert nodearcs a matrix
# We will need to augment this matrix for the actual solution -- so after augmentation will
  deallocate the memory
nodearcs <-matrix(nodearcs,nrow=nodeCount,ncol=pumpCount,byrow = TRUE)
pumps <-matrix(pumps,nrow=pumpCount,ncol=4,byrow=TRUE)
.....
```

Next we will have to compute the added head factor at each step, just like the friction factor, and we will overwrite the pipe that the pump replaces.³⁹

Listing 38 implements the computation of the added head factor, and the pump selection factor.

Listing 38. R Code to include pumps in a pipeline network

```

.....
# compute the current pump factor
if(pumpCount > 0){
  for (i in 1:pumpCount)
  {
    addedhead[i] <- p_factor(pumps[i,2],pumps[i,3],pumps[i,4],flowguess[pumps[i,1]])
  }
}
# build the function matrix
# operate on the lower left partition of the matrix
istart <- nodeCount+1
iend <- nodeCount+pipeCount
jstart <- 1
jend <- flowCount
for (i in istart:iend){
  for(j in jstart:jend){
    if ((i-istart+1) == j) {augmentedMat[i,j] <- -1*lossfactor[j];
      if(pumpCount > 0){
        for(ipump in 1:pumpCount) {
          if(j == pumps[ipump,1]) augmentedMat[i,j] <- addedhead[ipump]
        }
      }
    }
  }
}
# print(augmentedMat)
.....

```

The remainder of the code is unchanged. Listing 39 illustrates the changes in the input file. We have added a row to indicate how many pumps will be used as the first record in the file. The last record after the right-hand side vector is the pump characteristics; one row for each pump. The scripts also test if there are zero pumps and skip code as needed. Observe we still preserve Link #1 data because its part of the node-arc matrix, but the length and diameter of the link is irrelevant (but need to be non-zero because we compute friction factors as if there were a pipe, but never use them).

Listing 39. Input file with pumps at link#1 in a pipeline network

```

1 <== how many pumps
4
6
1.00 0.67 0.67 0.67 0.67 0.5 <== link #1 needs values as placeholders, but are not used
800 800 700 700 800 600
0.00001 0.00001 0.00001 0.00001 0.00001 0.00001
0.000011
1 1 1 1 1 1
1 -1 0 -1 0 0
0 1 -1 0 0 1
0 0 0 1 -1 -1
0 0 1 0 1 0
0 4 3 1 0 0 0 0 0
1 100.54 0.25 2.0 <== Pump Link ID, H\_shutoff, K\_pump, Exponent

```

Figure 90 is a screen capture of the example problem run in **R Studio**. The script produces the correct flow values, and the pump specified was intended to match the

³⁹This approach is decidedly a hack for illustration purposes. A more advanced program would probably just treat everything as a link and use a similar database build structure to determine if a link is a head loss or head add link. My reasoning is that there will be fewer pumps than pipes in any system, so overwriting a fictitious pipe is not too much trouble.

The screenshot shows the RStudio interface with a script editor and a console. The script editor contains R code for a pipe network simulation, and the console shows the output of the simulation.

```

211 }
212 test <- abs(gq)
213 if( t(test) %% tolerance2 ){
214   message("G(Q) close to zero -- exit loop and report current update")
215   message("Iteration count = ",iteration)
216   solveguess <- solvecnew
217   flowguess[1:flowCount] <- solveguess[1:flowCount]
218   break
219 }
220 solveguess <- solvecnew
221 flowguess[1:flowCount] <- solveguess[1:flowCount]
222 ##### END OF ITERATION OUTER LOOP #####
223 }
224 message("Current Results")
225 print(cbind(solveguess,gq,dq))
226 #print(cbind(friction,diameter,distance,velocity_pipe))
227 for (ipump in 1:pumpCount)
228   message("Pump # ",ipump," at Q = ",solveguess[pumps[ipump,1]])
229
230
226:56 END OF ITERATION OUTER LOOP

```

```

> source("~/Dropbox/1-CE-TTU-Courses/CE4333-PCH-R/1-Lectures/Lecture10/ScriptsInLecture/PipeLineModel-Pumps.R")
Node Count = 4
Pipe Count = 6
Pump Count = 1
Update not changing -- exit loop and report current update
Iteration count = 14
Current Results
      [,1]      [,2]      [,3]
[1,] 8.0000000 0.000000e+00 -1.110223e-16
[2,] 3.9008367 0.000000e+00 -8.422430e-15
[3,] 0.2640482 0.000000e+00 5.298161e-13
[4,] 4.0991633 -1.110223e-16 8.311408e-15
[5,] 0.7359518 0.000000e+00 -5.299271e-13
[6,] 0.3632115 1.136868e-13 5.382385e-13
[7,] 84.5400000 -8.384404e-13 -2.346456e-15
[8,] 55.7068640 -9.947598e-14 8.475925e-15
[9,] 56.8952666 1.904255e-12 -1.497464e-14
[10,] 55.5134505 -3.545608e-12 7.074308e-14
Pump # 1 at Q = 8
>

```

Figure 90. Pipe Network with a Pump.

previous problem closely in that it produces enough head so that node N1 has nearly the same head value as the problem without a pump.

10.1.2 Fitting (Minor) Losses

In addition to head loss in the conduit, other losses are created by inlets, outlets, transitions, and other connections in the system. In fact such losses can be used to measure discharge (think of the orifice plate in the fluids laboratory). The fittings create additional turbulence that generates heat and produces the head loss.

Equation 107 is the typical loss model

$$h_{minor} = K \frac{V^2}{2g} \quad (107)$$

where K is called a minor loss coefficient, and is tabulated (e.g. Table 6) for various kinds of fittings.

The use is straightforward, and multiple fittings are summed in the loss term in the energy equation. In practical computation, these losses make the most sense when

Table 6. Minor Loss Coefficients for Different Fittings.

Fitting Type	K
Tee, Flanged, Line Flow	0.2
Tee, Threaded, Line Flow	0.9
Tee, Flanged, Branched Flow	1.0
Tee, Threaded, Branch Flow	2.0
Union, Threaded	0.08
Elbow, Flanged Regular 90°	0.3
Elbow, Threaded Regular 90°	1.5
Elbow, Threaded Regular 45°	0.4
Elbow, Flanged Long Radius 90°	0.2
Elbow, Threaded Long Radius 90°	0.7
Elbow, Flanged Long Radius 45°	0.2
Return Bend, Flanged 180°	0.2
Return Bend, Threaded 180°	1.5
Globe Valve, Fully Open	10
Angle Valve, Fully Open	2
Gate Valve, Fully Open	0.15
Gate Valve, 1/4 Closed	0.26
Gate Valve, 1/2 Closed	2.1
Gate Valve, 3/4 Closed	17
Swing Check Valve, Forward Flow	2
Ball Valve, Fully Open	0.05
Ball Valve, 1/3 Closed	5.5
Ball Valve, 2/3 Closed	200
Diaphragm Valve, Open	2.3
Diaphragm Valve, Half Open	4.3
Diaphragm Valve, 1/4 Open	21
Water meter	7

associated with a particular pipe. If we rewrite the loss equation

$$h_{minor} = \frac{K}{2g} \frac{16Q^2}{\pi^2 D^4} \quad (108)$$

we see that these terms can be added to a pipe either as an additional loss term and placed in the augmented matrix in the same way as the other loss term.

10.1.3 Extracted Head — Turbines

The head recovered by a turbine is also an “added head” but appears on the loss side of the equation. Equation 109 is the power that can be recovered by a turbine (again using the concept of “water-to-wire” efficiency is

$$P = \eta Q \rho g h_t \quad (109)$$

An approach similar to pumps would be employed — the effort in all these cases is to represent the hydraulic components as a loss factor so the non-linear solver we have already built can be used.

11 Pipeline Transients — Water Hammer

Unsteady flow in closed conduits is important for estimating the forces involved from a sudden change in discharge from a pump failing (or starting), or the closing (or opening) of a valve. The flow variation will create a pressure wave traveling along the pipe.⁴⁰ The computational goal is to estimate the magnitude and timing of these extreme pressures to evaluate the safety of the conduit, or design a pump shutdown (or startup) or valve operation protocol to control these extreme pressures to some acceptable magnitude.

11.1 Analysis

If the conduit walls are treated as elastic material, and the liquid is compressible the velocity of a wave along the conduit is

$$c = \sqrt{\frac{1}{\rho\left(\frac{1}{E_f} + \frac{D}{E_c e}\right)}} \quad (110)$$

The value c is called the celerity, E_f is the fluid modulus of elasticity, D is the conduit diameter, E_c is the pipe material modulus of elasticity, e is the conduit wall thickness. Typical values for water are $E_f = 2.2 \times 10^9$ Pa and for steel $E_c = 160 \times 10^9$ Pa.

The pipe is approximated as a series of steel rings (all in line, negligible Poisson ratio) with uniform internal pressure in each ring (but can be different for adjacent rings). The relative pressure head inside the pipe at any given ring is related to the pipe diameter as

$$H = \frac{2c^2}{g} \times \frac{\frac{d}{2} - \frac{d_0}{2}}{\frac{d}{2}} \quad (111)$$

where $\frac{d}{2}$ is the radius under pressure head H and $\frac{d_0}{2}$ is the radius when pressure head is zero gage.⁴¹

The continuity equation is written for each ring, and a force balance is written between each ring.⁴²

⁴⁰These waves will travel in alternating directions as they find boundaries as each end of the disturbed discharge conduit – in some sense the pipeline becomes a resonant chamber. Over time the magnitude of the waves will decrease as friction dissipates the energy. During these transients high and low pressures are applied to the pipe walls and fixtures and can conceivably damage them.

⁴¹1 atmosphere absolute. It is usually easier to work with gage pressure in practice.

⁴²The time to drain problem is mildly similar the tanks represent a ring where the change in depth (head) is related to outflow velocity; the outflow velocity is related to force at the outlet – in that case the pressure force of the water above the outlet; Bernoulli's equation for that case simplified the work considerably. Here we will have a series of tanks (rings) that communicate pressure head and velocity between them. We will arrive at a staggered spatial and temporal grid.

The resulting coupled set of differential equations are

$$\begin{aligned}\frac{\partial V}{\partial t} + V \frac{\partial V}{\partial x} + g \frac{\partial H}{\partial x} &= -\frac{2\tau_0}{r} \\ \frac{\partial H}{\partial t} + V \frac{\partial H}{\partial x} + \frac{c^2}{g} \frac{\partial V}{\partial x} &= 0\end{aligned}\tag{112}$$

The shear wall stress is τ_0 and the entire term is usefully approximated as

$$\frac{2\tau_0}{r} = \frac{fV|V|}{2d}\tag{113}$$

where f is a friction factor (Darcy-Wiesbach). The finite-difference approximation of the above equations is challenging for a variety of reasons, however a useful simplification can be achieved by neglecting the non-linear terms and assuming friction terms are small relative to the time-scale of the problem at a particular location. The resulting simplified model is

$$\begin{aligned}\frac{\partial V}{\partial t} + g \frac{\partial H}{\partial x} + \frac{fV|V|}{2d} &= 0 \\ \frac{\partial H}{\partial t} + \frac{c^2}{g} \frac{\partial V}{\partial x} &= 0\end{aligned}\tag{114}$$

The model is completed by specification of initial and boundary conditions. The velocity and pressure heads are determined before the initiation of the flow change. For the development herein, we will use the simplest boundary conditions which are a specified head at one end of our pipeline (a reservoir) and a specified velocity at the other end of our pipeline. A reasonably simple model of a valve-type condition is

$$V(t) = \frac{A(t)}{A_0} \sqrt{2gH(t)}\tag{115}$$

where $A(t)$ is the time varying open area of the conduit at the valve, and A_0 is the cross sectional area of the conduit when fully open. For an instant close or open $A(t)$ is a step function.

11.2 Explicit Finite-Difference Model

An explicit finite-difference scheme (much like the tank-drain model, but there are coupled components of H and V) can be constructed using a staggered space-time grid like that depicted in Figure 91.

The figure depicts the spatial dimension along the horizontal axis. Each index value ($i, i+1, i+2, \dots$) labeled as a circle in the figure represents the interface between two “rings, and fluid flow values are computed at these locations. Each index value labeled as a triangle in the figure represents the space within a “ring” and the head values are computed at these locations. Thus in the figure, location i (solid vertical lines) is the interface between rings $i-1$ and i . The values of head in the $i-1$ and i ring




Figure 91. Staggered grid (in space and time).

determine the velocity that traverses the ring interface. For instance if the head in $i-1$ is larger than in ring i , flow should be from left to right in that time interval.

A similar staggered structure is depicted in the vertical time domain. Time levels are treated in the same fashion where all the head values are used to update velocity values, then these updated velocity values are in turn used to update the head values.

The boundary conditions are applied at each end of the computational domain – Typically node 1 where head is a fixed value, and velocity depends on the head in node 2, and node N , where velocity is determined from the valve model, and head at $N-1$ is computed.

The application of the finite-difference equations to the linearized set of partial differential equations for any interior node (of either type) is

$$\begin{aligned} \frac{V_i^{n+3/2} - V_i^{n+1/2}}{\Delta t} + g \frac{H_i^{n+1} - H_{i-1}^{n+1/2}}{\Delta x} + \frac{f V_i^{n+1/2} |V_i^{n+1/2}|}{2d} &= 0 \\ \frac{H_i^{n+1} - H_i^n}{\Delta t} + \frac{c^2}{g} \frac{V_{i+1}^{n+1/2} - V_i^{n+1/2}}{\Delta x} &= 0 \end{aligned} \tag{116}$$

11.3 An Example R Implementation

Figure 92 is a sketch showing a pipeline supplied by a reservoir on the left with constant head of 5 meters. The reservoir will be quite large relative to the volume in the pipeline – think of Lake Mead on Boulder Canyon (Hoover Dam), and the pipeline is a supply pipe to a water treatment plant for a nearby community.




Figure 92. Pipeline with a constant head reservoir at the left side, and a valve shutdown at the right side.

The pipeline itself is 6 kilometers long (about 3.5 miles), with a diameter of 0.5 meters (about 18 inches), and the pipe wall thickness is 4 millimeters (a little under 1/4 inch thick). The steady flow velocity with the valve open is 9.9 meters/second (30 feet per second – really fast for real pipes).

The valve is closed over a 2 second interval, with linear decrease in outlet area during the closure. Neglecting frictional losses in the system, compute and plot the pressures at the valve for a 250-second duration.

The script to accomplish these computations is shown in Listing 40. The code includes provision for the friction term, here we will just set it to zero. The code also sets time step size based on the Courant condition. The Courant number must equal unity for the algorithm to be stable.

Listing 40. R Code to compute for a pipeline transient example

```
#####
# Pipeline Transients using Explicit Finite Differences (linearized formulation) #
#####
rm(list=ls()) # deallocate memory
##### prototype functions #####
celerity <- function(density,elasticity_fluid,elasticity_solid,diameter,thickness){
  temp1<- 1.0/elasticity_fluid
  temp2<- diameter/(elasticity_solid*thickness)
  temp3<- temp1 + temp2
  temp4<- density*temp3
  celerity <- sqrt(1.0/temp4)
  return(celerity)
}
### Simulator Code #####
# Simulation Conditions (this section could be replaced with an input file)
# fluid properties
density <- 1000 #kg/m^3
elasticity_fluid <- 2.0e09 #Pa
elasticity_solid <- 160.0e09 #Pa
diameter <- 0.500 #m
thickness <- 0.004 #m
cc <- celerity(density,elasticity_fluid,elasticity_solid,diameter,thickness)
print(cc)
# simulation properties
deltax <- 600 #meters
courantRatio <- 0.99 #select courant ratio to set time step. If bigger than 1 unstable
deltat <- courantRatio*(deltax/cc) # force to be courant number for stability
# allocate head and velocity vectors, assign initial values
startHead <- 5.0 #meters
startVelo <- 9.9 #meters/second
pipeLength <- 6000 #meters
closeTime <- 2.0 #seconds
simulationDuration <- 250 #seconds
frictionFactor <- 0.015 #Moody Chart
cellCount <- as.integer((pipeLength/600)+1)
head <- numeric(0)
velocity <- numeric(0)
for(i in 1:cellCount){
  head[i]<-startHead
  velocity[i]<-startVelo
}
# useful constants
dtdx <- deltax/deltax
cc2g <- cc^2/9.81
do2 <- 1.0/(diameter*2) #used when friction included
# allocate some output vectors
etime<-numeric(0)
headvalve<-numeric(0)
velocitytank<-numeric(0)
# simulation values
etime[1] <- 0
headvalve[1] <- head[cellCount]
velocitytank[1] <-velocity[1]
#####
# Time Stepping Loop #
#####
maxiter <- 1+simulationDuration/deltat
for(itime in 2:maxiter){
  etime[itime]<-etime[itime-1]+deltat
  ##### valve closure model #####
  closeRatio <- 1.0 - (etime[itime]/closeTime)
  if(closeRatio >= 0.0){
    velocity[cellCount] <- closeRatio*sqrt(2*9.8*(head[1]))
  }
  else{
    velocity[cellCount] <- 0
  }
  ##### update velocity #####
  for(i in 1:(cellCount-1)){
    friction <- frictionFactor*velocity[i]*abs(velocity[i])*do2
    velocity[i]=velocity[i]-9.81*dtdx*(head[i+1]-head[i])-deltat*friction
  }
  ##### update head #####
  for(i in 2:(cellCount)){
    head[i]=head[i]-cc2g*dtdx*(velocity[i]-velocity[i-1])
  }
  headvalve[itime]<-head[cellCount]
  velocitytank[itime] <-velocity[1]
}
# report results
message("Maximum head at valve : ",max(headvalve))
message("Minimum head at valve : ",min(headvalve))
# plot results
par(mfrow=c(2,1))
plot(etime,headvalve,type="l",pch=19,lwd=3,tck=1,xlab="Time (seconds)",ylab="Head at Valve
(meters)",col="red")
plot(etime,velocitytank,type="l",pch=19,lwd=3,tck=1,xlab="Time (seconds)",ylab="Velocity at
Tank (meters/sec)",col="blue")
```

The listing seems elaborate, but it is relatively simple. First is a prototype function for the celerity. Next material properties and simulation characteristics are supplied. For the code to be more general, one would probably move a lot of these assignments into a file read structure.

Once everything is ready, the algorithm begins time-stepping. Within any single time step the following activities are performed:

1. Determine the downstream velocity at the valve based on the required closure model.
2. Update the velocity in each interior interface (1 to the section before the valve).
3. Update the head in each interior cell (2 to the last cell, which is upstream of the valve).
4. Record the head at the valve, and the velocity at the tank exit.

When the requisite time steps are completed, then report and plot the results.

Figure 93 is the example presented herein using **R Studio**. The friction factor was set to 0, and the courant ratio was set to 1. The head at the valve oscillates from $+\Delta H$ to $-\Delta H$, and the velocity repeatedly changes direction an value from $+9.9$ to -9.9 meters per second, exactly phase shifted relative to head by one-half cycle. The time for the values to change is twice the pipeline length divided by the wave speed.




Figure 93. Pipeline transient simulation.

Now if we add the friction term say $f = 0.015$ which would be a value for a smooth

pipe (especially at the flow velocities in the example) the effect of frictional damping is evident in the simulation and is displayed in Figure 94

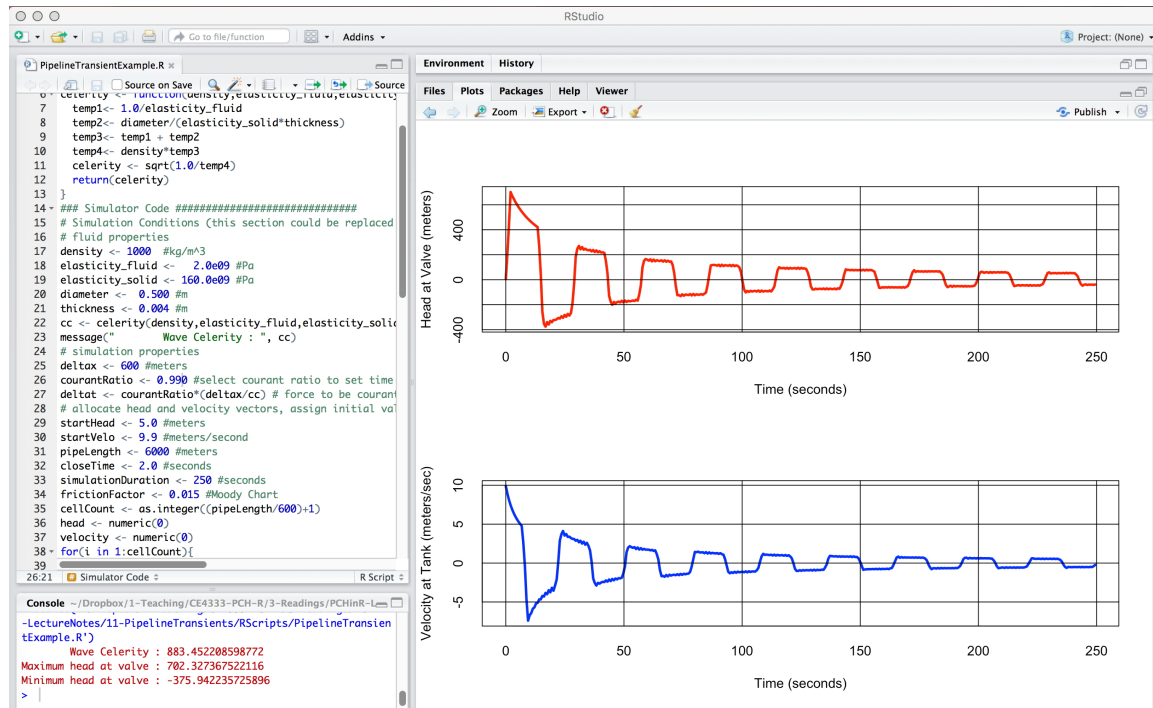


Figure 94. Pipeline simulation with friction.

Exercises

Exercise Set ES-9 On Class Server

12 Open Channel Flow

This chapter derives the St. Venant equations for one-dimensional (1-D) open channel flow. The equations were originally developed in the 1850's, so the concept is not very new. The tools have changed since that time; computational methods have greatly increased the utility of these equations.

12.1 St. Venant Equations

In general, 1-D unsteady flow would be considered state-of-practice computation; every engineer would be expected to be able to make such calculations (albeit using software). 2-D computation is becoming routine using general purpose software. 3-D computation as of this writing (circa 2009) is still in the realm of state-of-art, and would not be within the capability a typical consulting firm.

The conservation of mass, momentum, and energy in the context of the cell balance method is used to develop the mathematical and computational structure. The cell balance is a computational structure that is analogous to the Reynolds transport theorem, except the end result are difference equations that can be updates to approximate physical processes. We later employ the method for porous media flow. The philosophy is a hybrid approach – instead of developing the differential equations first, then numerical approximations, the numerical constructs are built directly and the limiting process is employed to demonstrate that the constructs indeed mimic the differential equations that describe our current understanding of the physics.

12.1.1 The Computational Cell

The cell balance method envisions the world as representable by a computational cell (or more typically a collection of cells) with some finite dimension, fixed in space about a cell centroid. Some dimensions are changeable – such as depth.⁴³

The fundamental computational element is a computational cell or a reach.⁴⁴ Figure 95 is a sketch of a portion of a channel. The left-most section is uphill (and upstream) of the right-most section. The section geometry is arbitrary, but is drawn to look like a channel cross section.

The length of the reach (distance between each section along the flow path) is Δx . The depth of liquid in the section is z , the width at the free surface is $B(z)$, the

⁴³This concept is distinct from a particle view of the world, which will be explored in future versions of this course.

⁴⁴Some professional software, in particular HEC-RAS, considers a reach to be a specific portion of a river system that may be comprised of several computational sub-reaches (cells). The engineer will need to consider the context and the tool used to decide which way to describe their problem — and it is quite possible that the author of this document may be wrong in terminology!

functional relationship established by the channel geometry. The flow into the reach on the upstream face is $Q - \frac{\partial Q}{\partial x} * \frac{\Delta x}{2}$. The flow out of the reach on the downstream face is $Q + \frac{\partial Q}{\partial x} * \frac{\Delta x}{2}$. The direction is strictly a sign convention and the development does not require flow in a single direction. The topographic slope is S_0 , assumed relatively constant in each reach, but can vary between reaches.




Figure 95. Reach/Computational Cell.

12.1.2 Assumptions

The development of the unsteady flow equations in this chapter uses several assumptions:

1. The pressure distribution at any section is hydrostatic — this assumption allows computation of pressure force as a function of depth.
2. Wavelengths are long relative to flow depth — this is called the shallow wave theory.
3. Channel slopes are small enough so that the topographic slope is roughly equal to the tangent of the angle formed by the channel bottom and the horizontal.
4. The flow is one-dimensional — this assumption implies that longitudinal dimension is large relative to cross sectional dimension. Generally river flows will meet this assumption, it fails in estuaries where the spatial dimensions (length and width) are roughly equal. Thus rivers that are hundreds of feet wide imply that reaches are miles long. If this assumption cannot be met, then 2-D

methods are more appropriate.

5. Friction is modeled by Chezy or Manning's type empirical models. The particular friction model does not really matter, but historically these equations have used the friction slope concept as computed from one of these empirical models.

The tools that are used to build the equations are conservation of mass and linear momentum.

12.1.3 Conservation of Mass

The conservation of mass in the cell is the statement that mass entering and leaving the cell is balanced by the accumulation or loss of mass within the cell. For pedagogical clarity, this section goes through each part of a mass balance then assembles into a difference equation of interest.

Mass Entering: Mass enters from the left of the cell in our sketch. This direction only establishes a direction convention and negative flux means the arrow points in the direction opposite of that in the sketch. In the notation of the sketch mass entering in a short time interval is:

$$\dot{M}_{in} = \rho * \left(Q - \frac{\partial Q}{\partial x} * \frac{\Delta x}{2} \right) * \Delta t \quad (117)$$

where ρ is the fluid density. Notice that the mass flux is evaluated at the cell interface and not the centroid, while by convention ρ is assumed to be defined as an average cell property.

Mass Leaving: Mass leaves from the right of the cell in our sketch. In the notation of the sketch mass leaving is:

$$\dot{M}_{out} = \rho * \left(Q + \frac{\partial Q}{\partial x} * \frac{\Delta x}{2} \right) * \Delta t \quad (118)$$

Mass Accumulating: Mass accumulating within the reach is stored in the prism depicted in the sketch by the dashed lines. The product of density and prism volume is the mass added to (or removed from) storage.

The rise in water surface in a short time interval is $\frac{\partial z}{\partial t} * \Delta t$. The plan view area of the prism is $B(z) * \Delta x$. The product of these two terms is the mass added to storage, expressed as:

$$\dot{M}_{storage} = \rho * \left(\frac{\partial z}{\partial t} * \Delta t \right) * B(z) * \Delta x \quad (119)$$

Equating the accumulation to the net inflow produces

$$\rho * \left(\frac{\partial z}{\partial t} * \Delta t \right) * B(z) * \Delta x = \rho * \left(Q - \frac{\partial Q}{\partial x} * \frac{\Delta x}{2} \right) * \Delta t - \rho * \left(Q + \frac{\partial Q}{\partial x} * \frac{\Delta x}{2} \right) * \Delta t \quad (120)$$

This is the mass balance equation for the reach. If the flow is isothermal, and essentially incompressible then the density is a constant and can be removed from both sides of the equation.

$$\left(\frac{\partial z}{\partial t} * \Delta t \right) * B(z) * \Delta x = \left(Q - \frac{\partial Q}{\partial x} * \frac{\Delta x}{2} \right) * \Delta t - \left(Q + \frac{\partial Q}{\partial x} * \frac{\Delta x}{2} \right) * \Delta t \quad (121)$$

Rearranging the right hand side produces

$$\left(\frac{\partial z}{\partial t} * \Delta t \right) * B(z) * \Delta x = -\frac{\partial Q}{\partial x} * \frac{\Delta x}{2} * \Delta t - \frac{\partial Q}{\partial x} * \frac{\Delta x}{2} * \Delta t = -\frac{\partial Q}{\partial x} * \Delta x * \Delta t \quad (122)$$

Dividing both sides by $\Delta x * \Delta t$ yields

$$\left(\frac{\partial z}{\partial t} \right) * B(z) = -\frac{\partial Q}{\partial x} \quad (123)$$

This equation is the conventional representation of the conservation of mass in 1-D open channel flow. If the equation includes lateral inflow the equation is adjusted to include this additional mass term. The usual lateral inflow is treated as a discharge per unit length added into the mass balance as expressed in Equation 151.

$$\left(\frac{\partial z}{\partial t} \right) * B(z) + \frac{\partial Q}{\partial x} = q \quad (124)$$

This last equation is one of the two equations that comprise the St. Venant equations. The other equation is developed from the conservation of linear momentum — the next section.

12.1.4 Conservation of Momentum

The conservation of momentum is the statement of the change in momentum in the reach is equal to the net momentum entering the reach plus the sum of the forces on the water in the reach. As in the mass balance, each component will be considered separately for pedagogical clarity.

Figure 96 is a sketch of the reach element under consideration, on some non-zero sloped surface.

Momentum Entering: Momentum entering on the left side of the sketch is




Figure 96. Equation of motion definition sketch.

$$\rho * QV = \rho * V^2 A \tag{125}$$

Momentum Leaving: Momentum leaving on the right side of the sketch is

$$\rho * QV + \frac{\partial}{\partial x}(\rho * QV)\delta x = \rho * V^2 A + \frac{\partial}{\partial x}(\rho * \rho * V^2 A)\delta x \tag{126}$$

Momentum Accumulating: The momentum accumulating is the rate of change of linear momentum:

$$\frac{dL}{dt} = \frac{d(mV)}{dt} = \frac{\partial}{\partial t}(\rho * AV * \delta x) = \rho * \delta x \frac{\partial}{\partial t}(\rho * AV) \tag{127}$$

Forces on the liquid in the reach:

Gravity forces: The gravitational force on the element is the product of the mass in the element and the downslope component of acceleration.

The mass in the element is $\rho * A\delta x$

The x -component of acceleration is $g \sin(\alpha)$, which is $\approx S_0$ for small values of α .

The resulting force of gravity is is the product of these two values:

$$F_g = \rho g * AS_0 \delta x \tag{128}$$

Friction forces: Friction force is the product of the shear stress and the contact area. In the reach the contact area is the product of the reach length and average wetted perimeter.

$$F_{fr} = \tau * P_w * \delta x \quad (129)$$

where $P_w = A/R$, R is the hydraulic radius. A good approximation for shear stress in unsteady flow is $\tau = \rho g R S_f$. S_f is the slope of the energy grade line at some instant and is also called the friction slope. This slope can be empirically determined by a variety of models, typically Chezy's or Manning's equation is used. In either of these two models, we are using a STEADY FLOW equation of motion to mimic unsteady behavior — nothing wrong, and it is common practice, but this decision does limit the frequency response of the model (the ability to change fast — hence the shallow wave theory assumption!).

The resulting friction model is

$$F_{fr} = \rho g A S_f * \delta x \quad (130)$$

Pressure forces: [Set the equations, backfill discussion next version]

$$F_p = \int_A dF \quad (131)$$




Figure 97. Pressure integral sketch.

$$dF = (z - h)\rho g\xi(h)dh \quad (132)$$

where $\xi(h)$ is the width of the panel at a given distance above the channel bottom (h) at any section.

$$F_{p \text{ net}} = F_{p \text{ up}} - F_{p \text{ down}} \quad (133)$$

$$F_{p \text{ net}} = F_p - (F_p + \frac{\partial F_p}{\partial x} * \delta x) = -\frac{\partial F_p}{\partial x} * \delta x \quad (134)$$

$$-\frac{\partial F_p}{\partial x} * \delta x = -\frac{\partial}{\partial x} [\int_0^Z \rho g(z - h)\xi(h)dh] \delta x \quad (135)$$

$$F_{p \text{ net}} = -\rho g [\frac{\partial z}{\partial x} \int_0^Z \xi(h)dh + \int_0^Z (z - h)\xi(h) \frac{\partial \xi(h)}{\partial x} dh] \delta x \quad (136)$$

The first term integrates to the cross sectional area, the second term is the variation in pressure with position along the channel.

The other pressure force to consider is the bank force (the pressure force exerted by the banks on the element). This force is computed using the same type of integral structure except the order is swapped.

$$F_{p \text{ bank}} = [\int_0^Z \rho g(z - h) \frac{\partial \xi(h)}{\partial x} \delta x] dh \quad (137)$$

Now we put everything together.

$$Momentum_{in} - Momentum_{out} + \sum F = \frac{d(mV)}{dt} \quad (138)$$

Substitution of the pieces:

$$Momentum_{in} - Momentum_{out} + F_{p \text{ net}} + F_{bank} + F_{gravity} - F_{friction} = \frac{d(mV)}{dt} \quad (139)$$

Now when the expressions for each expressions for each part

$$\begin{aligned}
& \rho * V^2 A - \rho * V^2 A - \frac{\partial}{\partial x}(\rho * V^2 A) \delta x \\
& - \rho g \frac{\partial z}{\partial x} \int_0^Z \xi(h) dh \delta x - \left[\int_0^Z \rho g (z - h) \frac{\partial \xi(h)}{\partial x} dh \right] \delta x \\
& + \left[\int_0^Z \rho g (z - h) \frac{\partial \xi(h)}{\partial x} \delta x \right] dh \\
& + \rho g * A S_0 \delta x \\
& - (\rho g R S_f * \delta x) \\
& = \rho * \delta x \frac{\partial}{\partial t} (\rho * g * AV)
\end{aligned} \tag{140}$$

Each row of Equation 140 is in order:

1. Net momentum entering the reach.
2. Pressure force differential at the end sections.
3. Pressure force on the channel sides.
4. Gravitational force.
5. Frictional force opposing flow.
6. Total acceleration in the reach (change in linear momentum).

Canceling terms and dividing by $\rho \delta x$ (isothermal, incompressible flow; reach has finite length) Equation 140 simplifies to

$$-\frac{\partial}{\partial x}(V^2 A) - g \frac{\partial z}{\partial x} \int_0^Z \xi(h) dh + g * A S_0 - (g R S_f *) = \frac{\partial}{\partial t}(g * AV) \tag{141}$$

The second term integral is the sectional flow area, so it simplifies to

$$-\frac{\partial}{\partial x}(V^2 A) - g \frac{\partial z}{\partial x} A + g A S_0 - g A S_f = \frac{\partial}{\partial t}(AV) \tag{142}$$

The term with the square of mean section velocity is expanded by the chain rule, and using continuity becomes (notice the convective acceleration term from the change in area with time)

$$\frac{\partial}{\partial t}(AV) = A \frac{\partial V}{\partial t} + V \frac{\partial A}{\partial t} = A \frac{\partial V}{\partial t} - V A \frac{\partial V}{\partial x} - V^2 \frac{\partial A}{\partial x} \tag{143}$$

Now expand and construct

$$-V^2 \frac{\partial A}{\partial x} - 2VA \frac{\partial V}{\partial x} - gA \frac{\partial z}{\partial x} + gA(S_0 - S_f) = A \frac{\partial V}{\partial t} - VA \frac{\partial V}{\partial x} - V^2 \frac{\partial A}{\partial x} \quad (144)$$

Cancel common terms and simplify

$$-VA \frac{\partial V}{\partial x} - gA \frac{\partial z}{\partial x} + gA(S_0 - S_f) = A \frac{\partial V}{\partial t} \quad (145)$$

Equation 145 is the final form of the momentum equation for practical use. It will be rearranged in the remainder of this essay to fit some other purposes, but this is the expression of momentum in the channel reach.

Divide by gA and obtain

$$-\frac{V}{g} \frac{\partial V}{\partial x} - \frac{\partial z}{\partial x} + (S_0 - S_f) = \frac{1}{g} \frac{\partial V}{\partial t} \quad (146)$$

Now rearrange to place the two slopes on the left side, and the remaining part of momentum to the right side. Equation 147 let's us examine the several flow regimes common in open channel flows.

$$S_0 - S_f = \frac{1}{g} \frac{\partial V}{\partial t} + \frac{\partial z}{\partial x} + \frac{V}{g} \frac{\partial V}{\partial x} \quad (147)$$

If the local acceleration (first term on the right) is zero, the depth taper (middle term on the right) is zero⁴⁵, and the convective acceleration (last term on the right) is zero, then the expression degenerates to the algebraic equation of *normal* flow ($S_0 = S_f$). If just the local acceleration term is zero, and all the remaining terms are considered, then the expression degenerates to the ordinary differential equation of gradually varied flow. Finally, if all the terms are retained, then the dynamic flow (shallow wave) conditions are in effect and the resulting model is a partial differential equation.

Re-iterating these typical flow regimes.

1. Uniform flow; algebraic equation.

$$S_f = S_0 \quad (148)$$

⁴⁵Zero depth taper means constant depth flow.

2. Gradually varied flow; ordinary differential equation.

$$S_f = S_0 - \frac{\partial z}{\partial x} - \frac{V}{g} \frac{\partial V}{\partial x} \quad (149)$$

3. Dynamic flow (shallow wave) conditions; partial differential equation.

$$S_f = S_0 - \frac{\partial z}{\partial x} - \frac{V}{g} \frac{\partial V}{\partial x} - \frac{1}{g} \frac{\partial V}{\partial t} \quad (150)$$

The coupled pair of equations, Equation 151 for continuity and Equation 152 for momentum are called the St. Venant equations and comprise a coupled hyperbolic differential equation system.

$$\left(\frac{\partial z}{\partial t}\right) * B(z) + \frac{\partial Q}{\partial x} = q \quad (151)$$

$$S_0 - S_f - \frac{\partial z}{\partial x} - \frac{V}{g} \frac{\partial V}{\partial x} - \frac{1}{g} \frac{\partial V}{\partial t} = 0 \quad (152)$$

Solutions $((z, t)$ and (V, t) functions) are found by a variety of methods including finite difference, finite element, finite volume, and characteristics methods.

In the next chapter we will examine solutions to the gradually varied flow equation, then proceed to a finite difference solution to the full dynamic equations in the following chapter.

13 Steady Water Surface Profiles

13.1 Steady Uniform Flow

In the previous chapter uniform flow was defined at the situation where friction slope and channel slope are the same. In this kind of flow the profile grade line (channel bottom), the hydraulic grade line (water surface), and the energy grade line are all parallel. The following would be implied

$$S_f = \frac{\Delta h_L}{\Delta x} = S_0 \quad (153)$$

Suppose we apply the Darcy-Weisbach head loss model (and adapt it for non-circular conduit)

$$h_L = f \frac{L}{4R_h} \frac{V^2}{2g} \quad (154)$$

13.2 Steady Gradually Varied Flow

Need a brief theory

13.2.1 Finite Difference Methods — Fixed Depth Change Step Method

The fixed-step refers to fixed changes in depth for which we solve to find the variable spatial steps. The method is a very simple method for computing water surface profiles in prismatic channels. A prismatic channel is a channel of uniform cross sectional geometry⁴⁶ with constant bed (topographic) slope.

In such channels with smooth (non-jump) steady flow the continuity and momentum equations are

$$Q = AV \quad (155)$$

where Q is volumetric discharge, A is cross sectional flow area, and V is the mean section velocity.

and

$$\frac{V}{g} \frac{dV}{dx} + \frac{dh}{dx} = S_o - S_f \quad (156)$$

where h is the flow depth (above the bottom), and x is horizontal the distance along the channel.

⁴⁶Channel geometry is same at any section, thus rectangular, trapezoidal, and circular channels if the characteristic width dimension is constant would be prismatic.

For the variable step method, the momentum equation is rewritten as a difference equation (after application of calculus to gather terms) then rearranged to solve for the spatial step dimension ⁴⁷.

$$\frac{\frac{V_{i+1}^2}{2g} - \frac{V_i^2}{2g}}{\Delta x} + \frac{h_{i+1} - h_i}{\Delta x} = S_o - \bar{S}_f \quad (157)$$

where \bar{S}_f is the average slope of the energy grade line between two sections (along a reach of length Δx , the unknown value).

Rearrangement to isolate Δx produces an explicit update equation that can be evaluated to find the different values of Δx associated with different flow depths. The plot of the accumulated spatial changes versus the sum of the flow depth and bottom elevation is the water surface profile.

$$\frac{(h_{i+1} + \frac{V_{i+1}^2}{2g}) - (h_i + \frac{V_i^2}{2g})}{S_o - \bar{S}_f} = \Delta x \quad (158)$$

The distance between two sections with known discharges is computed using the equation, all the terms on the left hand side are known values. The mean energy gradient (\bar{S}_f) is computed from the mean of the velocity, depth, area, and hydraulic radius for the two sections.

The friction slope can be computed using Manning's, Chezy, or the Darcy-Weisbach friction equations adapted for non-circular, free-surface conduits.

13.2.2 Coding the algorithm in R

Here the method is illustrated in **R** to illustrate the tool as a programming environment⁴⁸.

First we build a set of utility functions, these will be used later in the `backwater` function:

Listing 41 are utility functions for rectangular channels for flow area given channel depth and width for rectangular and wetted perimeter given depth and width. Different geometries will need different functions (probably by numerical methods rather than actual functional relationships).

⁴⁷The equation here is written moving upstream, direction matters for indexing. Thus position $i + 1$ is assumed upstream of position i in this essay. This directional convention is not generally true in numerical methods and analysts need to use care when developing their own tools or using other tools. A clever analyst need not rewrite code, but simple interchange of upstream and downstream depths will handle both backwater and front-water curves.

⁴⁸The **R** code in this essay is broken into parts and some unnecessary line feeds are included to fit the page.

Listing 41. R Code for prototype hydraulic support functions (for rectangular channels)

```

# Depth-Area Function
area<-function(depth,width){
  area<-depth*width;
  return(area)
}

# Wetted perimeter function for rectangular channel
perimeter<-function(depth,width){
  perimeter<-2*depth+width;
  return(perimeter)
}

# 2-point average - useful with friction slopes
avg2point<-function(x1,x2){
  avg2point<-0.5*(x1+x2);
  return(avg2point)
}

```

Listing 42 is a listing of the code for the hydraulic radius (ratio of the above results), this is a generic function, it does not need to know the flow geometry

Listing 42. R Code for prototype hydraulic radius function

```

# Hydraulic radius function
radius<-function(area,perimeter){
  radius<-area/perimeter;
  return(radius)
}

```

Listing 43 is a listing of code for the friction slope given Manning's n , discharge, hydraulic radius, and flow area. Notice that this function implicitly assumes SI units (the 1.49 constant in U.S. Customary units is not present). For U.S. Customary units either add the constant or convert the US units into equivalent SI units.

Listing 43. R Code for prototype friction slope function)

```

# Friction slope function
slope_f<-function(discharge,mannings_n,area,radius){
  slope_f<-((discharge^2)*(mannings_n^2)/((radius^(4/3))*(area^2)));
  return(slope_f)
}

```

The semi-colons in the functions are probably unnecessary, but have value because it forces the expression to its left to be evaluated and helps prevent ambiguous⁴⁹ code. Also notice the use of the scope { } delimiter, the delimiter is required, however there is no requirement to line feed — I simply find it easier to read my own code in this fashion (and count delimiters).

At this point, we would have 5 useful, testable functions (and we should test before the next step).

Listing 44 is the step-backwater method implemented as a function. This function computes the space steps, changes in depth, etc. as per the algorithm. The function is a FORTRAN port, so it is not a terribly efficient use of **R**, but it illustrates count controlled repetition (for loops), array indexing, and use of the utility functions to make the code readable as well as ensure that the parts work before the whole program is assembled. This concept is really crucial, if you can build a tool of parts that are known to work, it helps keep logic errors contained to known locations.

⁴⁹To the human operator; the computer will either accept the code or throw an error.

Listing 44. R Code for prototype friction slope function)

```

# Backwater curve function
backwater<-function(begin_depth,end_depth,how_many,discharge,width,mannings_n,slope){
#
## Example function call will be shown in class and on movie.
#
# Other functions must exist otherwise will spawn errors
#
# Prepare space for vectors
depth<-numeric(0)
bse<-numeric(0)
wse<-numeric(0)
# change in depth for finding spatial steps
delta_depth<-(begin_depth-end_depth)/(how_many)
# assign downstream value
depth[1]<-begin_depth
# depth values to evaluate
for (i in 2:how_many){depth[i]<-depth[1]-i*delta_depth}
#velocity for each depth
velocity<-discharge/area(depth,width)
# numeric vector for space steps (destination space)
deltax<-numeric(0)
# next for loop is very FORTRANesque!
for (i in 1:(how_many-1)){
#compute average depth in reach
depth_bar<-avg2point(depth[i],depth[i+1]);
#compute average area in reach
area_bar<-area(depth_bar,width);
#compute average wetted perimeter
perimeter_bar<-perimeter(depth_bar,width);
#compute average hydraulic radius
radius_bar<-radius(area_bar,perimeter_bar);
#compute friction slope
friction<-slope_f(discharge,mannings_n,area_bar,radius_bar)
# compute change in distance for each change in depth
deltax[i]<-((depth[i+1]+(velocity[i+1]^2)/(2*9.8))
- (depth[i] + (velocity[i]^2)/(2*9.8))) / (slope-friction);
}
# space for computing cumulative distances
distance<-numeric(0);
distance[1]<-0;
bse[1]<-0; # bottom elevation at origin
for (i in 2:(how_many)){
distance[i]<-distance[i-1]+deltax[i-1]; # spatial distances
bse[i]<-bse[i-1]-deltax[i-1]*slope; # bottom elevations
}
wse<-bse+depth # water surface elevations
# output
z<-cbind(distance,depth,bse,wse) # bind output into 4 columns
return(z)
#
}

```

13.2.3 Example 1 — Backwater curve

Figure 98 is a backwater curve⁵⁰ for a rectangular channel with discharge over a weir (on the right hand side — not depicted). The channel width is 5 meters, bottom slope 0.001, Manning's $n = 0.02$ and discharge $Q = 55.4 \frac{m^3}{sec}$.

Using the `backwater` function and some plot calls in **R** we can duplicate the figure (assuming the figure is essentially correct).

⁵⁰Page 85. Koutitas, C.G. (1983). Elements of Computational Hydraulics. Pentech Press, London 138p. ISBN 0-7273-0503-4




Figure 98. Example backwater curve.

Listing 45. R Console Output when script is run

```
> # here is the function call
> backwater(begin_depth=8,end_depth=5,how_many=31,
+ discharge=55.4,width=5,mannings_n=0.02,slope=0.001)
      distance  depth      bse      wse
[1,]  0.0000  8.000000  0.0000000  8.0000000
[2,] -283.3504  7.806452  0.2833504  8.089802
[3,] -428.0112  7.709677  0.4280112  8.137689
[4,] -574.8516  7.612903  0.5748516  8.187755
[5,] -724.0433  7.516129  0.7240433  8.240172
[6,] -875.7785  7.419355  0.8757785  8.295133
... Many Rows ...
[29,] -7186.0943  5.193548  7.1860943  12.379643
[30,] -8389.5396  5.096774  8.3895396  13.486314
[31,] -11393.2010  5.000000  11.3932010  16.393201
```

Figure 99 is the same situation computed and plotted using the script in this essay⁵¹.

⁵¹With some additions that will be demonstrated in class!




Figure 99. Backwater curve computed and plotted using **R**.

13.2.4 Example 2 — Front-water curve

Figure 100 is another illustrative case. Here the water discharges into a horizontal channel at a rate of $1 \frac{m^3}{sec}$ per meter width. Assuming Manning's $n \approx 0.01$ we wish to compute the profile downstream of the gate and determine if it will extend to the sharp edge⁵².

We would need to know the critical depth for the section ($\approx 0.47meters$), then compute the profile moving from the gate downstream (a frontwater curve with respect to the gate).

With the `backwater` function, all we really need to do is change the function arguments because it is already built for rectangular channels.

Observe that the distance is now incrementing forward (by choice of begin and end depths). Figure 101 is the situation computed and plotted using the script.

⁵²Obviously the profile will change a lot near the edge, but the question is will the profile continue to rise as depicted if the edge were further away?




Figure A

Figure 100. example caption.

Listing 46. R Console Output when script is run

```
> # The function --- note the changes in parameter values!
> backwater(begin_depth=0.1,end_depth=0.47,how_many=20,
+ discharge=1,width=1,mannings_n=0.01,slope=0.000)
  distance depth bse wse
[1,] 0.00000 0.1000 0 0.1000
... Rows ...
[13,] 79.00486 0.3405 0 0.3405
[14,] 82.82243 0.3590 0 0.3590
... Rows ...
```




Figure 101. Frontwater curve computed and plotted using R.

13.2.5 Finite Difference Methods — Fixed Space Change Method

The fixed-depth change, variable-space result is a useful tool, but not terribly practical because we mostly perform engineering hydraulics calculations to estimate values (depth, pressure, force) at prescribed locations in space, so we need another approach to the problem where we can prescribe the spatial locations, and solve for the depths.

First the gradually varied flow equation is rearranged for relating the change in specific energy between two section to the spatial difference and the slope differences as

$$\Delta h_s = \Delta \left(h + \frac{V^2}{2g} \right) = \Delta x (S_o - S_f) \quad (159)$$

The computation of h_{i+1}, V_{i+1} from h_i, V_i is performed by iteration. An initial value for h_{i+1} that is known to be too large is used in Equation 159 along with the known value of h_i to compute a trial value $h_{s(i+1)}$.

Then the trial value is used in the right hand side of Equation 160

$$h_s = \left(h + \frac{V^2}{2g} \right) \quad (160)$$

The two trial values are compared and the next value of h_{i+1} is computed by successively decreasing until the two values computed by the difference equation and the definition of specific energy coincide. The example below uses a method from Hamming to make the comparisons and adjust the guesses until they are sufficiently close enough.

Example (Non-Prismatic Channel), Fixed Spatial Steps)

A plan view of a rectangular channel of variable width as shown in Figure 102.




Figure 102. Non-Prismatic Rectangular Channel.

The channel conveys $Q = 100 \text{ m}^3/\text{sec}$, with a bottom slope of 0.001 and average Manning's n value of 0.033. A backwater curve is caused by a weir at the downstream end (to the right in the figure) by a 7 meter tall weir. Flow depth over the weir is at critical depth $h_c = 2.17$ meters. Normal flow in the upstream portion for 10-meter channel width is $h_n = 5.6$ meters. Using the fixed space step method determine and plot a profile view of the water surface and channel bottom.

Figure 103 is a screen capture of the result for the example problem.

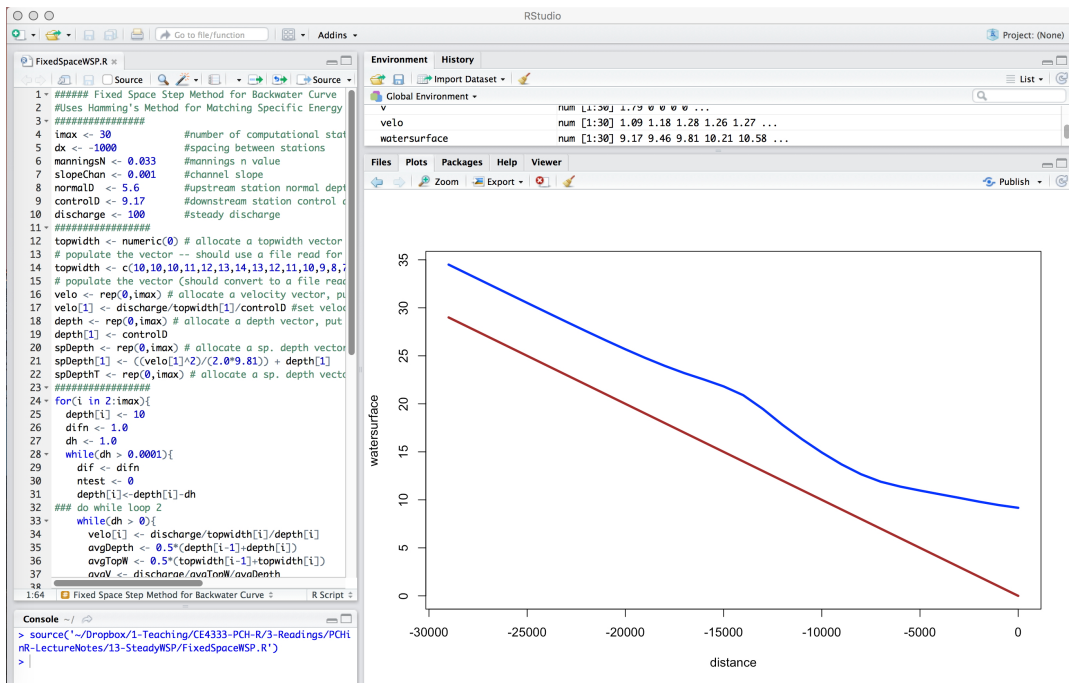


Figure 103. Backwater curve computed and plotted using R.

Listing 47 is an entire listing for the example problem. The listing introduces our first intricate use of `for(...)`, `while(...)`, and the use of `break` to control the program logic. The outer `for(...)` loop is used to step through the various cross sections while the two inner `while(...)` “loops” are used to perform the specific energy balancing logic and exit under different conditions.

In older languages where `goto ... <label>` was allowed logical control crossed over loopback paths. Computer scientists claimed that a `goto ... <label>` was unnecessary – this example here was ported from a FORTRAN code from the 1980's, and indeed the `goto ... <label>` statements were not required.

Listing 47. R script for Backwater Curve by Fixed Space Method

```
##### Fixed Space Step Method for Backwater Curve #####
#Uses Hamming's Method for Matching Specific Energy at Sections
#####
imax <- 30          #number of computational stations
dx <- -1000         #spacing between stations
manningsN <- 0.033  #mannings n value
slopeChan <- 0.001  #channel slope
normalD <- 5.6      #upstream station normal depth
controlD <- 9.17    #downstream station control depth
discharge <- 100    #steady discharge
#####
topwidth <- numeric(0) # allocate a topwidth vector
# populate the vector -- should use a file read for general program
topwidth <- c
  (10,10,10,10,11,12,13,14,13,12,11,10,9,8,7,8,9,10,10,10,10,10,10,10,10,10,10,10,10,10,10,10)
# populate the vector (should convert to a file read)
velo <- rep(0,imax) # allocate a velocity vector, put zeroes everywhere
velo[1] <- discharge/topwidth[1]/controlD #set velocity at control section
depth <- rep(0,imax) # allocate a depth vector, put zeroes everywhere
depth[1] <- controlD
spDepth <- rep(0,imax) # allocate a sp. depth vector, put zeroes everywhere
spDepth[1] <- ((velo[1]^2)/(2.0*9.81)) + depth[1]
spDepthT <- rep(0,imax) # allocate a sp. depth vector, put zeroes everywhere
#####
for(i in 2:imax){
  depth[i] <- 10
  difn <- 1.0
  dh <- 1.0
  while(dh > 0.0001){
    dif <- difn
    ntest <- 0
    depth[i] <- depth[i] - dh
  ## do while loop 2
    while(dh > 0){
      velo[i] <- discharge/topwidth[i]/depth[i]
      avgDepth <- 0.5*(depth[i-1]+depth[i])
      avgTopW <- 0.5*(topwidth[i-1]+topwidth[i])
      avgV <- discharge/avgTopW/avgDepth
      hydR <- avgDepth*avgTopW/(avgTopW+2.0*avgDepth)
      sFric <- (avgV^2)*(manningsN^2)/(hydR^(1.33))
      spDepth[i] <- spDepth[i-1]+(slopeChan-sFric)*dx
      spDepthT[i] <- depth[i]+(velo[i]^2/(2.0*9.81))
      difn <- spDepthT[i]-spDepth[i]
      #
      print(difn)
      #
      print(cbind(i,depth[i],spDepth[i],spDepthT[i],dh))
      if(ntest > 0){
        dh <- dh/10.0
        break #break from do while loop 2
      }
      if(dif*difn > 0){
        break #break from do while loop 2
      }
    }
    depth[i] <- depth[i] + dh
    ntest <- 1
  } #do while loop 2
} # do while loop 1
} # for loop
##### report results #####
# build x-vector
distance <- seq(0,(imax-1)*dx,dx)
bottom <- -distance*slopeChan
watersurface <- depth+bottom
plot(distance,watersurface,ylim=c(0,max(watersurface)),type="l",col="blue",lwd=3)
lines(distance,bottom,col="brown",lwd=3)
```

13.3 Exercise Set 10

1. Modify the (variable space step method) **R** code for U.S. Customary units. Use the modified function to compute the water surface profile for a wide rectangular channel with Manning's $n = 0.022$, bottom slope $S_0 = 0.0048$, and discharge per unit width of $\frac{Q}{W} = 5 \frac{ft^3}{sec}$. Determine how far along the channel $x = L$ does it take for the flow depth to rise from a value of $y = 3.0ft$ to $y = 4.0ft$. Is the 4-ft depth position upstream or downstream of the 3-ft depth position? Plot

the water surface profile.⁵³.

2. Modify the (variable space step method) `area` and `perimeter` functions as well as the `backwater` functions for a trapezoidal channel. Apply the modified functions for a trapezoidal channel with a [2 : 1] side slope, 3.5ft bottom width, 0.012 bed slope, that discharges from a reservoir at $Q = 185 \frac{ft^3}{sec}$. You should assume the upstream value is at critical depth ($\approx 2.7ft$) and compute the profile to within 2% of normal depth. Plot the profile⁵⁴.
3. Convert the fixed space step example problem into US customary units. Then modify the script for US customary units and repeat the example.

These exercises are also located on the class server as ES-10

⁵³Guideline: The profile should extend less than 400 feet for this problem — if your profiles are going further, there is probably something wrong with your functions or input values.

⁵⁴Guideline: The profile should extend less than 100 feet for this problem — if your profiles are going further, there is probably something wrong with your functions or input values.

14 Finite-Difference Method for Unsteady Open Channel Flow

Solving the St. Venant Equations is accomplished by mapping the physical system and the set of partial differential equations into an algebraic structure that a computer can manipulate. Finite-difference, finite-element, finite-volume, and marker-in-cell are the typical methods. The simplest form of solution that is conditionally stable and reasonably straightforward to program is called the Lax-Diffusion scheme.

This scheme is reasonably accurate and useful for practical problems as well as to learn what goes on under the hood of a professional tool like HEC-RAS.

14.1 Governing Difference Equations – Lax Scheme

The finite-difference analysis converts the two PDEs into an algebraic update structure and maps boundary conditions onto a computational domain. The two PDEs are continuity and momentum; they are coupled in the same sense as they were in the water-hammer part of pipeline flow – the resulting difference scheme will look quite similar.

14.1.1 Continuity

The continuity equation for a computational cell (reach) is

$$\frac{\partial y}{\partial t} = -\frac{A}{B} \frac{\partial V}{\partial x} - V \frac{\partial y}{\partial x} \quad (161)$$

A is the depth-area function (a function of x and y). B is the depth-topwidth function (a function of x and y). The Lax scheme uses spatial averaging to represent the A , B , and V terms that appear as coefficients on the partial derivatives on the right-hand side of the equation. The time derivative is accomplished with a conventional forward-in-time first order

finite difference model, and the spatial derivatives are conventional

first-order centered differences.

14.1.2 Momentum

The momentum equation is

$$\frac{\partial V}{\partial t} = g(S_0 - S_f) - V \frac{\partial V}{\partial x} - g \frac{\partial y}{\partial x} \quad (162)$$

and also uses spatial averages for the coefficients on the spatial derivatives in the right hand side of the equation as well as spatial averages for the friction and topographic slopes. Friction slope can be recovered using any resistance model, Chezy-Manning's is typical.

14.2 Mapping from the Physical to Computational Domain

The next important step is to map the physical world to the computer world.




Figure 104. Plan view of a stream. Flow in figure is from left to right.

Figure 104 is a schematic of a stream that is to be modeled. The stream has some width, depth, and path. The dashed line in the figure is the thalweg and is the pathline of the stream. Distances in the computational model are along this path. The conventional orientation is “looking downstream.” So when the cross sections are stationed the distances in a cross section are usually referenced as distanced from the left bank, looking downstream.

Figure 105 is a schematic that depicts the relationship of left-bank, cross section, elevations, and such – all referenced to the concept of “looking downstream.”




Figure 105. Schematic of relationship of cross-section, elevation, and left bank.

Figure 106 is a schematic of the next step of mapping into the computational domain.




Figure 106. Schematic of physical interpretation of a reach, cell, and node.

In the figure the stream is divided into cells called reaches (or cells, depends on context and author). The centroid of the reach is called the node, and most of the arithmetic is written with the understanding that all properties of the reach are somehow “averaged” and these averages are assigned to these nodes. Adjacent nodes are connected (in the computer) by links. The continuity and momentum equations collectively describe the node average behavior (such as depth) and link behavior (such as momentum flux).

Figure 107 is a schematic of three adjacent nodes that is used to develop the difference equations.

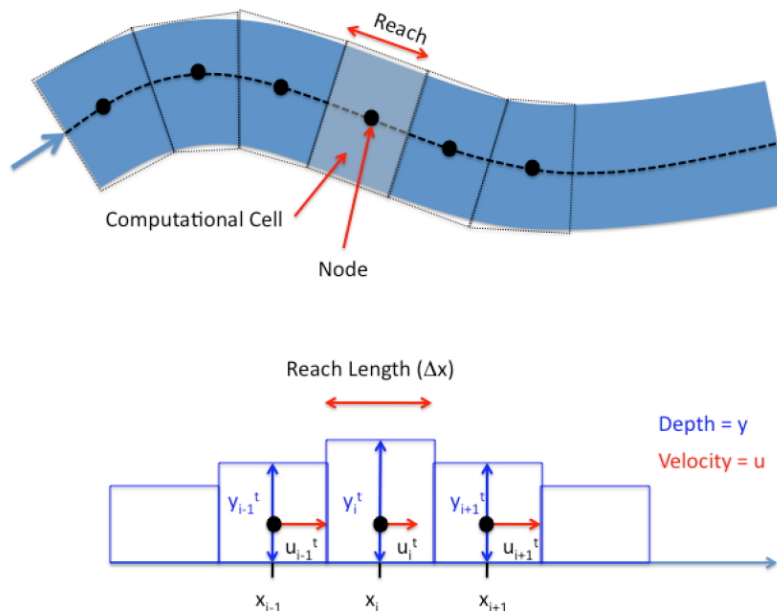


Figure 107. Schematic of three adjacent nodes, with average depth and section velocity depicted at the node. Different schemes map these quantities to different places in the cell – often velocities are at cell interfaces.

In the figure both the velocities and depths are mapped to the node (Lax-Diffusion scheme), but other schemes map the velocities to the interfaces. Again this decision affects the differencing scheme; the differencing scheme chooses the location. A kind of chicken and egg situation.

At this point the mapping has abstracted considerably from the physical world and the computer world loses sense of sinuosity. In this development, we will assume the reach lengths are all the same value, the velocities are all parallel to the local thalweg and perpendicular to the cross sections, and the depth is measured from the channel bottom. The differencing scheme then replaces the continuity and momentum PDEs with update equations to map the water surface position and mean section velocity at the nodes to different moments in time. The updating is called time-stepping.

14.3 Building the Difference Equations

The partial derivatives are replaced with difference quotients that approximate their behavior. The mapping in some sense influences the resulting difference scheme.

14.3.1 Time Differences

The naive time difference is

$$\frac{\partial y}{\partial t} \approx \frac{y_i^{t+\Delta t} - y_i^t}{\Delta t} \quad (163)$$

Lax replaced the known time-level term with its spatial average from adjacent cells. For the depth;

$$\frac{\partial y}{\partial t} \approx \frac{y_i^{t+\Delta t} - \frac{1}{2}(y_{i-1}^t + y_{i+1}^t)}{\Delta t} \quad (164)$$

Similarly for mean section velocity

$$\frac{\partial V}{\partial t} \approx \frac{V_i^{t+\Delta t} - \frac{1}{2}(V_{i-1}^t + V_{i+1}^t)}{\Delta t} \quad (165)$$

14.3.2 Space Differences

Lax used centered differences for the spatial derivatives

$$\frac{\partial y}{\partial x} \approx \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (166)$$

$$\frac{\partial V}{\partial x} \approx \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} \quad (167)$$

Lax also used spatial averages for the depth-area and slope functions

$$\frac{1}{2} \left(\frac{A}{B} \Big|_{i-1}^t + \frac{A}{B} \Big|_{i+1}^t \right) \quad (168)$$

$$\frac{1}{2} (S_{f,i-1}^t + S_{f,i+1}^t) \quad (169)$$

These difference formulations are substituted into continuity and momentum and then rearranged to isolate the terms at the $t + \Delta t$ time level.

14.3.3 Continuity

Starting with the PDE,

$$\frac{\partial y}{\partial t} = -\frac{A}{B} \frac{\partial V}{\partial x} - V \frac{\partial y}{\partial x} \quad (170)$$

first replace the time derivative

$$\frac{y_i^{t+\Delta t} - \frac{1}{2}(y_{i-1}^t + y_{i+1}^t)}{\Delta t} = -\frac{A}{B} \frac{\partial V}{\partial x} - V \frac{\partial y}{\partial x} \quad (171)$$

then replace the space derivatives

$$\frac{y_i^{t+\Delta t} - \frac{1}{2}(y_{i-1}^t + y_{i+1}^t)}{\Delta t} = -\frac{A}{B} \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} - V \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (172)$$

then the spatial averages for the remaining terms.

$$\frac{y_i^{t+\Delta t} - \frac{1}{2}(y_{i-1}^t + y_{i+1}^t)}{\Delta t} = -\frac{1}{2} \left(\frac{A}{B} \Big|_{i-1}^t + \frac{A}{B} \Big|_{i+1}^t \right) \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} - \frac{1}{2} (V_{f,i-1}^t + V_{f,i+1}^t) \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (173)$$

Next multiply by Δt

$$y_i^{t+\Delta t} - \frac{1}{2}(y_{i-1}^t + y_{i+1}^t) = -\Delta t \frac{1}{2} \left(\frac{A}{B} \Big|_{i-1}^t + \frac{A}{B} \Big|_{i+1}^t \right) \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} - \Delta t \frac{1}{2} (V_{f,i-1}^t + V_{f,i+1}^t) \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (174)$$

Move the time level t term to the right hand side

$$y_i^{t+\Delta t} = \frac{1}{2}(y_{i-1}^t + y_{i+1}^t) - \Delta t \frac{1}{2} \left(\frac{A}{B} \Big|_{i-1}^t + \frac{A}{B} \Big|_{i+1}^t \right) \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} - \Delta t \frac{1}{2} (V_{f,i-1}^t + V_{f,i+1}^t) \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (175)$$

Rename the constant $\frac{\Delta t}{2\Delta x} = r$ and simplify

$$y_i^{t+\Delta t} = \frac{1}{2}(y_{i-1}^t + y_{i+1}^t) - \frac{r}{2} \left(\frac{A}{B} \Big|_{i-1}^t + \frac{A}{B} \Big|_{i+1}^t \right) (V_{i+1}^t - V_{i-1}^t) - \frac{r}{2} (V_{f,i-1}^t + V_{f,i+1}^t) (y_{i+1}^t - y_{i-1}^t) \quad (176)$$

14.3.4 Momentum

Again, starting with the PDE, make the time substitution

$$\frac{V_i^{t+\Delta t} - \frac{1}{2}(V_{i-1}^t + V_{i+1}^t)}{\Delta t} = g(S_0 - S_f) - V \frac{\partial V}{\partial x} - g \frac{\partial y}{\partial x} \quad (177)$$

next the space derivatives

$$\frac{V_i^{t+\Delta t} - \frac{1}{2}(V_{i-1}^t + V_{i+1}^t)}{\Delta t} = g(S_0 - S_f) - V \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} - g \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (178)$$

then the spatial averages⁵⁵

$$\frac{V_i^{t+\Delta t} - \frac{1}{2}(V_{i-1}^t + V_{i+1}^t)}{\Delta t} = g(S_0 - \frac{1}{2}(S_{f,i-1}^t + S_{f,i+1}^t)) - \frac{1}{2}(V_{i-1}^t + V_{i+1}^t) \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} - g \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (179)$$

Multiply by Δt

$$V_i^{t+\Delta t} - \frac{1}{2}(V_{i-1}^t + V_{i+1}^t) = \Delta t g(S_0 - \frac{1}{2}(S_{f,i-1}^t + S_{f,i+1}^t)) - \Delta t \frac{1}{2}(V_{i-1}^t + V_{i+1}^t) \frac{V_{i+1}^t - V_{i-1}^t}{2\Delta x} - \Delta t g \frac{y_{i+1}^t - y_{i-1}^t}{2\Delta x} \quad (180)$$

Rename the constant $\frac{\Delta t}{2\Delta x} = r$ and isolate the $t + \Delta t$ term

$$V_i^{t+\Delta t} = \frac{1}{2}(V_{i-1}^t + V_{i+1}^t) + \Delta t g(S_0 - \frac{1}{2}(S_{f,i-1}^t + S_{f,i+1}^t)) - \frac{r}{2}(V_{i-1}^t + V_{i+1}^t)(V_{i+1}^t - V_{i-1}^t) - rg(y_{i+1}^t - y_{i-1}^t) \quad (181)$$

The pair of update equations, Equation 176 and 181 are the interior point update equations.

Figure 108 depicts the updating information transfer. At each cell the three known values of a variable (y or V) are projected to the next time line as depicted in the figure. Boundary conditions are the next challenge. These are usually handled using characteristic equations (unless the boundaries are really simple).

14.3.5 Example 1 – Steady flow in a rectangular channel

The backwater curve situation for a rectangular channel with discharge over a weir is repeated. The channel width is 5 meters, bottom slope 0.001, Manning's $n = 0.02$ and discharge $Q = 55.4 \frac{m^3}{sec}$. We will start with the flow depth artificially large and observe that the transient solver will eventually produce an equilibrium solution that is the same as the steady-flow solver. Generally such a simulation is a good idea to test a new algorithm – it should be stable enough to convereg to and maintain a steady solution.

⁵⁵If the channel slope is changing, then this would be subjected to a spatial averaging scheme too!




Figure 108. Relation of the linked reaches to solution of the equations in the XT-plane. Explicit updating (as used herein) uses the three values at the known time level to project (update) the unknown value at the next time level. Boundary behavior is an independent calculation, dependent on the evolution of the interior solution.

A script that implements these concepts is listed on the next several pages. The script is comprised of several parts, and for the sake of taking advantage of the ability of **R** to read and operate on files, the script will have several “libraries” that read by a main control program. The main program controls the overall solution process, while the library functions can be built and tested in advance.

First are prototype functions for the hydraulic components associated with the channel geometry, shown on listing 48.

Listing 48. R code demonstrating prototype hydraulic functions.

```
# hydraulic functions -- save in file named <lax-diffusion-hydraulics-lib.R>
# depth == flow depth
# bottom == bottom width of trapezoidal channel
# side == side slope (same value both sides) of trapezoidal channel
# computed values:
# bt == computed topwidth :: ar == flow area, used in fd update :: wp == wetted perimeter,
# used in fd update
bt <- function(depth,bottom,side){ # depth-topwidth function
  topwidth <- (bottom + 2.0*side*depth);
  return(topwidth);
} # tested 12MAR2015 TGC
ar <- function(depth,bottom,side){ # depth area function
  area <- (depth*(bottom+side*depth));
  return(area)
} # tested 12MAR2015 TGC
# depth perimeter
wp <- function(depth,bottom,side){
  perimeter <- (bottom+2.0*depth*sqrt(1.0+side*side));
  return(perimeter)
} # tested 12MAR2015 TGC
```

The next set of functions are prototype functions for reporting the output – it will be cleaner to build the output functions separate from the control program, and send

the necessary vectors when we want to actually print results. Listing 49 is a listing of such functions (albeit specific to the particular code).

Listing 49. R code demonstrating prototype display (printing) functions.

```
# display functions -- save in a file named <lax-diffusion-display-lib.R>
##### Message Functions #####
writenow <- function(t,dt,y,v,b0,s) { # printing functions
message("-----")
message("Time = ",t," seconds.", "Time step length = ",dt," seconds ")
aa <- ar(y,b0,s);
qq <- aa*v
brb <- bt(y,b0,s)
ww <- wp(y,b0,s)
message("--in write now----")
print(cbind(y,qq,v,qq,brb,ww,zz))
message("-----")
return() #observe a NULL return, this function messages to the output device, so there is
nothing to return.
}
##### Plotting Functions #####
plotnow<-function(t,x,y,v){
  mainlabel=c("Flow Depth at time = ",t," seconds")
  plot(x,y,main=mainlabel,xlab="distance (m)",ylab="depth (m)",xlim=c(0,30000),ylim=c
(0,15),type="l",lwd=3,col="blue")
  lines(x,v,lwd=3,col="red",type="l")
return() #observe a NULL return, this function plots to the graphics device, so there is
nothing to return.
}
```

Listing 50. R code demonstrating prototype finite-difference formulation/function.

```
##### Finite Difference Functions #####
finitedifference<-function(){
r <- 0.5*dt/dx;
##### LEFT BOUNDARY #####
## MODIFY TO A REFLECTION BOUNDARY ##
## THEN SUPPLY DISCHARGE TO THE FIRST CELL AND USE CONTINUITY TO FIND DEPTH ##
## qq is the interpolated hydrograph flow value
#debug message("left boundary ",qq$y)
#debug message("left boundary ",y[1])
#####
vp[1] <- qq$y/ar(y[1],b0,s);
ab <- ar(y[2],b0,s);
bb <- bt(y[2],b0,s);
cb <- sqrt(g*bb/ab);
rb <- ab/wp(y[2],b0,s);
sfb <- (mn2*v[2]*v[2])/(rb^(1.333));
#debug print(cbind(ab,bb,cb,rb,sfb));
cn <- v[2] -cb*y[2]+ g*(s0-sfb)*dt;
yp[1] <- (vp[1] - cn )/cb;
#####
## gives y,v at location 1 time level k+1
#debug message("--left bndry ----")
#debug print(cbind(y,yp,v,vp,sfb,cn,rb));
##### RIGHT BOUNDARY #####
## MODIFY TO A FREE OUTFALL
## USE A WEIR EQUATION BASED ON SOME REFERENCE DEPTH
vp[nn]<- (y[n]-yd)*sqrt(9.8*y[n]);
## USE CHARACTERISTIC LINE TO FIND ASSOCIATED DEPTH
aa <- ar(y[n],b0,s);
ba <- bt(y[n],b0,s);
ca <- sqrt(g*ba/aa);
ra <- aa/wp(y[n],b0,s);
sfa <- (mn2*v[n]*v[n])/(ra^(4.0/3.0));
cp <- v[n] + ca*y[n]+g*(s0-sfa)*dt;
yp[nn] <- (cp - vp[nn])/ca;
# ## fixed stage
# yp[nn] <- yd ;
# aa <- ar(y[n],b0,s);
# ba <- bt(y[n],b0,s);
# ca <- sqrt(g*ba/aa);
# ra <- aa/wp(y[n],b0,s);
# sfa <- (mn2*v[n]*v[n])/(ra^(4.0/3.0));
# cp <- v[n] + ca*y[n]+g*(s0-sfa)*dt;
# vp[nn] <- (cp - ca*yp[nn]); # check sign
# message("--right bndry ----")
# print(cbind(y,yp,v,vp));

# reflection boundary, find depth along a characteristic
# vp[nn] <-0 ;
# aa <- ar(y[n],b0,s);
# ba <- bt(y[n],b0,s);
# ca <- sqrt(g*ba/aa);
# ra <- aa/wp(y[n],b0,s);
# sfa <- (mn2*v[n]*v[n])/(ra^(4.0/3.0));
# cp <- v[n] + ca*y[n]+g*(s0-sfa)*dt;
```

```

# yp[nn] <- (cp - vp[nn])/ca;
##print(cbind(y,yp,v,vp));
##### INTERIOR NODES AND REACHES #####
# loop through the interior nodes
for (i in 2:n){ # begin interior node loop scope
aa <- ar(y[i-1],b0,s);
ba <- bt(y[i-1],b0,s);
pa <- wp(y[i-1],b0,s);
ra <- aa/pa;
sfa <- (mn2*v[i-1]*v[i-1])/(ra^(4.0/3.0));
ab <- ar(y[i+1],b0,s);
bb <- bt(y[i+1],b0,s);
pb <- wp(y[i+1],b0,s);
rb <- ab/pb;
sfb <- (mn2*v[i+1]*v[i+1])/(rb^(4.0/3.0));
# need averages of sf, hydraulic depth
dm <- 0.5*(aa/ba + ab/bb);
sfm <- 0.5*(sfa+sfb);
vm <- 0.5*(v[i-1]+v[i+1]);
ym <- 0.5*(y[i-1]+y[i+1]);
# update momentum
# note the double <<, this structure forces the
# value to be global and accessible
# to other functions when the script is run
vp[i] <- vm -r*g*(y[i+1] - y[i-1]) -r*vm*(v[i+1] - v[i-1]) + g*dt*(s0-sfm);
# update depth
yp[i] <- ym - r*dm*(v[i+1] - v[i-1]) -r*vm*(y[i+1] - y[i-1]);
} # end of interior node loop scope
} # end of function scope

```

Listing 51. R code demonstrating prototype updating function.

```

##### Solution Update Function #####
update<-function(y,yp,v,vp){
y <- yp;
v <- vp;
return()
}
### NEED ADAPTIVE TIME STEPPING FOR THE FLOOD WAVE
##### Adaptive Time Step Functions #####
## here is where we do adaptive time stepping
bestdt<-function(y,v){
bestdt <- dt # start with current time step
for (i in 1:nn){
a <- ar(y[i],b0,s);
b <- bt(y[i],b0,s);
c <- sqrt(g*a/b);
dtn <- dx/abs((v[i])+c)
# now test
if(dtn <= bestdt){bestdt <- dtn}
} # end loop scope
dt <- bestdt
} #end bestdt function
##### Solution Update Function #####
update<-function(y,yp,v,vp){
y <- yp;
v <- vp;
return()
}

```

[Present the problem, and important features – then to solution listed below. Uses the library]

```

# main program for st. venant
rm(list=ls())
setwd("~/Dropbox/CE5361-2015-1/Project2/Project2-Problem2")
# clear workspace and set directory
source('~/Dropbox/CE5361-2015-1/Project2/Project2-Problem2/Project2.2.Lib.R')

##### Problem Constants #####
# these are constants that define the problem
# change for different problems
# a good habit is to assign constants to names so the
# program is readable by people in a few years

```

```

g <- 9.81 # gravitational acceleration, obviously SI units
n <- 29 # number of reaches
#q0 <- 110 # initial discharge
q0 <- 11 # initial discharge
#yd <- 3.069 # initial flow depth in the model
yd <- 1.000 # initial flow depth in the model
#yu <- 3.069 # upstream constant depth
yu <- 1.1 # upstream constant depth
# mn <- 0.013 # Manning's n
mn <- 0.02 # Manning's n
# b0 <- 20 # bottom width
b0 <- 5 # bottom width
## modify for problem 1 s0 <- 0.0001 # longitudinal slope (along direction of flow)
s0 <- 0.0000001 # longitudinal slope (along direction of flow)
#s <- 2.0 # side slope (passed to calls to hydraulic variables)
s <- 0.0 # side slope (passed to calls to hydraulic variables)
l <- 30000.0 # total length (the length of computational domain)
tmax <- 640000 # total simulation time in seconds
iprt <- 8 # print every iprt time steps
nn <- n+1 # how many nodes, will jack with boundaries later
mn2 <- mn*mn # Manning's n squared, will appear a lot.
a <- ar(yd,b0,s) # flow area at beginning of time
v0 <- q0/a # initial velocity
##### Here we build vectors #####
y <- numeric(nn) # create nn elements of vector y
yp <- numeric(nn) # updates go in this vector, same length as y
v <- numeric(nn) # create nn elements of vector v
vp <- numeric(nn) # updates go in this vector, same length and v
ytmp <- numeric(nn)
vtmp <- numeric(nn)
y <- rep(yd,nn) # populate y with nn things, each thing has value yd
v <- rep(v0,nn) # populate v with nn things, each thing has value v0
b <- bt(yd,b0,s) # topwidth at beginning
c <- sqrt(g*a/b) # celerity at initial conditions
###
hydrograph <- numeric(1)
dummy<- read.csv(file="hydrograph.csv",header=F)
elapsedtime <- dummy$V1; # first column of hydrograph is time
hydrograph <- dummy$V2; # second column of hydrograph is flow
print(cbind(c))
dx <- l/n # delta x, length of a reach
xx <- dx*seq(1:nn)-1000 # Spatial locations of nodes, used for plotting
zz <- 30 - s0*xx # bottom channel elevation
dt <- dx/(v0 + c) # the time step that satisfies the courant conditions
kmax <- round(tmax/dt) # set maximum number of time steps

```

```
print(cbind(dx,dt))
```

```
### Run the simulation ###
k <- 0 # time counter
t <- 0.0 # elapsed time
pdf("junk2.2.plot.pdf") # graphics device for plots -- plotnow() sends data to plot
writenow(t,dt,y,v,b0,s) # Write initial conditions
plotnow(t,xx,y,v)
##### BEGIN TIME STEPPING #####
message("kmax =", kmax)
for (itime in 1:kmax){
## put in the hydrograph here -- use approx function to interpolate
## we are after the second value qq$y in the approx function
qq <-< approx(elapsedtime,hydrograph,t)
print (qq$x);
print (qq$y);
#debug print(qq$y)
##### NEED ADAPTIVE TIME STEPPING FOR THE HYDROGRAPH ROUTING #####
bestdt(y,v);
##### THIS IS A HACK TO GET STABILITY #####
# halve the time step
dt <- dt/16
finitedifference(); # Finite difference a single time step
#update(ytmp,yp,vtmp,vp); # Update vectors
#update(y,yp,v,vp); # Update vectors
#bestdt(yp,vp)
message("dt now = ",dt);
#finitedifference(); # Finite difference a single time step
ytmp <- (yp+ytmp)/2;
vtmp <- (vp+vtmp)/2;
update(y,yp,v,vp); # Update vectors
message(" dt is ",dt);
t <- t+dt; # Increment simulation time
k <- k+1; # Increment loop counter
if (k%%iprt == 0){writenow(t,dt,y,v,b0,s)}; # Write current conditions every iprt tim
if (k%%iprt == 0){plotnow(t,xx[seq(1,nn,1)],(y[seq(1,nn,1)]),(v[seq(1,nn,1)]))}; # Pl
# reset the time step
##### UNHACK TO KEEP ADAPTIVE TIME STEPPING WORKING #####
dt <- 16*dt
}
dev.off() # disconnect the pdf file.
```

14.3.6 Example 2 – Routing a storm hydrograph

The initial depth in a horizontal channel of rectangular cross section is 1 meter. The channel is 29 kilometers long and ends with a non-reflection boundary condition. The initial discharge in the channel is 0 cubic meters per second. The upstream input hydrograph is shown in Figure 109. The Manning friction factor is $n = 1/40$. Simulate the water surface elevation over time in the channel.




Figure 109. Upstream hydrograph for example.

14.3.7 Example 3 – Sudden gate closing in an aqueduct channel

To illustrate a script to implement these concepts consider the example problem: Flow in a 1000-m long trapezoidal channel with a bottom width of 20-m, side slope of 2H:1V, longitudinal slope $S_0=0.0001$, and Manning's resistance $n=0.013$. Initial discharge in the channel is 110 m³/s and initial flow depth is 3.069 m. Simulate the flow and depth at every 100-m station when a downstream gate is closed at $t=0$. Produce a graph of depth and velocity versus location for $t=0, 60, 360$ seconds.

14.3.8 Example 4 – Long waves in a tidal-influenced channel

14.4 Quasi-2D Methods

[Use Cleveland and Botkins report to make explain – then make adaptation]

14.5 Exercises

15 Using Method of Characteristics for Boundary Conditions

The method of characteristics is a technique used to find boundary conditions for the finite-difference model(s). The readings for this chapter (on the server) contain several examples of the method of characteristics to establish boundary values for the model

Starting again with continuity and momentum

$$\frac{\partial y}{\partial t} = -\frac{A}{B} \frac{\partial V}{\partial x} - V \frac{\partial y}{\partial x} \quad (182)$$

$$\frac{\partial V}{\partial t} = g(S_0 - S_f) - V \frac{\partial V}{\partial x} - g \frac{\partial y}{\partial x} \quad (183)$$

multiply the continuity equation by a term λ , then add the result to the momentum equation to obtain

$$\lambda \frac{\partial y}{\partial t} + \lambda \frac{A}{B} \frac{\partial V}{\partial x} + \lambda V \frac{\partial y}{\partial x} + g \frac{\partial y}{\partial x} + \frac{\partial V}{\partial t} + V \frac{\partial V}{\partial x} = g(S_0 - S_f) \quad (184)$$

Define the λ as

$$V + \lambda \frac{A}{B} = \frac{dx}{dt} = V + \frac{g}{\lambda} \quad (185)$$

where $\frac{dx}{dt}$ is from the total derivative

$$\frac{dF}{dt} = \frac{\partial F}{\partial t} + \frac{\partial F}{\partial x} \frac{\partial x}{\partial t} \quad (186)$$

[need sketches to define + and - characteristics then extrapolate back to the boundary.]
 [for Summer 2017 - use the boundary conditions in the code without explain]

16 Flow in Porous Materials

Flow in porous media is a topic that appears in many branches of engineering and science, e.g., ground water hydrology, reservoir engineering, soil science, soil mechanics, and chemical engineering (filtration). The aquifer, which is the porous medium domain of the hydrologist, or the oil reservoir, which is the porous medium of the petroleum engineer are typical examples.

Figure 110 is a sketch of different aquifer classifications. A confined aquifer (pressure aquifer) is one bounded above and below by impermeable formations. In a well penetrating such an aquifer, the water level will rise above the base of the confining formation. Water levels in wells that sample a certain aquifer define an imaginary surface called the piezometric surface.




Figure 110. Aquifer classifications.

An unconfined aquifer (water table aquifer; phreatic aquifer) is one with the water table as its upper boundary. The classifications are important because the equations of motion are different in different kinds of aquifers.

16.1 Storage

Storativity of an aquifer is the relationship between changes in head within the aquifer and the quantity of water stored in the aquifer. Figure 111 is a sketch showing the storage process in a confined, and unconfined aquifer.




Figure 111. Illustrative sketches of definition of storage coefficient.

The mechanism of storage is different for confined and unconfined aquifers. In a confined aquifer the water is stored or released by compression and decompression of water and the solid matrix (like a sponge squeezed while wrapped in plastic wrap). In an unconfined aquifer the water is stored or released from the pore space when the water table elevation changes.

The storage coefficient (confined) or specific yield (unconfined) is the volume of water added to (or removed from) storage per unit area of aquifer per unit change in head. The usual symbols are S , and S_y .

16.2 Permeability

Permeability is the material property that relates the resistance of flow through the porous medium to the hydraulic gradient.

16.3 Head Loss Models

Darcy's law (a linear flux model) is the head loss model used for porous media flow. Equation 187 is Darcy's law expressed as a head loss model.

$$h_L = \frac{QL}{KA} \quad (187)$$

where Q is the discharge in the aquifer, L is the length in the flow direction, A is the cross sectional area of aquifer (pore space and solid phase), K is the hydraulic conductivity.⁵⁶

A more useful (for computation) form of the head loss model, is to express it [the loss equation] as an equation of motion as in Equation 188.

$$Q = -KA \frac{\partial h}{\partial x} \quad (188)$$

where $-\frac{\partial h}{\partial x}$ is the hydraulic gradient (slope of the hydraulic grade line) in the aquifer.

Figure 112 is a diagram that illustrates the relationships expressed by Darcy's law.

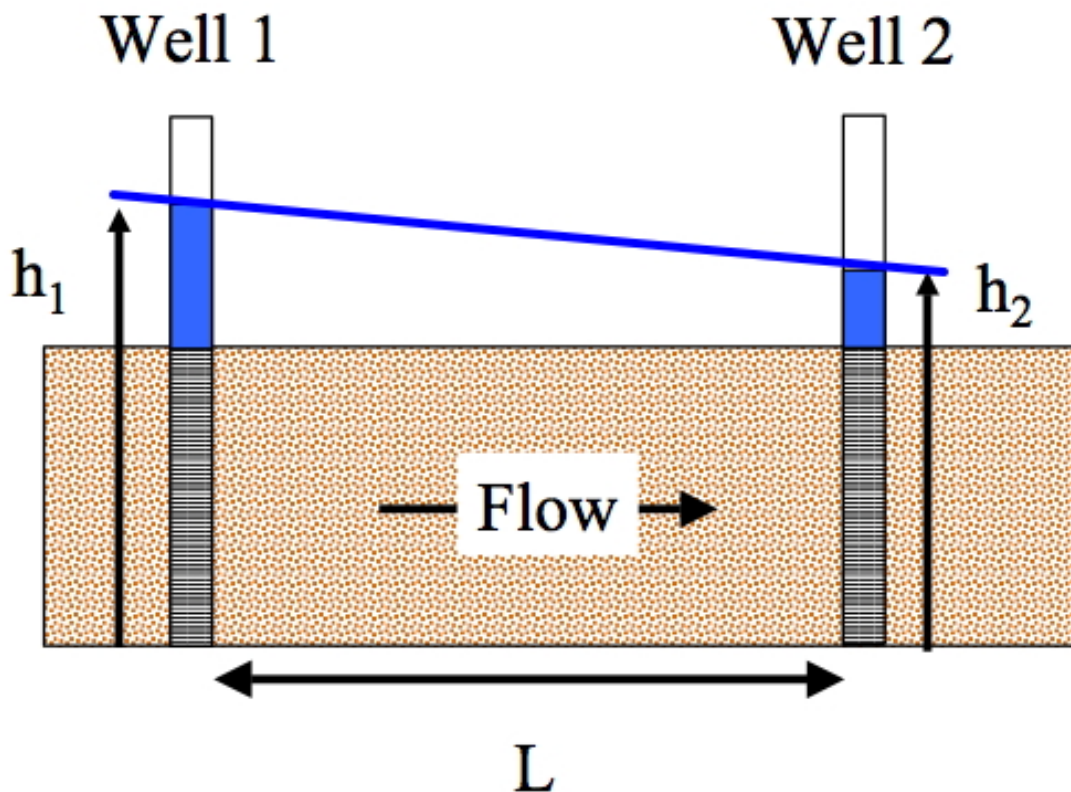


Figure 112. Schematic diagram of unidirectional flow in a generic aquifer, showing heads in two measuring wells located distance L apart..

The cross-sectional flow area, A , is the product of height of the aquifer block and its width [in this case the width is into the plane of the paper]. The distance between two measurement points is L . The head at the two points is h_1 and h_2 . The gradient of head, $\frac{\partial h}{\partial x}$, is $\frac{h_2-h_1}{L}$. The hydraulic gradient is $-\frac{\partial h}{\partial x}$, is $\frac{h_1-h_2}{L}$. Finally Darcy's law (for the drawing) is $Q = KA \frac{h_1-h_2}{L}$.

⁵⁶also called the permeability

16.4 Confined Aquifer Flow

Using Figure 112 as a starting point, we can develop a computational model of flow in a confined aquifer. Let's decide that the distance L in the figure is going to be divided into a series of connected, small blocks. The flow direction in the figure will be declared the x direction, the depth into the drawing is declared the y direction, and the height of the block is declared the z direction.

Figure 113 is a diagram of one such small block.




Figure 113. Single computational cell definition sketch.

Using this diagram we can now develop a set of expressions for the cell volume, solids volume in the cell, pore volume in the cell (where water actually can flow), and solids mass.

$$V_{cell} = \Delta x \times \Delta y \times \Delta z \quad (189)$$

$$V_{solid} = (1 - \omega)\Delta x \times \Delta y \times \Delta z \quad (190)$$

$$V_{pore} = \omega\Delta x \times \Delta y \times \Delta z \quad (191)$$

$$M_{solid} = \rho_s(1 - \omega)\Delta x \times \Delta y \times \Delta z \quad (192)$$

Next write a mass balance for water in the cell;

$$\frac{dM_{water}}{dt} = M_{Inflow} - M_{Outflow} \quad (193)$$

The left side of the expression is simply the storage term, and in the context of storage coefficients and aquifer head is replaced by

$$\left. \frac{dM_{water}}{dt} \right|_{cell} = \rho_w S_s \Delta x \Delta y \Delta z \frac{\partial h_i}{\partial t} \quad (194)$$

where h_i is the head in the i -th cell.

The right hand side of the expression is based on writing Darcy's law for the cell, using values in adjacent (hydraulically connected) cells.




Figure 114. Multiple computational cell definition sketch.

Figure 114 is a sketch showing three such cells. The i -th cell is the cell of interest, the cell to the left is cell ID $i - 1$, and the cell to the right is cell ID $i + 1$.

We now write Darcy's law for each face of cell i , treating the head in each of the cell centers as if they were the sampling wells of Figure 112.⁵⁷

Darcy's law for the left face is

$$M_{Inflow} = Q_{left} = \rho_w K \Delta y \Delta z \frac{h_{i-1} - h_i}{\Delta x} \quad (195)$$

⁵⁷In the context of Figure 112, the cell face is halfway between the two wells; the cell centers are at the wells.

Similarly for the right face,

$$M_{Outflow} = Q_{right} = \rho_w K \Delta y \Delta z \frac{h_i - h_{i+1}}{\Delta x} \quad (196)$$

Now combine these together in the mass balance

$$\rho_w S_s \Delta x \Delta y \Delta z \frac{\partial h_i}{\partial t} = (\rho_w K \Delta y \Delta z \frac{h_{i-1} - h_i}{\Delta x}) - (\rho_w K \Delta y \Delta z \frac{h_i - h_{i+1}}{\Delta x}) \quad (197)$$

Next divide by the water density ρ_w ,

$$S_s \Delta x \Delta y \Delta z \frac{\partial h_i}{\partial t} = (K \Delta y \Delta z \frac{h_{i-1} - h_i}{\Delta x}) - (K \Delta y \Delta z \frac{h_i - h_{i+1}}{\Delta x}) \quad (198)$$

The divide by cell width Δy ,

$$S_s \Delta x \Delta z \frac{\partial h_i}{\partial t} = (K \Delta z \frac{h_{i-1} - h_i}{\Delta x}) - (K \Delta z \frac{h_i - h_{i+1}}{\Delta x}) \quad (199)$$

Rewrite the right hand side into gradient of head form

$$S_s \Delta x \Delta z \frac{\partial h_i}{\partial t} = (K \Delta z \frac{h_{i+1} - h_i}{\Delta x}) - (K \Delta z \frac{h_i - h_{i-1}}{\Delta x}) = K \Delta z \frac{\partial h}{\partial x} \Big|_{i \rightarrow i+1} - K \Delta z \frac{\partial h}{\partial x} \Big|_{i-1 \rightarrow i} \quad (200)$$

Divide by cell distance, Δx ,

$$S_s \Delta z \frac{\partial h_i}{\partial t} = \frac{K \Delta z \frac{\partial h}{\partial x} \Big|_{i \rightarrow i+1} - K \Delta z \frac{\partial h}{\partial x} \Big|_{i-1 \rightarrow i}}{\Delta x} \quad (201)$$

Take limit as $\Delta x \rightarrow 0$,

$$S_s \Delta z \frac{\partial h}{\partial t} = \frac{\partial}{\partial x} (K \Delta z \frac{\partial h}{\partial x}) \quad (202)$$

Finally, stipulate that $S_s \Delta z = S$, the storage coefficient (for confined aquifer), and define the aquifer transmissivity as $K \Delta z = T$ and we have performed a back-handed way to get the partial differential equation of aquifer flow.

$$S \frac{\partial h}{\partial t} = \frac{\partial}{\partial x} (T \frac{\partial h}{\partial x}) \quad (203)$$

Ironically, the analysis actually provides an algorithm to approximate head in the aquifer at Equation 203 – which is the subject of the next two chapters.

16.5 Unconfined Aquifer Flow

Unconfined aquifer flow has an upper boundary of flow defined by the phreatic surface (water table). If we assume the flow lines are parallel to the aquifer bottom (they are not, but the assumption is often quite adequate) then the governing PDE can be modified with the Δz term replaced by the saturated thickness. Writing Darcy's law for each face of cell i in Figure 114, treating the head in each of the cell centers as if they were the sampling wells of Figure 112.⁵⁸ The primary difference is that the average head between two adjacent cells, \bar{h} , is used as the thickness term rather than Δz .

Darcy's law for the left face is

$$M_{Inflow} = Q_{left} = \rho_w K \Delta y \bar{h} \frac{h_{i-1} - h_i}{\Delta x} \quad (204)$$

Similarly for the right face,

$$M_{Outflow} = Q_{right} = \rho_w K \Delta y \bar{h} \frac{h_i - h_{i+1}}{\Delta x} \quad (205)$$

Now combine these together in the mass balance, replacing the term $S_s \Delta z$ with S_y the specific yield of the unconfined aquifer we obtain

$$\rho_w S_y \Delta x \Delta y \frac{\partial h_i}{\partial t} = (\rho_w K \Delta y \bar{h} \frac{h_{i-1} - h_i}{\Delta x}) - (\rho_w K \Delta y \bar{h} \frac{h_i - h_{i+1}}{\Delta x}) \quad (206)$$

Next divide by the water density ρ_w ,

$$S_y \Delta x \Delta y \frac{\partial h_i}{\partial t} = (K \Delta y \bar{h} \frac{h_{i-1} - h_i}{\Delta x}) - (K \Delta y \bar{h} \frac{h_i - h_{i+1}}{\Delta x}) \quad (207)$$

The divide by cell width Δy ,

$$S_y \Delta x \frac{\partial h_i}{\partial t} = (K \bar{h} \frac{h_{i-1} - h_i}{\Delta x}) - (K \bar{h} \frac{h_i - h_{i+1}}{\Delta x}) \quad (208)$$

Rewrite the right hand side into gradient of head form

$$S_y \Delta x \frac{\partial h_i}{\partial t} = (K \bar{h} \frac{h_{i+1} - h_i}{\Delta x}) - (K \bar{h} \frac{h_i - h_{i-1}}{\Delta x}) = K \bar{h} \frac{\partial h}{\partial x} \Big|_{i \rightarrow i+1} - K \bar{h} \frac{\partial h}{\partial x} \Big|_{i-1 \rightarrow i} \quad (209)$$

Divide by cell distance, Δx ,

$$S_y \frac{\partial h_i}{\partial t} = \frac{K \bar{h} \frac{\partial h}{\partial x} \Big|_{i \rightarrow i+1} - K \bar{h} \frac{\partial h}{\partial x} \Big|_{i-1 \rightarrow i}}{\Delta x} \quad (210)$$

⁵⁸In the context of Figure 112, the cell face is halfway between the two wells; the cell centers are at the wells.

Take limit as $\Delta x \rightarrow 0$,

$$S_y \frac{\partial h}{\partial t} = \frac{\partial}{\partial x} \left(K h \frac{\partial h}{\partial x} \right) \quad (211)$$

Equation ?? is the equation for unconfined aquifer flow in a horizontal unconfined aquifer. In the next chapter we will start with Equation 209 and build our numerical method using it as the base construct for the difference equations.

17 Steady Groundwater Flow

This chapter presents methods for the approximation of heads in steady groundwater flow where the time derivative vanishes.

17.1 Confined Aquifer Flow

We will address approximating heads in a confined aquifer, first in one spatial dimension, then in two spatial dimensions. The 2D example is a vertical aquifer slice, but the method described would also work just fine for a horizontal layer. The extension to a third spatial dimension is relatively straightforward if one can picture each 2D horizontal slice as a layer.

17.1.1 Finite-Difference Methods – 1 Spatial Dimension

Using Equation 203 as a starting point for simulating aquifer behavior, the steady flow condition means that the left hand side vanishes (there is no change in storage).

$$0 = (K \Delta y \Delta z \frac{h_{i-1} - h_i}{\Delta x}) - (K \Delta y \Delta z \frac{h_i - h_{i+1}}{\Delta x}) \quad (212)$$

Next we will use the arithmetic mean values of the material properties (K) at the cell interfaces, so the difference equation becomes

$$0 = (\frac{1}{2}(K_{i-1} + K_i) \Delta y \Delta z \frac{h_{i-1} - h_i}{\Delta x}) - (\frac{1}{2}(K_i + K_{i+1}) \Delta y \Delta z \frac{h_i - h_{i+1}}{\Delta x}) \quad (213)$$

Now let's group some constants:

$$\begin{aligned} A_i &= \frac{1}{2\Delta x} (K_{i-1} + K_i) \Delta y \Delta z \\ B_i &= \frac{1}{2\Delta x} (K_i + K_{i+1}) \Delta y \Delta z \end{aligned} \quad (214)$$

Now substitute into the difference equation

$$0 = A_i(h_{i-1}) - (A_i + B_i)(h_i) + B_i(h_{i+1}) \quad (215)$$

Now all that remains is to specify boundary conditions, and then implement an algorithm to solve the resulting system of algebraic equations.⁵⁹

Some of the plausible boundary conditions are:

1. Specified head boundary (pretty easy to specify in a computer representation).

⁵⁹I have assumed that the spatial step, Δx is the same for each cell – it doesn't have to be, but relaxing that assumption complicates the specifications of the constants A and B .

2. Zero-flux boundary (also easy to specify using the cell-centered formulation herein).
3. Specified flux boundary (using a modeling trick not too hard to specify).

These three types of conditions should handle the majority of practical situations where one would need to model an aquifer system.

Lets examine the difference equation a little bit – assume we have the correct values then

$$h_i = \frac{A_i(h_{i-1}) + B_i(h_{i+1})}{(A_i + B_i)} \quad (216)$$

which suggests a nice algorithm. We will simply apply boundary conditions, then evaluate the expression for each cell, and after we do the expression for al the cells, we will repeat the process until the solution stops changing. Computationally, we are employing a Jacobi iteration scheme, which will work nicely for this particular problem structure. An alternative, equally valid, would be to construct the linear system of equations (in this case it will be a three-banded matrix), and apply an appropriate row reduction technique to find the solution.

Here are a few examples:

Example 1: 1D Steady Flow in a Confined Aquifer using Jacobi Iteration

First the iterative approach.

Consider the aquifer depicted in Figure 112. Suppose that the two wells are located 400 meters apart, and we wish to approximate the head distribution in the aquifer using the multiple-cell-balance model. Further suppose that the desired cell spacing is $\Delta x = 100$ meters. Additionally suppose the aquifer is 100 meters wide ($\Delta y = 100$ meters), and the thickness is 50 meters ($\Delta z = 50$ meters). The hydraulic conductivity everywhere in the aquifer is 0.2 meters per second.

Figure 115 is the original sketch, annotated with cell spacing, material properties, and the head at the two wells. In the figure, Well 1 has a head measurement of 100 meters, whereas Well 2 has a head measurement of 60 meters. Because the flow is steady, and the material properties are spatially invariant, the EGL/HGL would be a line (the blue one in the figure) connecting the wells sloping from Well 1 to Well 2.

We will use the cell balance model to estimate the EGL/HGL in the aquifer (i.e. find the piezometric surface⁶⁰ between the two wells).

Listing 52. R code demonstrating an Aquifer Flow Simulator for Steady Flow

This fragment of code contains

```
# 1D-confined-aquifer-steady-flow
```

⁶⁰Hydraulic head or piezometric head is a specific measurement of liquid pressure above a geodetic datum. It is usually measured as a liquid surface elevation, expressed in units of length, at the entrance (or bottom) of a piezometer. In an aquifer, it can be calculated from the depth to water in a piezometric well, and given information of the piezometer's elevation and screen depth.




Figure 115. Schematic diagram of unidirectional flow in a generic aquifer, showing heads in two measuring wells. The diagram is annotated with the cell spacing for the example (in magenta), as well as the cell ID. The aquifer material properties are constant ($K=0.2$ m/s). The head in the left well (Well 1) is 100 meters. The head in the right well (Well 2) is 60 meters..

```
# Implements Finite-Difference Porous Medium Flow using Jacobi Iteration
# Assumes boundary cells 1 and ncells are fixed head cells.
zz <- file("input1.dat", "r") # Open a connection named zz to file named input.dat
#
# read the simulation conditions
#
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
ncells <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
hydhead <-(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcond <-(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
#
# split the multiple column strings into numeric components for a vector
#
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcond <-as.numeric(unlist(strsplit(hydcond,split=" ")))
#
# built a position array for plotting
#
distance<-numeric(ncells)
distance[1]<-deltax/2.0
for (i in 2:ncells){distance[i]<-distance[i-1]+deltax}
#
# Plot Distance vs. Head before calculations
#
plot(distance,hydhead,col="red",xlim=c(0,deltax*ncells),ylim=c(0,max(hydhead)*2.0),pch=21,
tck=1)
lines(distance,hydhead,col="red",type="l",lwd=3)
#
# built the transmissivity arrays
#
amat<-numeric(ncells) # make an ncells long array
bmat<-numeric(ncells) # make an ncells long array
for(irow in 2:(ncells-1)){
amat[irow]<-((hydcond[irow-1]+hydcond[irow ])*deltax*deltaz)/(2.0*deltax)
bmat[irow]<-((hydcond[irow ]+hydcond[irow+1])*deltax*deltaz)/(2.0*deltax)
}
#
# #se Jacobi iteration to find a solution to the difference equations
#
```

```

headold<-hydhead # copy the head array, used to test for stopping
maxit <- 100 # set the maximum number of iterations (to prevent infinite loop)
for (iter in 1:maxit){
  for (irow in 2:(ncells-1)){
    hydhead[irow]<-(amat[irow]*hydhead[irow-1]+bmat[irow]*hydhead[irow+1])/(amat[irow]+bmat
      [irow])
  }
  # test for stopping iterations
  percentdiff <- sum((hydhead-headold)^2)
  if (percentdiff < tolerance){break}
  headold<-hydhead
}
#
# Add the steady-flow solution to the graph -- both points and a line.
#
lines(distance,hydhead,col="blue",type="p")
lines(distance,hydhead,col="blue",type="l",lwd=3)

```

Listing 52 is a **R** script that reads an input file, computes the head distribution (the “blue line”) and plots the result. The script also plots the initial values of the heads used in the computation engine. Only the left most value and right most value remain unchanged as they represent the boundary conditions for the problem.

The iterative method is elementary for this case – the replication of the head array at each step is used to test for stopping when the solution meets some tolerance. The compute, test, update part of the script is a common structure in iterative problems, and when we modify the program for transient (time-varying) behavior it will come in handy. There is no error trapping in the example – so it is quite possible to supply an input file that would not work.

Listing 53 is the contents of the input file (`input1.dat`) that is read by the script. To model different conditions, we would change the input file contents and leave the script alone.

Listing 53. Input File for Example Problem

This fragment of code contains

```

100
100
50
5
1e-12
100 100 100 100 60
0.2 0.2 0.2 0.2 0.2

```

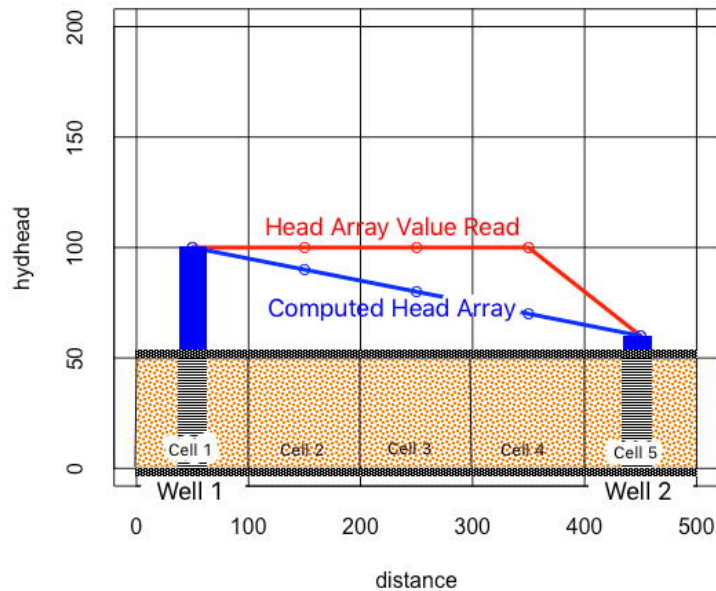


Figure 116. Plot of computed head in each cell using Jacobi iteration. The head in the left well (Well 1) is 100 meters. The head in the right well (Well 2) is 60 meters. The steady flow solution is the blue line labelled “Computed Head Array”..

Finally, for this problem Figure 116 is the output graphics from the script. The red line is the value of heads supplied in the input file (these can be any values as long as the left and right values represent the conditions at the two wells); the blue line is the computed piezometric surface. The markers are the computed head for each cell.

Example 2: 1D Steady Flow in a Confined Aquifer using Gaussian Reduction The second example is the identical physical situation, but instead of iterative computations in our code, we will instead construct a simultaneous linear system ($Ax = b$) where the coefficient matrix A is constructed from the difference equations and solved simultaneously.

First we have to specify how to construct the matrices – essentially because h_1 and h_5 are known, there are only three unknowns in the 5-cell example. So we will have a linear system with 3 equations and 3 unknowns. Here is the linear system in the context of the development of the difference equations.

$$\begin{array}{rclcl}
 -(A_2 + B_2) \times h_2 & B_2 \times h_3 & 0 \times h_4 & = & -A_3 \times h_1 \\
 A_3 \times h_2 & -(A_3 + B_3) \times h_3 & B_3 \times h_4 & = & 0 \\
 0 \times h_2 & A_4 \times h_3 & -(A_4 + B_4) \times h_4 & = & -B_4 \times h_5
 \end{array}$$

A more compact representation is

$$\begin{pmatrix} -(A_2 + B_2) & B_2 & 0 \\ A_3 & -(A_3 + B_3) & B_3 \\ 0 & A_4 & -(A_4 + B_4) \end{pmatrix} \begin{pmatrix} h_2 \\ h_3 \\ h_4 \end{pmatrix} = \begin{pmatrix} -A_3 h_1 \\ 0 \\ -B_4 h_5 \end{pmatrix}$$

Recall that the values for the coefficient matrix are known, as are the values h_1 and h_5 . Now we will modify the script, to use the linear system solver and dispense with the iteration entirely.

The modification(s) are to read in the values for the material and spatial properties – just use the same code. Then construct the compact matrix-vector equation by careful looping on the index values. In the 5 cell example, only the inner three cells [2-4] are part of the linear system. So we use indexing in the arithmetic to build the coefficient matrix and right-hand-side.

Then when we have the solution, we need to put the computed values back into their correct position in the head vector. Listing 54 is a code listing that performs the various tasks. The script is intentionally built to use the same input file as the prior example.

The computed results are identical (as anticipated). The next step (and the point of building such tools) is to try different spatial sizing (partly to be sure the code is generic and automatically adapts to such changes), and to apply the tool to aquifers with different material properties.

Listing 54. R code demonstrating an Aquifer Flow Simulator for Steady Flow

This version constructs coefficient matrix then solves the linear system. The program uses the same input file.

```
# 1D-confined-aquifer-steady-flow
# Implements Finite-Difference Porous Medium Flow using Gaussian reduction
# Assumes boundary cells 1 and ncells are fixed head cells.
zz <- file("input1.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditions
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
ncells <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
hydhead <-readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcond <-readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcond <-as.numeric(unlist(strsplit(hydcond,split=" ")))

# built a position array for plotting
distance<-numeric(ncells)
distance[1]<-deltax/2.0
for (i in 2:ncells){distance[i]<-distance[i-1]+deltax}
plot(distance,hydhead,col="red",xlim=c(0,deltax*ncells),ylim=c(0,max(hydhead)*2.0),pch=21,
tck=1)
lines(distance,hydhead,col="red",type="l",lwd=3)
# built the transmissivity arrays
amat<-numeric(ncells) # make an ncells long array
bmat<-numeric(ncells) # make an ncells long array
for(irow in 2:(ncells-1)){
amat[irow]<-((hydcond[irow-1]+hydcond[irow ])*deltay*deltaz)/(2.0*deltax)
bmat[irow]<-((hydcond[irow ]+hydcond[irow+1])*deltay*deltaz)/(2.0*deltax)
}
amatrix <- matrix(0,ncells-2,ncells-2) # prefill matrix with zeros
rhs <- matrix(0,ncells-2) # prefill vector with zeros
#####
## build matrices here -- there are array indexing things going on #
## because we only have three equations to deal with. #
## hydhead[1] is the left boundary #
## hydhead[ncells] is the right boundary #
## Only build the non-zero elements using the structure of the #
## difference-equation stencil (mask, computational molecule ...) #
#####
## first row special
for (irow in 1:1){
amatatrix[irow,1]=-1.0*(amat[irow+1]+bmat[irow+1])
amatatrix[irow,2]=bmat[irow+1]
rhs[irow] = -amat[irow+1]*hydhead[irow]
}
## interior rows
for (irow in 2:(ncells-3)){
amatatrix[irow,irow-1]=amat[irow+1]
amatatrix[irow,irow]= -1.0*(amat[irow+1]+bmat[irow+1])
amatatrix[irow,irow+1]= bmat[irow+1]
}
## last row special
for (irow in (ncells-2):(ncells-2)){
amatatrix[irow,ncells-3]=amat[irow+1]
amatatrix[irow,ncells-2]=-1.0*(amat[irow+1]+bmat[irow+1])
rhs[irow] = -bmat[irow+1]*hydhead[irow+2]
}
#####
## Now solve the linear system amatrix*unkvector=rhs for unkvector #
## then put values from unkvector into correct position hydhead #
#####
unkvector<-solve(amatrix,rhs)
for (irow in 2:(ncells-1)){
hydhead[irow] <- unkvector[irow-1]
}
#####
## Plot the computed values in blue #
#####
lines(distance,hydhead,col="blue",type="p")
lines(distance,hydhead,col="blue",type="l",lwd=3)
```


The next part of the example is what is the cell size is changed? The whole point of input files and such to keep the script generic. So to check this situation, lets halve the space size (50 meters spacing, instead of 100 meters). Thus instead of 5 cells of 100 meters each, we now have 5 cells of 50 meters each. Listing 55 is a listing of what the input file looks like, other than twice as many initial heads and K values, essentially the same input.

Listing 55. Input File for Example Problem

```
50 << changed cell size
100
50
9 << more cells
1e-12
100 100 100 100 100 100 100 100 60 << more cells
0.2 0.2 0.2 0.2 0.2 0.2 0.2 0.2 0.2 0.2 << more cells
```




Figure 117. Plot of computed head in each cell using different cell spacing. The head in the left well (Well 1) is 100 meters. The head in the right well (Well 2) is 60 meters. The steady flow solution is the blue line labelled “Computed Head Array”..

Lastly, before we move to 2-Dimensional cases, lets examine when the material properties change. In this change, we will leave the first third of the aquifer with the same properties, but decrease hydraulic conductivity in the next third by $1/2$ and the last third by $1/4$ again.

Listing 56 is the input file for this case. The anticipated result is that the HGL/EGL will change slope twice (starting out shallow, and getting steeper as we move to the right).

Listing 56. Input File for Example Problem

```

50
100
50
9 <<
1e-12
100 100 100 100 100 100 100 100 60 << more cells
0.2 0.2 0.2 0.1 0.1 0.1 0.05 0.05 0.05 << more cells

```

Figure 118 is the results with this different input file and indeed as anticipated there are three different slopes (it looks like curvature on the plot, but its really just three different line segments with different slopes.).




Figure 118. Plot of computed head in each cell using different cell spacing. The head in the left well (Well 1) is 100 meters. The head in the right well (Well 2) is 60 meters. The steady flow solution is the blue line. The apparent curvature is really the three different slopes anticipated as the material properties change..

17.1.2 Finite-Difference Methods – 2 Spatial Dimensions

If we perform an analysis in the same way as we did to arrive at Equation 203 except now include another direction (the y -direction) we will have an aquifer in two spatial dimensions. The governing equation becomes

$$S \frac{\partial h}{\partial t} = \frac{\partial}{\partial x} (T_x \frac{\partial h}{\partial x}) + \frac{\partial}{\partial y} (T_y \frac{\partial h}{\partial y}) \quad (217)$$

The meanings of the terms are the same, except the transmissivity terms now have subscripts to indicate they can have different values depending on direction.

Then as before we will construct the difference-equation model from a multiple-cell balance model of the aquifer at a cell of interest, then extend the equations to cover the entire model domain.

Figure 119 is a plan view schematic of a aquifer with flow to be computed in two directions (x and y). The cell indexing convention in the sketch is that rows are indexed by the letter j and columns are indexed by the letter i . This naming convention is arbitrary; in some instances it may be preferable to reverse the convention. The schematic also shows the assumed flow directions; for column i , flows are upward in the drawing, and for row j , flows are from left to right. If indeed the opposite is true for a given set of boundary conditions and material properties, then the flows will be computed as negative numbers – hence the convention here is that “positive flow” is right and up.




Figure 119. Plan view schematic of 2-dimensional multiple cell balance computational stencil.

As in the one-dimensional development the storage term is

$$\left. \frac{dM_{water}}{dt} \right|_{cell} = \rho_w S_s \Delta x \Delta y \Delta z \frac{\partial h_i}{\partial t} \quad (218)$$

where h_i is the head in the i -th cell.

The mass flows entering the i -th, j -th cell are:

$$M_{Inflow} = Q_{left} + Q_{bottom} = \rho_w K_x \Delta y \Delta z \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \rho_w K_y \Delta x \Delta z \frac{h_{i,j-1} - h_{i,j}}{\Delta y} \quad (219)$$

The mass flows leaving the i -th, j -th cell are:

$$M_{Outflow} = Q_{right} + Q_{top} = \rho_w K_x \Delta y \Delta z \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \rho_w K_y \Delta x \Delta z \frac{h_{i,j} - h_{i,j+1}}{\Delta y} \quad (220)$$

Now write the entire balance equation

$$\rho_w S_s \Delta x \Delta y \Delta z \frac{\partial h_i}{\partial t} = [\rho_w K_x \Delta y \Delta z \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \rho_w K_y \Delta x \Delta z \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] -$$

$$[\rho_w K_x \Delta y \Delta z \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \rho_w K_y \Delta x \Delta z \frac{h_{i,j} - h_{i,j+1}}{\Delta y}]$$
(221)

Next replace $S_s \Delta z$ with the storage coefficient S , and the $K_{x,y} \Delta z$ with the transmissivity $T_{x,y}$ terms, and divide by the density of the fluid and the cell plan view area $\Delta x \Delta y$ to obtain a more compact form of the difference equation.

$$S \frac{\partial h_i}{\partial t} = [\frac{1}{\Delta x} T_x \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] -$$

$$[\frac{1}{\Delta x} T_x \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j} - h_{i,j+1}}{\Delta y}]$$
(222)

As in the one-dimensional case, lets again consider steady flow (we will do transient flows later on)

$$0 = [\frac{1}{\Delta x} T_x \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] -$$

$$[\frac{1}{\Delta x} T_x \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j} - h_{i,j+1}}{\Delta y}]$$
(223)

Also as in the one-dimensional case, we will approximate the spatial variation of the material properties (transmissivity) as arithmetic mean values between two cells, so making the following definitions:

$$A_{i,j} = \frac{1}{2\Delta x^2} (T_{x,(i-1,j)} + T_{x,(i,j)})$$

$$B_{i,j} = \frac{1}{2\Delta x^2} (T_{x,(i,j)} + T_{x,(i+1,j)})$$

$$C_{i,j} = \frac{1}{2\Delta y^2} (T_{y,(i,j-1)} + T_{y,(i,j)})$$

$$D_{i,j} = \frac{1}{2\Delta y^2} (T_{y,(i,j)} + T_{y,(i,j+1)})$$
(224)

Substitution into the difference equation yields

$$0 = A_{i,j} h_{i-1,j} + B_{i,j} h_{i+1,j} - (A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j}) h_{i,j} + C_{i,j} h_{i,j-1} + D_{i,j} h_{i,j+1}$$
(225)

As before we can explicitly write the cell equation for $h_{i,j}$ as

$$h_{i,j} = \frac{[A_{i,j}h_{i-1,j} + B_{i,j}h_{i+1,j} + C_{i,j}h_{i,j-1} + D_{i,j}h_{i,j+1}]}{[A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j}]} \quad (226)$$

This difference equation represents an approximation to the governing flow equation, the accuracy depending on the cell size. Boundary conditions are applied directly into the analogs (another name for the difference equations) at appropriate locations on the computational grid. Also as in the one-dimensional case we can generate solutions either by iteration or solving the resulting linear system.

Example 1: 2D Steady Flow in a Confined Aquifer using Jacobi Iteration

As before we will start the example with a simple physical condition and use Jacobi iteration⁶¹ to obtain a solution. We will also incorporate two kinds of boundary conditions (fixed head as before, and no-flow boundaries).

Figure 120 is a schematic of this example. The right panel of the figure shows the index naming convention. The known material properties are transmissivity (in each direction, at each cell center, and the thickness of the aquifer ($b == \Delta z$)). Our task




Figure 120. Schematic of 2-dimensional aquifer. The left and right boundaries are constant head boundaries, whereas the upper and lower boundaries are no-flow.

is to simulate the aquifer with the 5 x 5 model shown. The left and right boundaries will be treated as specified head boundaries. The upper and lower boundary will be treated as no flow boundaries.

The head on the left is 100 meters and the right is 60 meters (same as before). The transmissivity ($K_{x,y}\Delta z=10$) square meters per second (but to be consistent with the earlier models we will supply a value of K and Δz ; keeping with the earlier

⁶¹Jacobi iteration for large domain problems (say 200x200) or bigger, is pretty inefficient – better iterative methods are available; however they represent clever changes to the basic iteration methods explained here, hence Jacobi is a good place to start.

examples the values are $K = 0.2$ meters per second, and $\Delta z = 50$ meters. The spatial dimensions are $\Delta x = 100$ meters and $\Delta y = 100$ meters.

Listing 57 is a listing that implements the method – in this case there are changes to the data reading (to read and build matrices), and notice how the no-flow boundary conditions are implemented.

Listing 57. R code demonstrating an Aquifer Flow Simulator for 2D Steady Flow. This code fragment implements the Jacobi iteration. A subsequent listing shows the contour plot syntax. In the example the two fragments are joined and run as a single source file

```
# Implements 2D-Confined Aquifer Steady Flow Finite-Difference using Jacobi Iteration
# Assumes no-flow boundary row=1, and nrows. Assumes fixed head boundary col=1, and ncols
# Head boundary conditions are entered in input file
zz <- file("input1.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
nrows <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
ncols <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
hydhead <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
hydcondx <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
hydcondy <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <-matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltaz)/(2.0*deltax^2)
    bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltaz)/(2.0*deltax^2)
    cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltay)/(2.0*deltay^2)
    dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltay)/(2.0*deltay^2)
  }
}
headold<-hydhead # copy the head array, used to test for stopping
maxit <- 100 # set the maximum number of iterations (to prevent infinite loop)
for (iter in 1:maxit){
  # first and last row special == no flow boundaries
  for(jcol in 1:ncols){
    hydhead[1,jcol]=hydhead[2,jcol]
    hydhead[nrows,jcol]=hydhead[nrows-1,jcol]
  }
  for (irow in 2:(nrows-1)){
    for (jcol in 2:(ncols-1)){
      hydhead[irow,jcol] =
        (amat[irow,jcol]*hydhead[irow-1,jcol ] +
         bmat[irow,jcol]*hydhead[irow+1,jcol ] +
         cmat[irow,jcol]*hydhead[irow ,jcol-1] +
         dmat[irow,jcol]*hydhead[irow ,jcol+1] )/
        (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
    }
  }
  # test for stopping iterations
  percentdiff <- sum((hydhead-headold)^2)
  if (percentdiff < tolerance){
    message("Exit iterations because tolerance met")
    break}
  headold<-hydhead #update the current head vector
}
```

Instead of line plots, the built-in contouring algorithm in **R** is used to render the output plot(s). Listing 58 is the script that generates a contour plot. The rows are actually plotted in the vertical, and columns in the horizontal, so the plot is rotated relative to the definition sketch.⁶²

Listing 58. R code demonstrating building a contour plot from the computed head distribution

```
#####
###      built position arrays for contour plotting      ###
#####
distancex<-numeric(nrows)
distancey<-numeric(ncols)
distancex[1]<-50
distancey[1]<-50
for (irow in 2:nrows){distancex[irow]<-distancex[irow-1]+deltax}
for (jcol in 2:ncols){distancey[jcol]<-distancey[jcol-1]+deltay}
#####
###      contour plot of head -- note axes are rotated      ###
#####
contour(distancex,distancey,hydhead,
        plot.title = title(main = "Steady 2D Aquifer Model (Head in Meters)",
                           xlab = "Meters (Y axis) ==>>>",
                           ylab = "Meters (X axis) ==>>>"))
```

Listing 59 is a listing of the input file. The only major change from the one-dimensional examples is the entire head and transmissivity arrays are supplied.

Listing 59. Input File for Example Problem

```
100
100
 50
 5
 5
1e-12
100 100 100 100 60
100 100 100 100 60
100 100 100 100 60
100 100 100 100 60
100 100 100 100 60
100 100 100 100 60
0.2 0.2 0.2 0.2 0.2
0.2 0.2 0.2 0.2 0.2
0.2 0.2 0.2 0.2 0.2
0.2 0.2 0.2 0.2 0.2
0.2 0.2 0.2 0.2 0.2
```

Figure 121 is a screen capture of the **R** output (using the R studio IDE). The plot is on the right side of the screen image. Figure 122 is a screen capture of just the plot, rotated, and with the boundaries identified (no-flow top and bottom; fixed head left and right).

Inspection of the script shows that there are still some parts of the script that could use generalization (namely the graphics portion, and a more sophisticated approach to boundary conditions), but otherwise we have the beginnings of a useful tool. To demonstrate the utility, lets change the input file to represent a different problem.

⁶²This kind of plotting is a hold-over from line-printer days, where the long axis would be oriented to the vertical so that it could print to its heart's content and still fit on the paper. Older tractor-feed line-printers could print 135 characters wide, and as long as the paper roll held out. The paper was called green-bar; each perforated sheet was 11x17 and the sheets were connected. Essentially the paper length was functionally infinite, but the width was fixed at 17 inches.

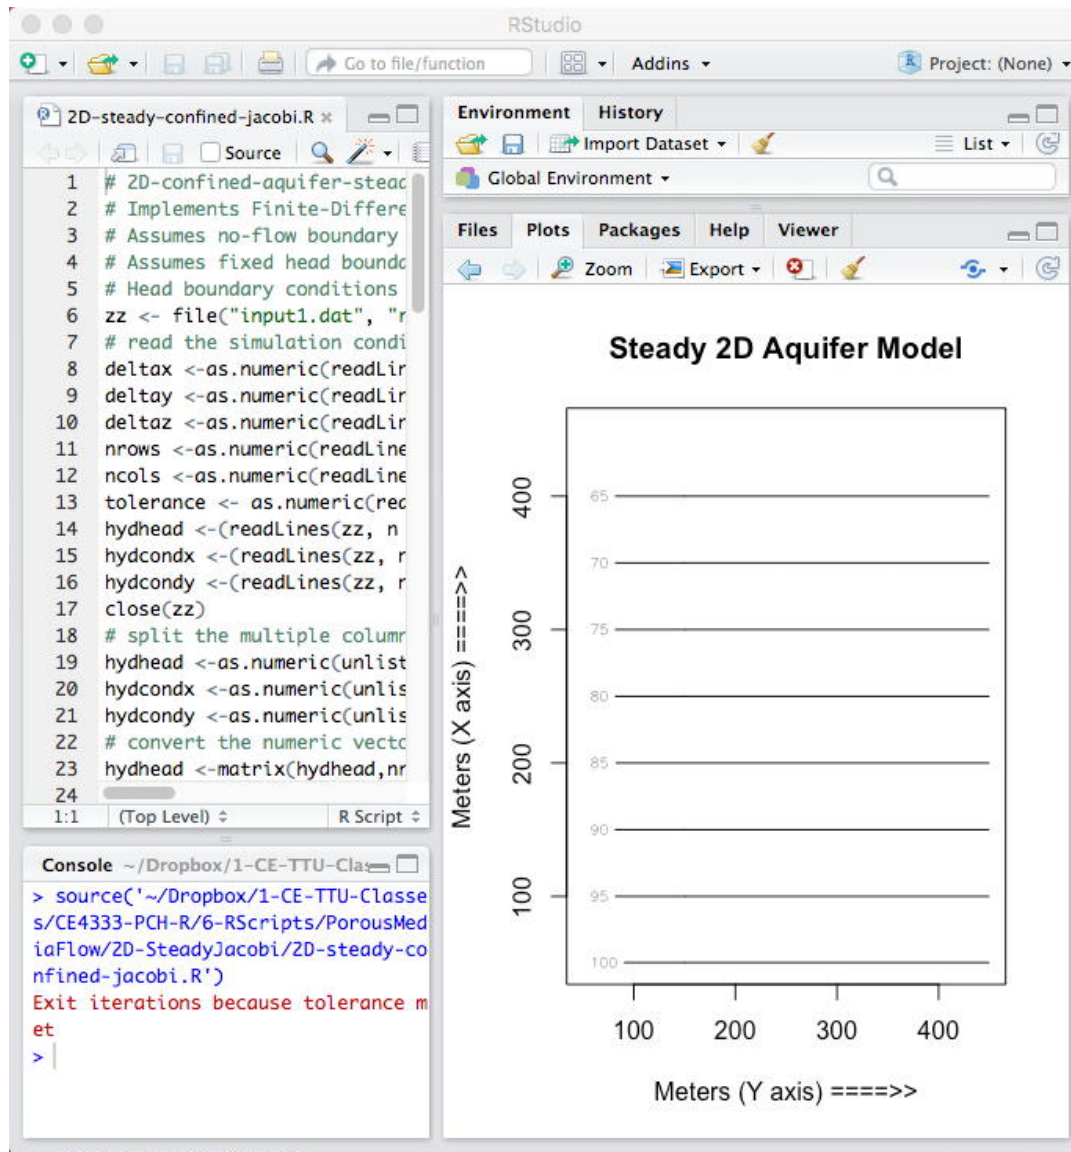


Figure 121. Output from R script for 2D model. Observe how the X-axis is plotted upward, and the Y-axis is plotted left to right. The plotting is because of how ROWS and COLUMNS were indexed in the script. Rather than alter the script, I find it easier to rotate the problem for practical application..




Figure 122. Output from R script for 2D model, rotated and annotated to fit the original problem statement.

Example 3: 2D Vertical Slice in a Confined Aquifer using Jacobi Iteration with Low Permeability Inclusion

Figure 146 is a schematic of a different situation that now only requires us to change the contents of the data file, and re-run the scripts unchanged. Some additional modifications have been added, mostly messages to the user because of anticipated long run times. The data file is changed a bit and two lines are added to help with the plotting – they represent the axis labels used in the contour plots. The boundary conditions are still directly coded into the algorithm, and that would be the next modification to the code - general boundary condition information.⁶³




Figure 123. Schematic of vertical slice in an aquifer with low permeability inclusion. Values are indicated on the schematic .

The following pages contain the code fragments (listings) for the head calculations, and the contour plotting. As above, these listings are combined into a single file (the fragmentation herein is to fit the printed page layout) and then run as a script.

Listing 60 is the listing for the head calculations. Listing 61 is the is the listing for the plotting calculations. Listing 84 is the input file for the problem. The file in this case is named `input2.dat`.

In addition the generalized boundary conditions, the reader should consider making the program prompt the user for the file name, so that the program is more deployable. Once a generalized boundary condition component is incorporated as well as a file prompt, the script would be useful in many practical situations. When we add wells to the governing PDE the tool would be quite relevant for engineering calculations.

⁶³That modification is left for homework.

Listing 60. 2D Steady Confined Aquifer Flow Script.

```

# 2D Aquifer Flow Model using Jacobi Iteration
# deallocate memory
rm(list=ls())
zz <- file("input2.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
nrows <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
ncols <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
maxiter <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
distancex <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
distancey <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydhead <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondx <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondy <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
distancex <-as.numeric(unlist(strsplit(distancex,split=" ")))
distancey <-as.numeric(unlist(strsplit(distancey,split=" ")))
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <- matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
# here we perform the velocity potential calculations
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltaz)/(2.0*deltax
^2)
    bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltaz)/(2.0*deltax
^2)
    cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltaz)/(2.0*deltay
^2)
    dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltaz)/(2.0*deltay
^2)
  }
}
headold <- hydhead # copy the head array, used to test for stopping
tolflag <- FALSE
for (iter in 1:maxiter){
  # first and last row special == no flow boundaries
  for(jcol in 1:ncols){
    hydhead[1,jcol]<-hydhead[2,jcol]
    hydhead[nrows,jcol]<-hydhead[nrows-1,jcol]
  }
  for (irow in 2:(nrows-1)){
    for (jcol in 2:(ncols-1)){
      hydhead[irow,jcol] <-
        (amat[irow,jcol]*hydhead[irow-1,jcol ] +
         bmat[irow,jcol]*hydhead[irow+1,jcol ] +
         cmat[irow,jcol]*hydhead[irow ,jcol-1] +
         dmat[irow,jcol]*hydhead[irow ,jcol+1 ] )/
        (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
    }
  }
  # test for stopping iterations
  percentdiff <- sum((hydhead-headold)^2)
  if (percentdiff < tolerance){
    message("Exit iterations in velocity potential because tolerance met");
    message("Iterations =", iter);
    tolflag <- TRUE
    break}
  headold<-hydhead #update the current head vector
  if( iter %% 1000 == 0){message("Calculating in Potential Function ",iter," of ",maxiter, "
iterations")}
}
if (tolflag == FALSE){message("Exit iterations in potential function at max iterations")}
}

```





Figure 124. Vertical slice in an aquifer with low permeability inclusion using Jacobi iteration scripts.

system matrix is non-trivial.⁶⁴ The head “array” would need to be addressed as a vector (we can use pointers to accomplish this task — not too hard, but we would need to build the coefficient matrix, solve the system, and de-construct the result). This task is saved for a future edition.

17.1.3 Generalizing the Boundary Conditions

In the prior examples the boundary conditions for the problems were kind of glossed over. We applied a fixed head boundary on the left and right edges of the rectangular domain, and zero-flux boundary at the top and bottom edges. A useful improvement is to allow the user to choose which type by supplying information in the input file. I find the easiest way (as we are just learning) is to assume the entire model is always surrounded by a constant head condition and use a mask to tell the program when that is not true.

The code fragments for making this change are pretty straightforward, and are displayed in Listing 63. We need to read in boundary indicators for the top, bottom, left, and right boundaries. Then convert them into numeric values for later. Here I choose to use a zero to indicate a zero-flux boundary and any non-zero (usually a 1) to indicate a fixed head boundary.

⁶⁴The system will be penta-diagonal, and knowing this structure we could code a very efficient solver – in this case we will use the built-in function, although we will tell the program that the matrix is sparse (e.g. `sparse = TRUE`) in the function call.

Next we apply the conditions within the solver loop by assigning the computed value of the adjacent cell (either above, below, left, or right as appropriate). The remainder of the script is unchanged, of course we need to include the new input values in the data file as in Listing 64.

Listing 63. Script fragments for implementing generalized boundary conditions.

```
### Up in the input read section of the script #####
# add boundary conditions 0= fixed head, 1= no flow
boundarytop <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
boundarybottom <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
boundaryleft <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
boundaryright <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
hydhead <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
.....
### Convert to numeric #####
boundarytop <-as.numeric(unlist(strsplit(boundarytop,split=" ")))
boundarybottom <-as.numeric(unlist(strsplit(boundarybottom,split=" ")))
boundaryleft <-as.numeric(unlist(strsplit(boundaryleft,split=" ")))
boundaryright <-as.numeric(unlist(strsplit(boundaryright,split=" ")))
.....
#### Apply the conditions ####
for (iter in 1:maxiter){
# Boundary Conditions
# Top and Bottom
for(jcol in 1:ncols){
if(boundarytop[jcol] == 0){hydhead[1,jcol]<-hydhead[2,jcol]} #no-flow at top
if(boundarybottom[jcol] == 0){hydhead[nrows,jcol]<-hydhead[2,jcol]} #no-flow at bottom
# otherwise values are fixed head
}
for(irow in 1:nrows){
if(boundaryleft[irow] == 0){hydhead[irow,1]<-hydhead[irow,2]} #no-flow at left
if(boundaryright[irow] == 0){hydhead[irow,ncols]<-hydhead[irow,ncols-1]} #no-flow at
right
# otherwise values are fixed head
}
for (irow in 2:(nrows-1)){
.....
```

Listing 64. Input file for 2D vertical slice confined aquifer with low permeability inclusion using a generalized boundary mask. The comments in the listing should be removed for running the program.

```
1
10
1
13
23
1e-8
5000
5 15 25 35 45 55 65 75 85 95 105 115 125 135 145 155 165 175 185 195 205 215 225
0.5 1.5 2.5 3.5 4.5 5.5 6.5 7.5 8.5 9.5 10.5 11.5 12.5
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 # Boundary Conditions along Top
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 # Boundary Conditions along Bottom
1 1 1 1 1 1 1 1 1 1 1 1 1 # Boundary Conditions along Left
1 1 1 1 1 1 1 1 1 1 1 1 1 # Boundary Conditions along Right
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100 100
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10
```


(shown as locations 1,2,3,4, and 5 in the sketch), which in turn requires computation of head under the dam. Thus the questions are answered by finding the head distribution under the dam.

The flow field (mathematically) extends an infinite distance upstream and downstream, but as a practical matter the contribution to seepage far upstream of the dam is negligible, and hence is approximated by the finite domain depicted.

Using the tools we have already built we can simply build an input file, run our script and determine the head distribution (and thus compute the discharges under the dam. There are two ways to conceptualize the model domain, we will examine both.

The first is to represent the domain as shown, and make the following specifications in the boundary condition information, and we will treat the sheetpile cutoff wall as a low permeability inclusion (much like the the prior example). The boundary conditions are:

1. The segment from A to B is a constant head boundary with value equal to 10.
2. The segment from B to F is a zero-flux boundary.
3. The segment from B to C to D to U to E to F should be treated as a zero-flux boundary, but our mask does not extend into the interior – however the sheetpile itself can be approximated by providing a very small permeability. Alternately we could (should) modify the code to handle interior boundaries – but that is outside the scope of this chapter.
4. The segment from F to G is a constant head boundary with value equal to 0.
5. The segment form G to H is a constant head boundary with value equal to 0.
6. The segment from H to I is a zero-flux boundary.
7. The segment from I to A is a constant head boundary with value equal to 10.

Listing 65 is the input file that is used that incorporates the boundary conditions stated above. Observe how all we did was change the Δx and Δy specifications, modified the boundary condition arrays, and set the row and column counts.

Listing 65. Input file for 2D vertical slice for Dam Seepage Example.

```

10
10
1
9
31
1e-12
5000
0 10 20 30 40 50 60 70 80 90 100 110 120 130 140 150 160 170 180 190 200 210 220 230 240
    250 260 270 280 290 300
0 10 20 30 40 50 60 70 80
1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 10 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

```

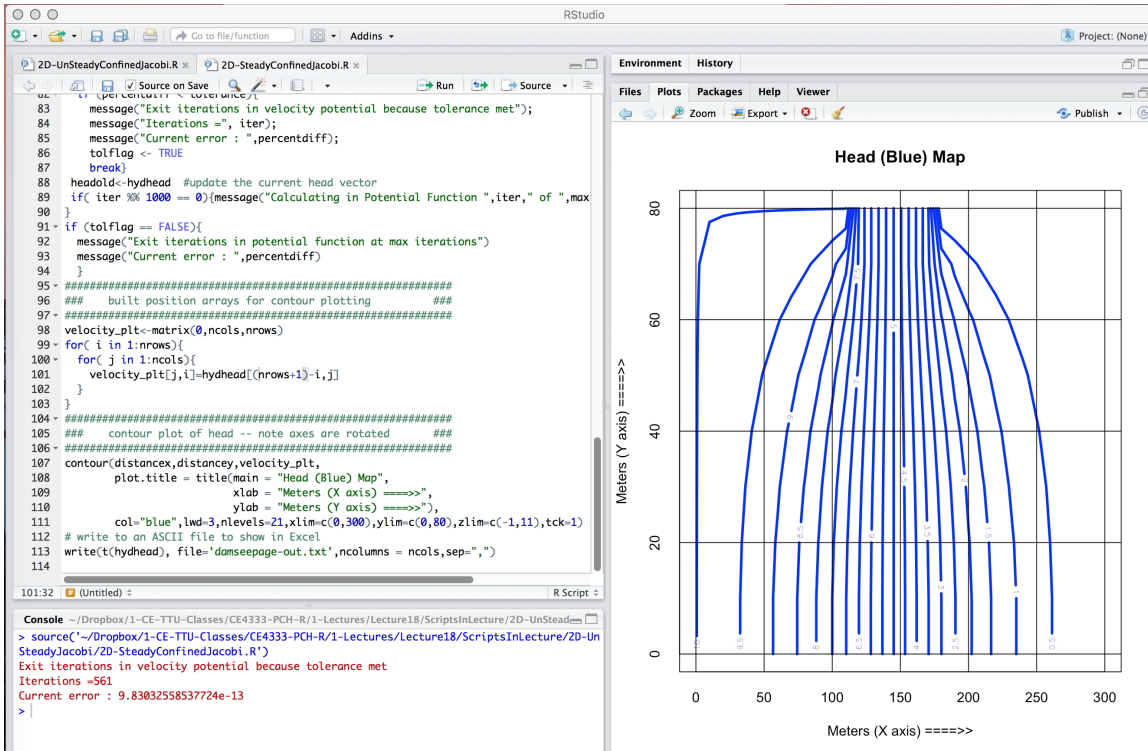



Figure 126. Screen capture of the script run with the different input file.

If we apply Darcy’s law to the orange portion in the vertical direction (where flow must be vertical to seep under the dam) the result is about 3.76 m³ per day, per meter of width (perpendicular to Figure Figure 125). These computations are shown in the spreadsheet beneath the output data.

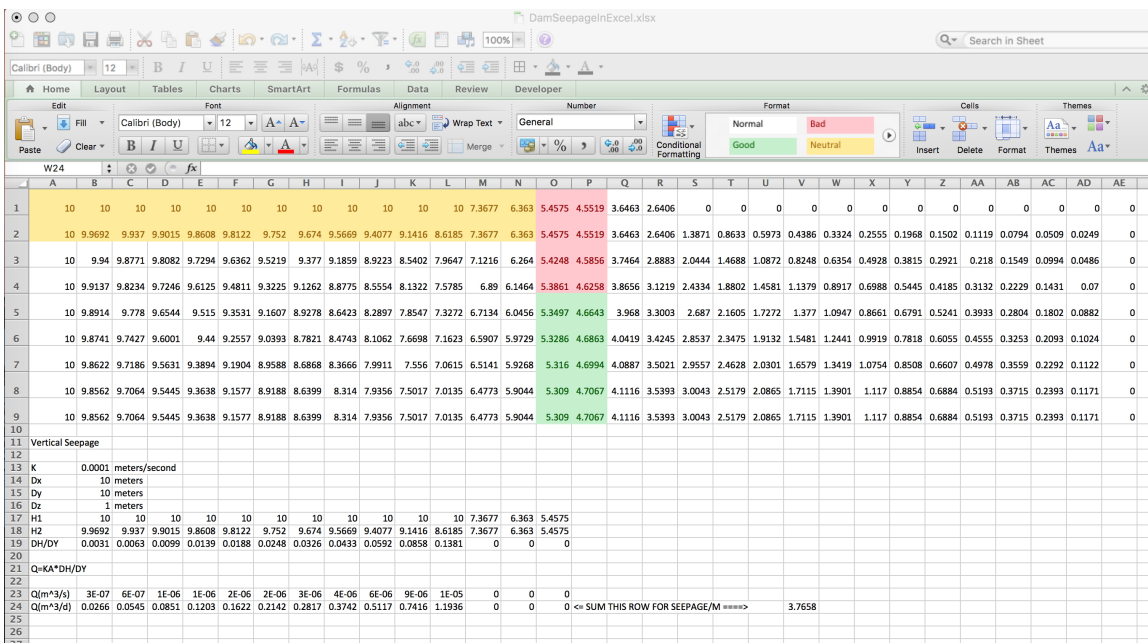


Figure 127. Output file loaded into Excel for further analysis.

This example (like most roll-your-own) requires some specific knowledge of hydraulics to interpret the results, in this case knowing to apply Darcy's law along the bottom of the reservoir to determine the flow rate, and to integrate (sum up) the individual cell flows to obtain total flow.

The second way to solve this example, and perhaps better is to take advantage of the symmetry and cut the domain in half, as in Figure 128. In this method we can actually specify the sheetpile as a boundary, and we will obtain the same results, but only need to supply half as much input data.




Figure 128. Output file loaded into Excel for further analysis.

Listing 66 is the input file for this symmetry based approach.

Listing 66. Input file for 2D vertical slice for Dam Seepage Example using Symmetry.

```

10
10
1
9
15
1e-16
5000
0 10 20 30 40 50 60 70 80 90 100 110 120 130 140
0 10 20 30 40 50 60 70 80
1 1 1 1 1 1 1 1 1 1 1 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 1 1 1 1 1 1 1
0 0 0 0 1 1 1 1
10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 10
10 10 10 10 10 10 10 10 10 10 10 10 10 5
10 10 10 10 10 10 10 10 10 10 10 10 10 5
10 10 10 10 10 10 10 10 10 10 10 10 10 5
10 10 10 10 10 10 10 10 10 10 10 10 10 5

```




Figure 129 is the result of using the half-domain approach.

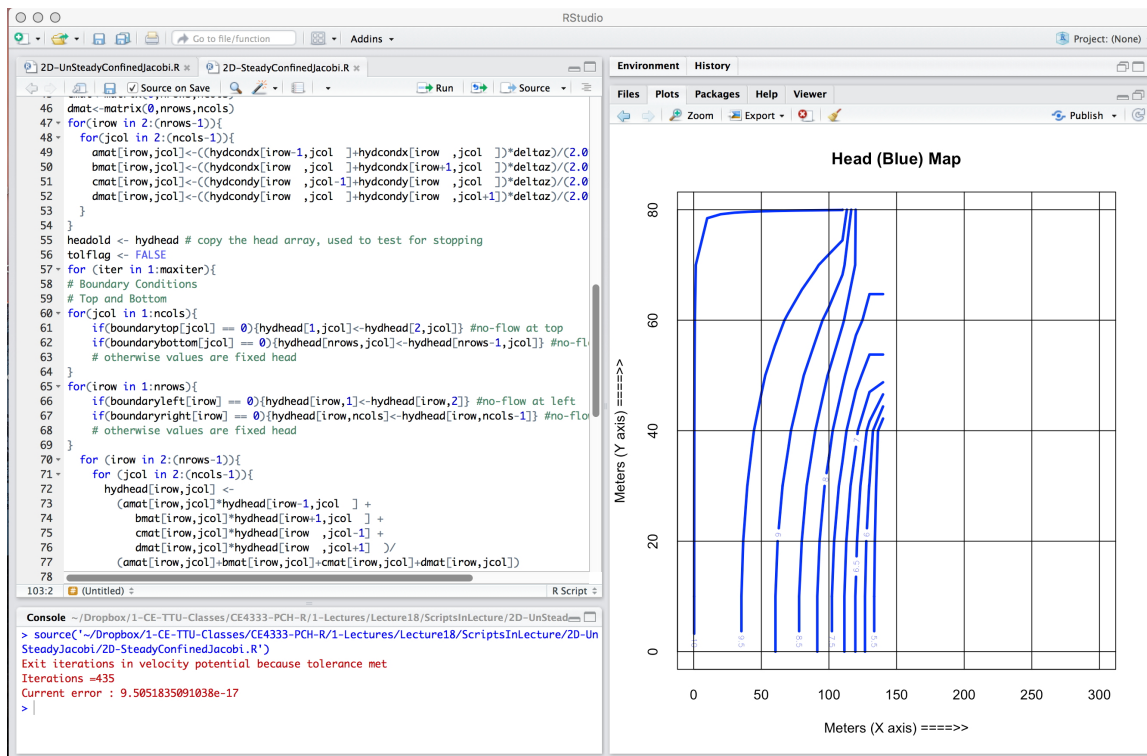


Figure 129. Output file loaded into Excel for further analysis.

If we repeat the flow calculations, this time along the portion at the bottom of the

sheetpile, we obtain nearly the same result, 3.74 m³ per day, per meter of width. This second way of solving the problem is more correct because we didn't have to choose an artificial value to mimic the effect of the sheetpile.

17.1.4 Internal sources/sinks to handle a discharge (well) at a specific location

The last concept we will add here is to illustrate a means to approximate the effect of a localized source (recharge or an injection well) or sink (pumping well). To incorporate these kinds of inputs we add the terms to Equation 217 to obtain Equation 227

$$S \frac{\partial h}{\partial t} = \frac{\partial}{\partial x} \left(T_x \frac{\partial h}{\partial x} \right) + \frac{\partial}{\partial y} \left(T_y \frac{\partial h}{\partial y} \right) + R - Q \quad (227)$$

where R and Q are recharge and pumping expressed in dimensions of $\frac{L^3}{T}$.

The resulting difference equation is

$$0 = \left[\frac{1}{\Delta x} T_x \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j-1} - h_{i,j}}{\Delta y} \right] - \left[\frac{1}{\Delta x} T_x \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j} - h_{i,j+1}}{\Delta y} \right] + \frac{R_{i,j}}{\Delta x \Delta y} - \frac{Q_{i,j}}{\Delta x \Delta y} \quad (228)$$

These additions are incorporated into our program by only adding a few code fragments in certain places. Listing 67 shows the various added components to handle a well field. Observe we can incorporate recharge as if it were a well with a negative pumping rate so for this document it is not considered separately, although in many practical cases that might be a preferable way to incorporate the process.

Listing 67. Input file for 2D Steady flow with generalized boundary conditions and wells.

```
# 2D Steady Confined -- With Boundary Arrays -- And Pumping Array
# 2D Aquifer Flow Model using Jacobi Iteration
# deallocate memory
rm(list=ls())
zz <- file("wellfield.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
deltay <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
deltax <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
nrows <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
ncols <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
maxiter <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
distancex <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown", skipNul =
FALSE))
distancey <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown", skipNul =
FALSE))
# add boundary conditions 0= fixed head, 1= no flow
boundarytop <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown", skipNul =
FALSE))
```

```

boundarybottom <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
boundaryleft <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
  = FALSE))
boundaryright <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
  = FALSE))
hydhead <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
# hydhead is now the initial condition array #
hydcondx <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
hydcondy <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
# add pumping array
pumping <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
distancex <- as.numeric(unlist(strsplit(distancex,split=" ")))
distancey <- as.numeric(unlist(strsplit(distancey,split=" ")))
boundarytop <- as.numeric(unlist(strsplit(boundarytop,split=" ")))
boundarybottom <- as.numeric(unlist(strsplit(boundarybottom,split=" ")))
boundaryleft <- as.numeric(unlist(strsplit(boundaryleft,split=" ")))
boundaryright <- as.numeric(unlist(strsplit(boundaryright,split=" ")))
hydhead <- as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcondx <- as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <- as.numeric(unlist(strsplit(hydcondy,split=" ")))
pumping <- as.numeric(unlist(strsplit(pumping,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <- matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <- matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <- matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
pumping <- matrix(pumping,nrow=nrows,ncol=ncols,byrow = TRUE)
# here we perform the velocity potential calculations
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
qrat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltax)/(2.0*deltax
      ^2)
    bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltax)/(2.0*deltax
      ^2)
    cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltay)/(2.0*deltay
      ^2)
    dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltay)/(2.0*deltay
      ^2)
    qrat[irow,jcol]<-(-1.0)*pumping[irow,jcol]/(deltax*deltay)
  }
}
headold <- hydhead # copy the head array, used to test for stopping
tolflag <- FALSE
for (iter in 1:maxiter){
# Boundary Conditions
# Top and Bottom
for(jcol in 1:ncols){
  if(boundarytop[jcol] == 0){hydhead[1,jcol]<-hydhead[2,jcol]} #no-flow at top
  if(boundarybottom[jcol] == 0){hydhead[nrows,jcol]<-hydhead[nrows-1,jcol]} #no-flow at
    bottom
  # otherwise values are fixed head
}
for(irow in 1:nrows){
  if(boundaryleft[irow] == 0){hydhead[irow,1]<-hydhead[irow,2]} #no-flow at left
  if(boundaryright[irow] == 0){hydhead[irow,ncols]<-hydhead[irow,ncols-1]} #no-flow at
    right
  # otherwise values are fixed head
}
for (irow in 2:(nrows-1)){
  for (jcol in 2:(ncols-1)){
    hydhead[irow,jcol] <-
      (
        qrat[irow,jcol] +
        amat[irow,jcol]*hydhead[irow-1,jcol ] +
        bmat[irow,jcol]*hydhead[irow+1,jcol ] +
        cmat[irow,jcol]*hydhead[irow ,jcol-1] +
        dmat[irow,jcol]*hydhead[irow ,jcol+1] )/
      (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
  }
}
# test for stopping iterations
percentdiff <- sum((hydhead-headold)^2)
if (percentdiff < tolerance){
  message("Exit iterations in velocity potential because tolerance met");
  message("Iterations =", iter);
  message("Current error : ",percentdiff);
  tolflag <- TRUE
  break}
headold<-hydhead #update the current head vector
if( iter % 1000 == 0){message("Calculating in Potential Function ",iter," of ",maxiter, "
  iterations")}

```

```

}
if (tolflag == FALSE){
  message("Exit iterations in potential function at max iterations")
  message("Current error : ",percentdiff)
}
#####
###    built position arrays for contour plotting    ###
#####
velocity_plt<-matrix(0,ncols,nrows)
for( i in 1:nrows){
  for( j in 1:ncols){
    velocity_plt[j,i]=hydhead[(nrows+1)-i,j]
  }
}
#####
###    contour plot of head -- note axes are rotated    ###
#####
contour(distancex,distancey,velocity_plt,
  plot.title = title(main = "Head (Blue) Map",
    xlab = "Meters (X axis) ==>>",
    ylab = "Meters (Y axis) ==>>"),
  col="blue",lwd=3,nlevels=21,xlim=c(0,300),ylim=c(0,300),zlim=c(-1,40),tck=1)
# write to an ASCII file to show in Excel
write(t(hydhead), file='damseepage-out.txt',ncolumns = ncols,sep=",")
message("min head : ",min(hydhead))

```

Example 5: 4 Wells in a rectangular aquifer Figure 130 is a rectangular aquifer with 4 wells as shown. The aquifer thickness is 1 meters. The aquifer is surrounded with a constant head boundary of 30 meters. The hydraulic conductivity is $K = 0.033 \text{ m/day}$.

Using a 10 meter \times 10 meter grid spacing estimate the pumping rate in each well so that the head within the rectangular area defined by the well field is no greater than 15 meters.




Figure 130. Rectangular aquifer with 4 wells.


```
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.430 0.000 0.000 0.000 0.000
0.430 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.430 0.000 0.000 0.000 0.000
0.430 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
```

Figure 131 is a screen capture of the script run with the datafile above. The contour plot shows the “cone of depression” around the well field. The script has an added message that reports the smallest head, and that report was used to adjust the pumping rates at the four wells until the value was at the prescribed 15 meters.

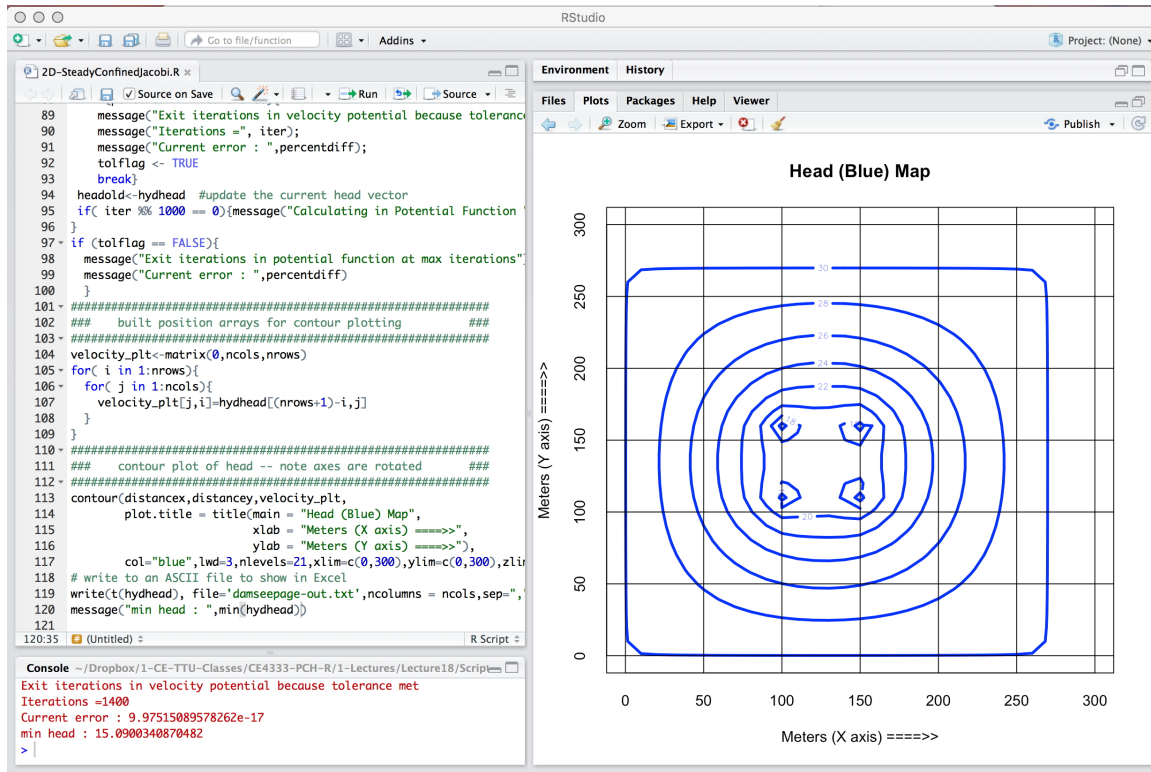


Figure 131. Rectangular aquifer with 4 wells.

The value for a single well is $0.430 \frac{m^3}{day}$ (using the units of hydraulic conductivity as a guide). If the goal was then to specify particular pumps, that flow value and the required lift would be used to specify and purchase pumps for the system.

We will repeat the problem in the next section of the workbook, except the aquifer will be unconfined — we expect different results, but will be able to re-use the input file.

17.2 Unconfined Aquifer Flow

Equation 229 is the meaningful part of Equation 209 from the previous chapter, which is our starting point for building unconfined flow computations.

$$S_y \Delta x \frac{\partial h_i}{\partial t} = (K \bar{h} \frac{h_{i+1} - h_i}{\Delta x}) - (K \bar{h} \frac{h_i - h_{i-1}}{\Delta x}) \quad (229)$$

17.2.1 Finite-Difference Methods

Here we will start with 2D flow – if you have a 1D problem to solve, just make the model 3 cells wide in the un-needed axis and it will accomplish the same goal (and avoid having separate codes for different dimensions).

Here is the unconfined flow equation in 2 dimensions

$$S_y \Delta x \Delta y \frac{\partial h_i}{\partial t} = [K_x \Delta y \bar{h} \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + K_y \Delta x \bar{h} \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] - [K_x \Delta y \bar{h} \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + K_y \Delta x \bar{h} \frac{h_{i,j} - h_{i,j+1}}{\Delta y}] \quad (230)$$

Next divide by the cell plan view area $\Delta x \Delta y$ to obtain a more compact form of the difference equation.

$$S \frac{\partial h_i}{\partial t} = [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] - [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j} - h_{i,j+1}}{\Delta y}] \quad (231)$$

Again consider steady flow (we will do transient flows later on)

$$0 = [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] - [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j} - h_{i,j+1}}{\Delta y}] \quad (232)$$

We will approximate the spatial variation of the $K \bar{h}$) as arithmetic mean values between two cells, so making the following definitions:

$$\begin{aligned} A_{i,j} &= \frac{1}{2\Delta x^2} (K_{x,(i-1,j)} h_{i-1,j} + K_{x,(i,j)} h_{i,j}) \\ B_{i,j} &= \frac{1}{2\Delta x^2} (K_{x,(i,j)} h_{i,j} + K_{x,(i+1,j)} h_{i+1,j}) \\ C_{i,j} &= \frac{1}{2\Delta y^2} (K_{y,(i,j-1)} h_{i,j-1} + K_{y,(i,j)} h_{i,j}) \\ D_{i,j} &= \frac{1}{2\Delta y^2} (K_{y,(i,j)} h_{i,j} + K_{y,(i,j+1)} h_{i,j+1}) \end{aligned} \quad (233)$$

Substitution into the difference equation yields

$$0 = A_{i,j} h_{i-1,j} + B_{i,j} h_{i+1,j} - (A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j}) h_{i,j} + C_{i,j} h_{i,j-1} + D_{i,j} h_{i,j+1} \quad (234)$$

As before we can explicitly write the cell equation for $h_{i,j}$ as, however the system is a bit non-linear as the A, B, C, D coefficients are themselves functions of the cell values.

$$h_{i,j} = \frac{[A_{i,j} h_{i-1,j} + B_{i,j} h_{i+1,j} + C_{i,j} h_{i,j-1} + D_{i,j} h_{i,j+1}]}{[A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j}]} \quad (235)$$

This difference equation represents an approximation to the governing flow equation, the accuracy depending on the cell size. Boundary conditions are applied directly into the analogs (another name for the difference equations) at appropriate locations on the computational grid. Also as in the one-dimensional case we can generate solutions either by iteration or solving the resulting non-linear system. Fortunately these kinds of systems are pretty well conditioned and one can usually compute the A, B, C, D coefficients from the most recent value (in an iterative sense) of h and proceeds as if the system is linear and it usually eventually converges.⁶⁵

The method is again demonstrated by example.

Example 6: 4 Wells in a rectangular unconfined aquifer

Figure 132 is a rectangular aquifer with 4 wells as shown – the drawing is identical to the previous example, except the aquifer is now an unconfined aquifer. The saturated thickness before the pumps are turned on is 30 meters. The aquifer is surrounded with a constant head boundary of 30 meters. The hydraulic conductivity is $K = 0.033 \text{ m/day}$.

Using a $10 \text{ meter} \times 10 \text{ meter}$ grid spacing estimate the pumping rate in each well so that the head within the rectangular area defined by the well field is no greater than 15 meters. In this example, the pumping rates will be quite a bit larger because the saturated thickness at the desired drawdown is 15 times larger than the previous example.

Listing 69 is an **R** script that performs the indicated computations. Upon initial inspection it appears nearly identical to the confined flow script, except the A, B, C, D computations are moved into the iteration loop, so they are updated after each iteration.

Observe that the Δz value is still read into the script, but never used. I am being somewhat lazy here; however it illustrates a code development protocol to re-use as much code as possible until the program functions as desired, then we could go back and delete unused components. I choose to retain the Δz so I would only have to change the input file pumping array and not create an entirely new input file structure. As stated earlier, as the problems get big, we would write scripts that build the input files before we run the computation engines.

⁶⁵The approach here is called the variable transmissivity approach as that's all I really did in the difference equation development – as long as there is not too much curvature in the head distribution or dramatic changes in material properties in adjacent cells the Jacobi iteration method already employed will function, albeit slowly compared to more elaborate methods. In this course we are satisfied with these relatively primitive constructs, and more sophisticated approaches are presented in the groundwater hydrology class taught by others.




Figure 132. Rectangular aquifer with 4 wells.

Listing 69. R script implementing unconfined aquifer flow computations with wells and generalized boundary conditions..

```
# 2D Steady UnConfined -- With Boundary Arrays -- And Pumping Array
# 2D Aquifer Flow Model using Jacobi Iteration
# deallocate memory
rm(list=ls())
zz <- file("wellfield2.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
deltay <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
deltaz <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
# here we retain deltaz but will not use it, we assume head is the deltaz
nrows <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
ncols <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
maxiter <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE, encoding = "unknown",
skipNul = FALSE))
```

```

distancex <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
distancey <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
# add boundary conditions 0= fixed head, 1= no flow
boundarytop <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
boundarybottom <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
boundaryleft <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
boundaryright <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
hydhead <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
# hydhead is now the initial condition array #
hydcondx <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondy <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
# add pumping array
pumping <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
distancex <-as.numeric(unlist(strsplit(distancex,split=" ")))
distancey <-as.numeric(unlist(strsplit(distancey,split=" ")))
boundarytop <-as.numeric(unlist(strsplit(boundarytop,split=" ")))
boundarybottom <-as.numeric(unlist(strsplit(boundarybottom,split=" ")))
boundaryleft <-as.numeric(unlist(strsplit(boundaryleft,split=" ")))
boundaryright <-as.numeric(unlist(strsplit(boundaryright,split=" ")))
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
pumping <-as.numeric(unlist(strsplit(pumping,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <- matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
pumping <-matrix(pumping,nrow=nrows,ncol=ncols,byrow = TRUE)
# here we perform the velocity potential calculations
# allocate the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
# allocate and build the pumping array
qratt<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    qratt[irow,jcol]<-(-1.0)*pumping[irow,jcol]/(deltax*deltay)
  }
}
headold <- hydhead # copy the head array, used to test for stopping
tolflag <- FALSE
for (iter in 1:maxiter){
# built the variable transmissivity arrays
  for(irow in 2:(nrows-1)){
    for(jcol in 2:(ncols-1)){
      amat[irow,jcol]<-((hydcondx[irow-1,jcol ]*hydhead[irow-1,jcol ]
+hydcondx[irow ,jcol ]*hydhead[irow ,jcol ]))/(2.0*deltax^2)

      bmat[irow,jcol]<-((hydcondx[irow ,jcol ]*hydhead[irow ,jcol ]
+hydcondx[irow+1,jcol ]*hydhead[irow+1,jcol ]))/(2.0*deltax^2)

      cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]*hydhead[irow ,jcol-1]
+hydcondy[irow ,jcol ]*hydhead[irow ,jcol ]))/(2.0*deltay^2)

      dmat[irow,jcol]<-((hydcondy[irow ,jcol ]*hydhead[irow ,jcol ]
+hydcondy[irow ,jcol+1]*hydhead[irow ,jcol+1]))/(2.0*deltay^2)

    }
  }

# Boundary Conditions
# Top and Bottom
for(jcol in 1:ncols){
  if(boundarytop[jcol] == 0){hydhead[1,jcol]<-hydhead[2,jcol]} #no-flow at top
  if(boundarybottom[jcol] == 0){hydhead[nrows,jcol]<-hydhead[nrows-1,jcol]} #no-flow at
  bottom
  # otherwise values are fixed head
}
for(irow in 1:nrows){
  if(boundaryleft[irow] == 0){hydhead[irow,1]<-hydhead[irow,2]} #no-flow at left
  if(boundaryright[irow] == 0){hydhead[irow,ncols]<-hydhead[irow,ncols-1]} #no-flow at
  right
  # otherwise values are fixed head
}
for (irow in 2:(nrows-1)){
  for (jcol in 2:(ncols-1)){
    hydhead[irow,jcol] <-
      (
        qratt[irow,jcol] +

```

```

amat[irow,jcol]*hydhead[irow-1,jcol ] +
bmat[irow,jcol]*hydhead[irow+1,jcol ] +
cmat[irow,jcol]*hydhead[irow ,jcol-1] +
dmat[irow,jcol]*hydhead[irow ,jcol+1] )/
(amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
}
}
# test for stopping iterations
percentdiff <- sum((hydhead-headold)^2)
if (percentdiff < tolerance){
  message("Exit iterations in velocity potential because tolerance met");
  message("Iterations =", iter);
  message("Current error : ",percentdiff);
  tolflag <- TRUE
  break}
headold<-hydhead #update the current head vector
if( iter %% 1000 == 0){message("Calculating in Potential Function ",iter," of ",maxiter, "
iterations")}
}
if (tolflag == FALSE){
  message("Exit iterations in potential function at max iterations")
  message("Current error : ",percentdiff)
}
#####
### built position arrays for contour plotting ###
#####
velocity_plt<-matrix(0,ncols,nrows)
for( i in 1:nrows){
  for( j in 1:ncols){
    velocity_plt[j,i]=hydhead[(nrows+1)-i,j]
  }
}
#####
### contour plot of head -- note axes are rotated ###
#####
contour(distancex,distancey,velocity_plt,
        plot.title = title(main = "Head (Blue) Map",
                           xlab = "Meters (X axis) ==>>",
                           ylab = "Meters (Y axis) ==>>"),
        col="blue",lwd=3,nlevels=21,xlim=c(0,300),ylim=c(0,300),zlim=c(-1,40),tck=1)
# write to an ASCII file to show in Excel
write(t(hydhead), file='damseepage-out.txt',ncolumns = ncols,sep=",")
message("min head : ",min(hydhead))

```

Listing 70 is the relevant portion of the input file for the example. The remainder of the file is unchanged from the previous example (hence the reason I retained the Δz variable).

Listing 70. Relevant portion of input file for unconfined flow example. Only the pumping values are different..

```

.....
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 9.500 0.000 0.000 0.000 0.000 0.000
9.500 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 9.500 0.000 0.000 0.000 0.000 0.000
9.500 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000
.....

```

Figure 133 is a screen capture of the example implemented in **R**. The plot is practically the same as in the previous example (in fact of the labels were not moved on the contour lines, the figures would look identical) however the pumping rate is 20 times larger for apparently the same head distribution — so if the dewatering effort were in an unconfined aquifer, perhaps for an excavation for a foundation, choosing the correct conditions would be kind of important.⁶⁶

⁶⁶Well more than “kind-of” — absolutely vital to size the dewatering points and pumps.

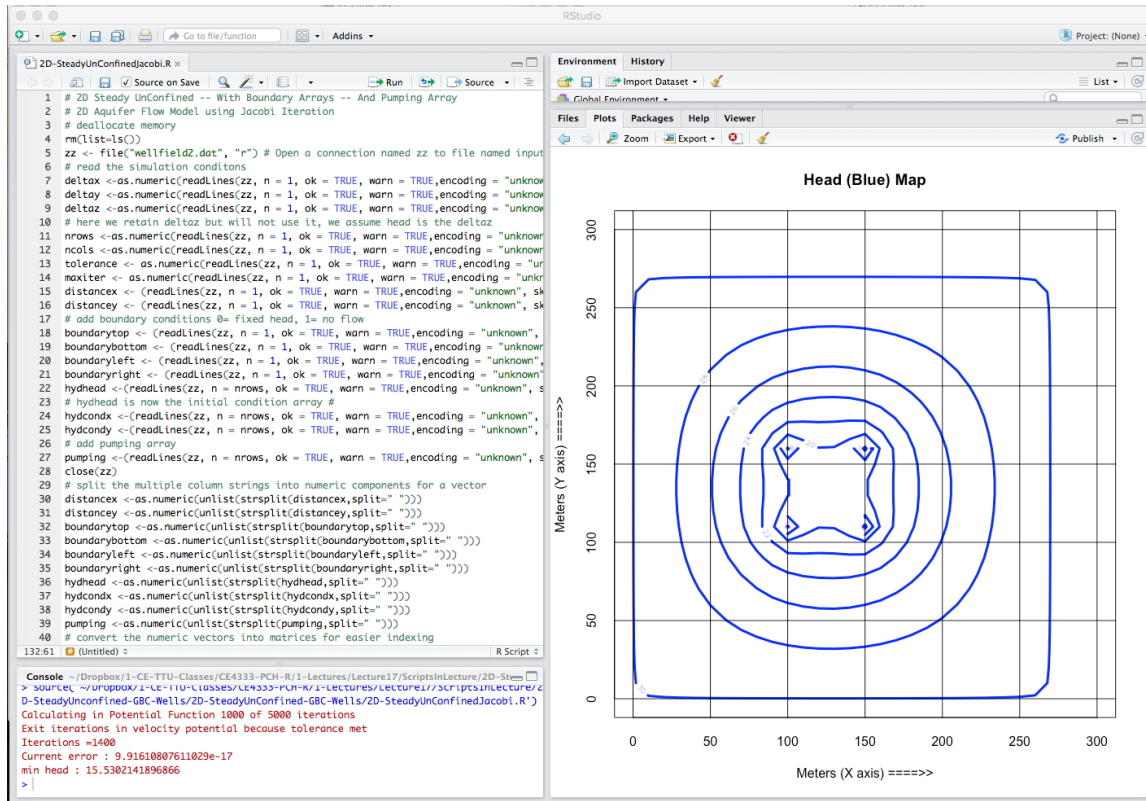


Figure 133. Rectangular unconfined aquifer with 4 wells. Solution and plotting using R.

As a quick summary of this chapter we started with a 1D situation and built a couple of finite-difference models to illustrate the use of `solve(...)` and our homebrew Jacobi iteration method. Next we progressed to a 2D situation (it can either be a slice or a layer) and used the Jacobi iteration method. Using a formal linear solver would also work quite well, but the matrix assembly is a bit tedious – so that was left for future editions of this book. Once we had our 2D computation engine working we extended capability by adding generalized boundary condition capability to the exterior boundaries (the same can be done for interior cells – the matrix assembly logic is awkward, so that too is left for future editions). Once the generalized boundary modification was built, we then added a recharge/pumping capability. This last model is functionally useful (although it needs testing against known solutions to detect any coding errors). To complete the chapter, a simple modification using the variable transmissivity construct allowed us to extend our capability to unconfined aquifers. The next chapter will illustrate unsteady flow where the left hand side of the PDEs is no longer zero.

17.3 Exercises

1. Repeat Example 5 (4 Wells in a Confined Aquifer) with an aquifer thickness of 10 meters and determine:
 - (a) The pumping rate in each well that produces a minimum aquifer head of 15 meters.
 - (b) Produce a contour plot of this solution and report the required pumping rate in cubic meters per day.
 - (c) Plot the head in a elevation/profile view (vertical slice into the aquifer – section a–a in the sketch) that passes through the middle of the aquifer (between the wells) running from left to right.
2. Figure 134 is a plan view map of the hydrologic setting of the OBLEO aquifer. The aquifer is an unconfined aquifer with a single well, P-1, in its geographic center. The aquifer is bounded to the North and South by mountains, and the East and West by two large lakes. The water elevation in each lake is 100.0 meters above some datum, while the ground elevation at the well is 200.0 meters above the same datum.

Rainfall records indicate that recharge to the subsurface is approximately 250 mm/yr. The aquifer is homogeneous and isotropic ($K_x = K_y$) with a hydraulic conductivity value of $K = 100m/day$, and a specific yield $S_y = 0.25$. The example constructs a computer model to help determine the safe yield of the aquifer - assume minimum permissible saturated thickness is 50 m.

- (a) Write/Modify an **R** script to include a recharge array (in addition to a pumping array). Take care that you use the correct sign when adding this array into the computation engine.
- (b) What kind of boundary condition makes sense for the two lakes?
- (c) What kind of boundary condition makes sense for the two mountain ranges?
- (d) Build an input file in the structure used in the examples. Use a grid size of 100km X 100km⁶⁷.
- (e) Compute the head distribution and produce a contour plot with just recharge (no pumping).
- (f) Plot the head in a elevation/profile view (vertical slice into the aquifer) that passes through the well running from West to East. (Lake to Lake).
- (g) Compute the head distribution with recharge **and** pumping at a rate that produces a minimum saturated thickness of about 50 meters.
- (h) Produce a contour plot of this solution and report the required pumping

⁶⁷Yes, i am aware that this aquifer is over 600 miles between the lakes!




Figure 134. Obleo Aquifer Setting (Plan View).

rate in cubic meters per day.

- (i) Plot the head in a elevation/profile view (vertical slice into the aquifer) that passes through the well running from West to East. (Lake to Lake).
- (j) Capture these plots into a brief report where you compare the aquifer conditions before pumping and during pumping.

18 Unsteady Groundwater Flow

18.1 Confined Aquifers

18.1.1 Finite-Difference Methods – 1 Spatial Dimension

Using Equation 203 as a starting point for simulating aquifer behavior, the unsteady flow condition means that the left hand side remains.

$$S\Delta x\Delta y\frac{h_i^{t+\Delta t} - h_i^t}{\Delta t} = (K\Delta y\Delta z\frac{h_{i-1} - h_i}{\Delta x}) - (K\Delta y\Delta z\frac{h_i - h_{i+1}}{\Delta x}) \quad (236)$$

Next we will use the arithmetic mean values of the material properties (K) at the cell interfaces, so the difference equation becomes

$$S\Delta x\Delta y\frac{h_i^{t+\Delta t} - h_i^t}{\Delta t} = (\frac{1}{2}(K_{i-1} + K_i)\Delta y\Delta z\frac{h_{i-1} - h_i}{\Delta x}) - (\frac{1}{2}(K_i + K_{i+1})\Delta y\Delta z\frac{h_i - h_{i+1}}{\Delta x}) \quad (237)$$

Now divide both sides by $\Delta x\Delta y$ to obtain

$$S\frac{h_i^{t+\Delta t} - h_i^t}{\Delta t} = (\frac{1}{2}(K_{i-1} + K_i)\Delta z\frac{h_{i-1} - h_i}{\Delta x^2}) - (\frac{1}{2}(K_i + K_{i+1})\Delta z\frac{h_i - h_{i+1}}{\Delta x^2}) \quad (238)$$

Multiply by $\frac{\Delta t}{S_i}$ to obtain

$$h_i^{t+\Delta t} - h_i^t = \frac{\Delta t}{S_i} \times [(\frac{1}{2}(K_{i-1} + K_i)\Delta z\frac{h_{i-1} - h_i}{\Delta x^2}) - (\frac{1}{2}(K_i + K_{i+1})\Delta z\frac{h_i - h_{i+1}}{\Delta x^2})] \quad (239)$$

Now lets group some constants:

$$A_i = \frac{\Delta t}{2S_i\Delta x^2}(K_{i-1} + K_i)\Delta z \quad (240)$$

$$B_i = \frac{\Delta t}{2S_i\Delta x^2}(K_i + K_{i+1})\Delta z$$

Now substitute into the difference equation.

$$h_i^{t+\Delta t} = A_i(h_{i-1}^*) - (A_i + B_i)(h_i^*) + B_i(h_{i+1}^*) + h_i^t \quad (241)$$

Next we have to decide what time level to evaluate the right hand side terms that have the * time superscript.

18.1.2 Explicit Formulation

If we choose to evaluate at the current time step, the method is called an explicit method. It works fine, but stability requirements force us to use ratios of space and time steps that are often limiting — this behavior is called “conditional stability.”

$$h_i^{t+\Delta t} = A_i(h_{i-1}^t) - (A_i + B_i)(h_i^t) + B_i(h_{i+1}^t) + h_i^t \quad (242)$$

The advantage of this construct is simplicity for programming – in fact one can actually solve these kind of problem graphically if you choose a proper space and time step.

18.1.3 Implicit Formulation

If instead we choose to evaluate at the new time step, the method is called an implicit method.

$$h_i^{t+\Delta t} = A_i(h_{i-1}^{t+\Delta t}) - (A_i + B_i)(h_i^{t+\Delta t}) + B_i(h_{i+1}^{t+\Delta t}) + h_i^t \quad (243)$$

A slight rearrangement of the equation produces a structure very amenable to Jacobi iteration (with all its faults!)

$$0 = A_i(h_{i-1}^{t+\Delta t}) - (A_i + B_i + 1)(h_i^{t+\Delta t}) + B_i(h_{i+1}^{t+\Delta t}) + h_i^t \quad (244)$$

If indeed we decide to use Jacobi iteration, then the system to be solved is

$$h_i^{t+\Delta t} = \frac{A_i(h_{i-1}^{t+\Delta t}) + B_i(h_{i+1}^{t+\Delta t}) + h_i^t}{A_i + B_i + 1} \quad (245)$$

18.2 Finite-Difference Methods – 2 Spatial Dimensions

Rather than build a 1D example, in this section we will just build a 2D example. We will use a fully-implicit method to take advantage of our Jacobi solver we have already built, and incorporate both the boundary condition masks and a pumping array. The result will be a functional tool, that with some further modifications could be used for actual groundwater engineering problems, albeit the computations would seem quite slow. However, a few hours of run time is a lot less costly than drilling wells in the wrong places — for a professional problem we would probably use MODFLOW which has far more efficient solvers (but interestingly works about the same).⁶⁸

⁶⁸At the conclusion of the example, we will compare the results to MODFLOW and demonstrate that the two tools produce the same answers. Generally we won't often have to build our own codes – however if we do, they should produce the same results as a professional code for the same problem. Then we would use our homebuilt code for whatever special need we have that the professional code cannot conveniently address.

Starting with

$$\begin{aligned}
 S \frac{\partial h_i}{\partial t} = & \left[\frac{1}{\Delta x} T_x \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j-1} - h_{i,j}}{\Delta y} \right] - \\
 & \left[\frac{1}{\Delta x} T_x \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j} - h_{i,j+1}}{\Delta y} \right] \\
 & + \frac{R_{i,j}}{\Delta x \Delta y} - \frac{Q_{i,j}}{\Delta x \Delta y}
 \end{aligned} \tag{246}$$

We apply the finite difference approach to the time derivative and obtain

$$\begin{aligned}
 S \frac{h_i^{t+\Delta t} - h_i^t}{\Delta t} = & \left[\frac{1}{\Delta x} T_x \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j-1} - h_{i,j}}{\Delta y} \right] - \\
 & \left[\frac{1}{\Delta x} T_x \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j} - h_{i,j+1}}{\Delta y} \right] \\
 & + \frac{R_{i,j}}{\Delta x \Delta y} - \frac{Q_{i,j}}{\Delta x \Delta y}
 \end{aligned} \tag{247}$$

Proceeding as above in the 1D case, we will isolate the time levels, multiply through by Δt , and divide by S we obtain

$$\begin{aligned}
 h_i^{t+\Delta t} = h_i^t + \frac{\Delta t}{S_i} & \left[\frac{1}{\Delta x} T_x \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j-1} - h_{i,j}}{\Delta y} \right] - \\
 \frac{\Delta t}{S_i} & \left[\frac{1}{\Delta x} T_x \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j} - h_{i,j+1}}{\Delta y} \right] \\
 & + \frac{\Delta t}{S_i} \frac{R_{i,j}}{\Delta x \Delta y} - \frac{\Delta t}{S_i} \frac{Q_{i,j}}{\Delta x \Delta y}
 \end{aligned} \tag{248}$$

Next if we choose the time level on the right-hand side as t , we would have an explicit scheme⁶⁹, however if we choose to evaluate at the new time level $t + \Delta t$, we have an implicit scheme. We are going to choose an implicit scheme to leverage our work on the Jacobi solver we have already built so the difference equation becomes

$$\begin{aligned}
 h_i^{t+\Delta t} = h_i^t + \frac{\Delta t}{S_i} & \left[\frac{1}{\Delta x} T_x \frac{h_{i-1,j}^{t+\Delta t} - h_{i,j}^{t+\Delta t}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j-1}^{t+\Delta t} - h_{i,j}^{t+\Delta t}}{\Delta y} \right] - \\
 \frac{\Delta t}{S_i} & \left[\frac{1}{\Delta x} T_x \frac{h_{i,j}^{t+\Delta t} - h_{i+1,j}^{t+\Delta t}}{\Delta x} + \frac{1}{\Delta y} T_y \frac{h_{i,j}^{t+\Delta t} - h_{i,j+1}^{t+\Delta t}}{\Delta y} \right] \\
 & + \frac{\Delta t}{S_i} \frac{R_{i,j}^t}{\Delta x \Delta y} - \frac{\Delta t}{S_i} \frac{Q_{i,j}^t}{\Delta x \Delta y}
 \end{aligned} \tag{249}$$

Equation 249 is the difference equation we will code into our **R** script with some modifications to allow time stepping and plotting results. I set the pumping and recharge at time level t , but one could just as well choose to average over two time levels. As

⁶⁹Easy to program, but conditionally stable and usually requiring a very small time step for stability. Here we will use an implicit scheme, which will be unconditionally stable – we will still use pretty small time steps for accuracy, but not need to worry too much about stability.

long as the time steps are small relative to changes in pumping and recharge⁷⁰, it won't make much practical difference.

Observe a very important part of the difference equation, h_i^t is known at the beginning of a time step and does not change during the computation effort for that time period – it is just another constant like recharge or pumping. However, once we have the update, then it is changed and the next time step is begun. The observation is important to remember because we are simply going to wrap a time stepping loop around our already tested and working Jacobi solver.

Employing the same A, B, C, D structure as the unsteady 2D solvers we built the update model is

Also as in the one-dimensional case, we will approximate the spatial variation of the material properties (transmissivity) as arithmetic mean values between two cells, so making the following definitions:

$$\begin{aligned}
 A_{i,j} &= \frac{\Delta t}{2S_i \Delta x^2} (T_{x,(i-1,j)} + T_{x,(i,j)}) \\
 B_{i,j} &= \frac{\Delta t}{2S_i \Delta x^2} (T_{x,(i,j)} + T_{x,(i+1,j)}) \\
 C_{i,j} &= \frac{\Delta t}{2S_i \Delta y^2} (T_{y,(i,j-1)} + T_{y,(i,j)}) \\
 D_{i,j} &= \frac{\Delta t}{2S_i \Delta y^2} (T_{y,(i,j)} + T_{y,(i,j+1)}) \\
 E_{i,j} &= \frac{\Delta t}{S_i \Delta x \Delta y}
 \end{aligned} \tag{250}$$

Substitution into the difference equation and writing the cell equation for $h_{i,j}$ yields

$$h_{i,j}^{t+\Delta t} = \frac{[A_{i,j}h_{i-1,j}^{t+\Delta t} + B_{i,j}h_{i+1,j}^{t+\Delta t} + C_{i,j}h_{i,j-1}^{t+\Delta t} + D_{i,j}h_{i,j+1}^{t+\Delta t}]}{[A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j} + 1]} + h_i^t + E_{i,j}[R_{i,j}^t - Q_{i,j}^t] \tag{251}$$

This equation is quite amenable to the Jacobi iteration technique, with the addition of the two left-most terms for head at the start of the time step and the net recharge less pumping. Notice the +1 in the denominator of the weighting term – it reflects that $h_{i,j}^{t+\Delta t}$ appears on both the left and right hand side of Equation 249.

Listing 71 is an **R** script that implements these changes. In the script the plotting has been moved into the prototype function, so we can generate a plot when we wish. The input data reading portion has a line added to capture the time step length, the maximum number of time steps to take, a printing frequency (make a plot every `iprint` time steps), and a storage coefficient array. Observe that the printing test makes use of the `MODULO` construct in **R**.

⁷⁰Stress periods on the order of months for time steps in days

The actual script use will be demonstrated by the next example.

Listing 71. R code demonstrating an Aquifer Flow Simulator for 2D Unsteady Confined Aquifer Flow. This code fragment implements the Jacobi iteration to solve the linear system at **each time step**. A plotting prototype function is used so we can plot to a graphics device at specified intervals (multiples of the total number of time steps). A graphics device is used (rather than plotting to the **R Studio** environment) so multiple plots at different simulation times can be rendered.

```
# 2D Unsteady Confined -- With Boundary Arrays and Wells using Jacobi Iteration
# deallocate memory
rm(list=ls())
#####
# prototype plotting function
#####
plotnow<-function(headnow,ncols,nrows,etime,deltat){
###    built position arrays for contour plotting
velocity_plt<-matrix(0,ncols,nrows)
for( i in 1:nrows){
  for( j in 1:ncols){
    velocity_plt[j,i]=headnow[i,j]
  }
}
###    contour plot of head -- note axes are rotated
#debug  message("elapsed time : ",maxtime * deltat)
mytitle <- paste("Head (Blue) Map at Time = :",etime*deltat, " seconds")
contour(distancex,distancey,velocity_plt,
        plot.title = title(main = mytitle,
                           xlab = "Meters (X axis) ===>>",
                           ylab = "Meters (Y axis) ===>>"),
        col="blue",lwd=3,nlevels=20)
}
#####
#####
zz <- file("input2.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
deltat <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
nrows <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
ncols <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
maxiter <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
maxtime <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
iprint <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
distancex <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
distancey <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
# add boundary conditions 0= fixed head, 1= no flow
boundarytop <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
boundarybottom <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
boundaryleft <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
boundaryright <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
hydhead <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
# hydhead is now the initial condition array #
hydcondx <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondy <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
storage <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
pumping <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
distancex <-as.numeric(unlist(strsplit(distancex,split=" ")))
distancey <-as.numeric(unlist(strsplit(distancey,split=" ")))
boundarytop <-as.numeric(unlist(strsplit(boundarytop,split=" ")))
boundarybottom <-as.numeric(unlist(strsplit(boundarybottom,split=" ")))
boundaryleft <-as.numeric(unlist(strsplit(boundaryleft,split=" ")))
boundaryright <-as.numeric(unlist(strsplit(boundaryright,split=" ")))
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
```



```

hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
storage <-as.numeric(unlist(strsplit(storage,split=" ")))
pumping <-as.numeric(unlist(strsplit(pumping,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <- matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
storage <-matrix(storage,nrow=nrows,ncol=ncols,byrow = TRUE)
pumping <-matrix(pumping,nrow=nrows,ncol=ncols,byrow = TRUE)
# allocate a graphics device for sequential plotting
pdf("2d-unsteady-confined-junk.plot.pdf") # send any plots to file
# here we perform the velocity potential calculations
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
qrat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltat*deltaz)
      /(2.0*storage[irow,jcol]*deltax^2)
    bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltat*deltaz)
      /(2.0*storage[irow,jcol]*deltax^2)
    cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltat*deltaz)
      /(2.0*storage[irow,jcol]*deltay^2)
    dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltat*deltaz)
      /(2.0*storage[irow,jcol]*deltay^2)
    qrat[irow,jcol]<-(-1.0)*deltat*pumping[irow,jcol]/(deltax*deltay*storage[irow,jcol])
  }
}
##### initial conditions #####
headnow <- hydhead
for (tstep in 1:maxtime){ ## this is the outer time-stepping loop
##### use jacobi to solve for a single time step #####
hydhead <- headnow
headold <- hydhead # copy the head array, used to test for stopping
tolflag <- FALSE
for (iter in 1:maxiter){
# Boundary Conditions
# Top Row
for(jcol in 1:ncols){
  if(boundarytop[jcol] == 0){hydhead[1,jcol]<-hydhead[2,jcol]} #no-flow at top
  if(boundarybottom[jcol] == 0){hydhead[nrows,jcol]<-hydhead[2,jcol]} #no-flow at bottom
  # otherwise values are fixed head
}
for(irow in 1:nrows){
  if(boundaryleft[irow] == 0){hydhead[irow,1]<-hydhead[irow,2]} #no-flow at left
  if(boundaryright[irow] == 0){hydhead[irow,ncols]<-hydhead[irow,ncols-1]} #no-flow at
  right
}
  for (irow in 2:(nrows-1)){
    for (jcol in 2:(ncols-1)){
      hydhead[irow,jcol] <- ( headnow[irow,jcol] +
        qrat[irow,jcol] +
        amat[irow,jcol]*hydhead[irow-1,jcol ] +
        bmat[irow,jcol]*hydhead[irow+1,jcol ] +
        cmat[irow,jcol]*hydhead[irow ,jcol-1] +
        dmat[irow,jcol]*hydhead[irow ,jcol+1] )/
      (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol]+1)
    }
  }
# test for stopping iterations
percentdiff <- sum((hydhead-headold)^2)
if (percentdiff < tolerance){
  message("Exit iterations in velocity potential because tolerance met");
  message("Iterations =", iter);
  message("Current error : ",percentdiff);
  tolflag <- TRUE
  break
}
  headold<-hydhead #update the current head vector
# if( iter %% 1000 == 0){message("Calculating in Potential Function ",iter," of ",maxiter,
  " iterations")}
}
if (tolflag == FALSE){
# message("Exit iterations in potential function at max iterations")
  message("Current error : ",percentdiff)
}
# perform an update
headnow <- hydhead
##### end of a time step #####
# check for plotnow
if(tstep%%iprint==0){plotnow(headnow,ncols,nrows,tstep,deltat)}
}
# exit the time-stepping loop
# output ending conditons
write(t(hydhead), file='2d-unsteady-confined.txt',ncolumns = ncols,sep=",")
# deallocate (disconnect) graphics device
dev.off() # disconnect the pdf file

```

Example 1: Rectangular Aquifer with Pumping – Unsteady Flow Figure 135 is a plan view of a confined aquifer with a Well field located as shown, near the Red River. The head in the river is 1000 meters as is the head at the South boundary, the Green Swamp. The aquifer thickness is about 10 meters. The hydraulic conductivity of the water producing zone is 20 meters per day. Determine the head distribution every 100 days of pumping at a rate of 20,000 cubic meters per day for 1000 days (about 3 years). The storage coefficient of the aquifer is $S = 0.01$.




Figure 135. Rectangular aquifer with wellfield. North and South boundaries are fixed head. East and West boundaries are Zero-Flux.

Each cell in the sketch is 1000X1000 meters.

Listing 72 is the input file for the example problem. Literally it only has two new “lines.”

Listing 72. Input file for the Hubbleville Aquifer Example. Annotations in the listing need to be removed for actual running of the file and are included to illustrate the various parts of the input data.

```

1000 <== Delta x
1000 <== Delta y
10 <== Delta z (thickness)
0.1 <== Time step length (in days)
10 <== row count
10 <== column count
1e-8 <== tolerance for linear solver
1000 <== how many trials per linear system
10000 <== how many time steps
1000 <== print every ## time step
0 1000 2000 3000 4000 5000 6000 7000 8000 9000
0 1000 2000 3000 4000 5000 6000 7000 8000 9000
1 1 1 1 1 1 1 1 1 1 <== boundary conditions
1 1 1 1 1 1 1 1 1 1
0 0 0 0 0 0 0 0 0 0

```

```

0 0 0 0 0 0 0 0 0 0
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000 <== initial head array
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
1000 1000 1000 1000 1000 1000 1000 1000 1000 1000
20 20 20 20 20 20 20 20 20 20 <== Kx array
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
20 20 20 20 20 20 20 20 20 20
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 <== Storage coefficient
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.11
0 0 0 0 0 0 0 0 0 0 <== pumping array
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
0 0 0 0 20000 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0

```

Figures 136 and 137 are plots of the head distribution after 300 and 900 days (in the first year and the third year of pumping). In the plots one can observe that the cone of depression is growing (outward from the well field), and the numerical values of the contours are reflecting the progressively declining head elevations.

At some point the head distribution would stop changing (and equilibrium solution) and would agree with a steady flow simulation for the same aquifer conditions. If one were to perform a “cut-and-fill” analysis of the contour maps, the amount of aquifer dewatered (difference in head multiplied by the storage coefficient, integrated over the entire area shown) would equal the volume pumped from the system.⁷¹

The next section will extend the concept to an unconfined aquifer – again using the variable transmissivity approach of the prior chapter.

⁷¹Numerically very close – there might be some differences because of arithmetic precision issues.




Figure 136. Plot of head after 300 days of pumping.




Figure 137. Plot of head after 900 days of pumping.

18.3 Unconfined Aquifers

As in the previous chapter, we will start with the unconfined flow equation in 2 dimensions

$$S_y \Delta x \Delta y \frac{\partial h_i}{\partial t} = [K_x \Delta y \bar{h} \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + K_y \Delta x \bar{h} \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] - [K_x \Delta y \bar{h} \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + K_y \Delta x \bar{h} \frac{h_{i,j} - h_{i,j+1}}{\Delta y}] \quad (252)$$

Next divide by the cell plan view area $\Delta x \Delta y$ to obtain a more compact form of the difference equation.

$$S \frac{\partial h_i}{\partial t} = [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] - [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j} - h_{i,j+1}}{\Delta y}] \quad (253)$$

Now finite-difference the left hand side and add pumping and recharge

$$S \frac{h_i^{t+\Delta t} - h_i^t}{\Delta t} = [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] - [\frac{1}{\Delta x} K_x \bar{h} \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j} - h_{i,j+1}}{\Delta y}] + \frac{R_{i,j}}{\Delta x \Delta y} - \frac{Q_{i,j}}{\Delta x \Delta y} \quad (254)$$

Multiply through by Δt and divide by S_y (specific yield – analogous to storage coefficient).

$$h_i^{t+\Delta t} = h_i^t + [\frac{\Delta t}{S_{y,i}} \frac{1}{\Delta x} K_x \bar{h} \frac{h_{i-1,j} - h_{i,j}}{\Delta x} + \frac{\Delta t}{S_{y,i}} \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j-1} - h_{i,j}}{\Delta y}] - [\frac{\Delta t}{S_{y,i}} \frac{1}{\Delta x} K_x \bar{h} \frac{h_{i,j} - h_{i+1,j}}{\Delta x} + \frac{\Delta t}{S_{y,i}} \frac{1}{\Delta y} K_y \bar{h} \frac{h_{i,j} - h_{i,j+1}}{\Delta y}] + \frac{\Delta t}{S_{y,i}} \frac{R_{i,j}}{\Delta x \Delta y} - \frac{\Delta t}{S_{y,i}} \frac{Q_{i,j}}{\Delta x \Delta y} \quad (255)$$

Next we will approximate the spatial variation of the $K \bar{h}$ as arithmetic mean values between two cells, so making the following definitions:

$$\begin{aligned}
A_{i,j} &= \frac{\Delta t}{2S_{y,i} \Delta x^2} (K_{x,(i-1,j)} h_{i-1,j} + K_{x,(i,j)} h_{i,j}) \\
B_{i,j} &= \frac{\Delta t}{2S_{y,i} \Delta x^2} (K_{x,(i,j)} h_{i,j} + K_{x,(i+1,j)} h_{i+1,j}) \\
C_{i,j} &= \frac{\Delta t}{2S_{y,i} \Delta y^2} (K_{y,(i,j-1)} h_{i,j-1} + K_{y,(i,j)} h_{i,j}) \\
D_{i,j} &= \frac{\Delta t}{2S_{y,i} \Delta y^2} (K_{y,(i,j)} h_{i,j} + K_{y,(i,j+1)} h_{i,j+1}) \\
E_{i,j} &= \frac{\Delta t}{S_{y,i} \Delta x \Delta y}
\end{aligned} \tag{256}$$

Collect the terms

$$\begin{aligned}
h_i^{t+\Delta t} &= h_i^t + \\
A_{i,j} h_{i-1,j}^{t*} + B_{i,j} h_{i+1,j}^{t*} - (A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j}) h_{i,j}^{t*} + \\
C_{i,j} h_{i,j-1}^{t*} + D_{i,j} h_{i,j+1}^{t*} + E_{i,j} [R_{i,j}^{t*} - Q_{i,j}^{t*}]
\end{aligned} \tag{257}$$

When we eventually decide whether to use explicit or implicit we will specify the time levels (t^*). As a practical matter the computation of the A, B, C, D terms will have to be at a known or intermediate time level so it is not superscripted with a time level index – although it is clear in the definition of the terms that a time level is needed to perform the requisite multiplications.

18.3.1 Explicit Formulation

The explicit formulation would specify the time levels on the right hand side as t . The resulting difference equation is

$$\begin{aligned}
h_i^{t+\Delta t} &= h_i^t + \\
A_{i,j} h_{i-1,j}^t + B_{i,j} h_{i+1,j}^t - (A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j}) h_{i,j}^t + \\
C_{i,j} h_{i,j-1}^t + D_{i,j} h_{i,j+1}^t + E_{i,j} [R_{i,j}^t - Q_{i,j}^t]
\end{aligned} \tag{258}$$

Everything on the right-hand side is known, so other than being conditionally stable (and using a small time step) we could start from known conditions and evolve the solution forward in time. There is non-linearity imbedded in the A, B, C, D terms, but they are evaluated at known time levels. As with the other unsteady case, the required time steps are small and dependent on the spatial and time step sizes, as well as the material properties. Unlike the previous case, the stability is also dependent on the solution itself and hence unpredictable in advance (although we can make very good guesses).

18.3.2 Implicit Formulation

The implicit formulation would specify the time levels on the right hand side as $t + \Delta t$. We would also have to decide how to handle the recharge and pumping – either an average over two adjacent times, or at the old (or new) time level. Here we choose the old time level, mostly for simplicity in building the **R** script.

The resulting difference equation is

$$h_i^{t+\Delta t} = h_i^t + A_{i,j}h_{i-1,j}^{t+\Delta t} + B_{i,j}h_{i+1,j}^{t+\Delta t} - (A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j})h_{i,j}^{t+\Delta t} + C_{i,j}h_{i,j-1}^{t+\Delta t} + D_{i,j}h_{i,j+1}^{t+\Delta t} + E_{i,j}[R_{i,j}^t - Q_{i,j}^t] \quad (259)$$

Observe that $h_i^{t+\Delta t}$ appears on both sides of the equation, so some additional manipulation is needed, first collect all the $h_i^{t+\Delta t}$ terms on the left side

$$(A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j} + 1)h_i^{t+\Delta t} = h_i^t + A_{i,j}h_{i-1,j}^{t+\Delta t} + B_{i,j}h_{i+1,j}^{t+\Delta t} + C_{i,j}h_{i,j-1}^{t+\Delta t} + D_{i,j}h_{i,j+1}^{t+\Delta t} + E_{i,j}[R_{i,j}^t - Q_{i,j}^t] \quad (260)$$

Then divide both sides by the constant to get an update formula

$$h_i^{t+\Delta t} = \frac{h_i^t + A_{i,j}h_{i-1,j}^{t+\Delta t} + B_{i,j}h_{i+1,j}^{t+\Delta t} + C_{i,j}h_{i,j-1}^{t+\Delta t} + D_{i,j}h_{i,j+1}^{t+\Delta t} + E_{i,j}[R_{i,j}^t - Q_{i,j}^t]}{(A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j} + 1)} \quad (261)$$

The formula above is now structured so we can use our existing Jacobi iteration solver we developed in the previous chapter. As we did for the steady case, we will simply move where the A, B, C, D terms are computed.⁷²

Listing 73 is a listing that implements the revised finite-difference equation structure for the unconfined flow equation. The new listing literally involved moving the portion of the steady, unconfined flow code that creates the variable transmissivity arrays to just before the assignment of boundary conditions in the unsteady confined code, and removing the existing instance of the unsteady confined code portion that created the A, B, C, D portions.

The script was then tested (just to see if it runs) with the same Hubbleville input file. In the revised version, Δz is still read into the program, but not actually used. Again this is an act of laziness as well as code re-usability.

⁷²They will go inside the Jacobi solver loop. We will pay a steep computational cost, but still be able to use our existing toolkit, handle the non-linearity by lagging the material properties a single iteration step, and still have an unconditionally stable solution method.

Listing 73. Listing for R implementation for unsteady unconfined flow using Jacobi iteration.

```

# 2D Unsteady Unconfined -- With Boundary Arrays and Wells
# 2D Aquifer Flow Model using Jacobi Iteration
# deallocate memory
rm(list=ls())
#####
# prototype plotting function
#####
plotnow<-function(headnow,ncols,nrows,etime,deltat){
###   built position arrays for contour plotting
velocity_plt<-matrix(0,ncols,nrows)
for( i in 1:nrows){
  for( j in 1:ncols){
    velocity_plt[j,i]=headnow[i,j]
  }
}
###   contour plot of head -- note axes are rotated
#debug  message("elapsed time : ",maxtime * deltat)
mytitle <- paste("Head (Blue) Map at Time = :",etime*deltat," time units")
contour(distancex,distancey,velocity_plt,
        plot.title = title(main = mytitle,
                           xlab = "Meters (X axis) ==>>>",
                           ylab = "Meters (Y axis) ==>>>"),
        col="blue",lwd=3,nlevels=10)
}
#####
#####
zz <- file("hubbleville.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
deltat <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
nrows <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
ncols <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
skipNul = FALSE))
maxiter <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
maxtime <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
iprint <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                                 skipNul = FALSE))
distancex <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
distancey <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
# add boundary conditions 0= fixed head, 1= no flow
boundarytop <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
boundarybottom <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
                             skipNul = FALSE))
boundaryleft <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
boundaryright <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul
= FALSE))
hydhead <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
# hydhead is now the initial condition array #
hydcondx <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondy <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
storage <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
pumping <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
distancex <-as.numeric(unlist(strsplit(distancex,split=" ")))
distancey <-as.numeric(unlist(strsplit(distancey,split=" ")))
boundarytop <-as.numeric(unlist(strsplit(boundarytop,split=" ")))
boundarybottom <-as.numeric(unlist(strsplit(boundarybottom,split=" ")))
boundaryleft <-as.numeric(unlist(strsplit(boundaryleft,split=" ")))
boundaryright <-as.numeric(unlist(strsplit(boundaryright,split=" ")))
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
storage <-as.numeric(unlist(strsplit(storage,split=" ")))
pumping <-as.numeric(unlist(strsplit(pumping,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <- matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
storage <-matrix(storage,nrow=nrows,ncol=ncols,byrow = TRUE)

```



```

pumping <-matrix(pumping,nrow=nrows,ncol=ncols,byrow = TRUE)
# allocate a graphics device for sequential plotting
pdf("2d-unsteady-confined-junk.plot.pdf") # send any plots to file
# here we perform the velocity potential calculations
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
qr<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    # amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltat*
    #   deltax)/(2.0*storage[irow,jcol]*deltax^2)
    # bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltat*deltax)
    #   /(2.0*storage[irow,jcol]*deltax^2)
    # cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltat*deltax)
    #   /(2.0*storage[irow,jcol]*deltax^2)
    # dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltat*deltax)
    #   /(2.0*storage[irow,jcol]*deltax^2)
    qr[irow,jcol]<-(-1.0)*deltat*pumping[irow,jcol]/(deltax*deltax*storage[irow,jcol])
  }
}
# for(irow in 2:(nrows-1)){
#   for(jcol in 2:(ncols-1)){
#     amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltax)/(2.0*
#       deltax^2)
#     bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltax)/(2.0*
#       deltax^2)
#     cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltax)/(2.0*
#       deltax^2)
#     dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltax)/(2.0*
#       deltax^2)
#   }
# }
##### initial conditions #####
print(hydhead)
#maxtime <- 1
headnow <- hydhead
for(tstep in 1:maxtime){
##### use jacobi to solve for a single time step #####
hydhead <- headnow
headold <- hydhead # copy the head array, used to test for stopping
tolflag <- FALSE
for(iter in 1:maxiter){
# built the variable transmissivity arrays for current time-step
  for(irow in 2:(nrows-1)){
    for(jcol in 2:(ncols-1)){
      amat[irow,jcol]<-((hydcondx[irow-1,jcol ]*hydhead[irow-1,jcol ]
        +hydcondx[irow ,jcol ]*hydhead[irow ,jcol ]))/(2.0*deltax^2)

      bmat[irow,jcol]<-((hydcondx[irow ,jcol ]*hydhead[irow ,jcol ]
        +hydcondx[irow+1,jcol ]*hydhead[irow+1,jcol ]))/(2.0*deltax^2)

      cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]*hydhead[irow ,jcol-1]
        +hydcondy[irow ,jcol ]*hydhead[irow ,jcol ]))/(2.0*deltax^2)

      dmat[irow,jcol]<-((hydcondy[irow ,jcol ]*hydhead[irow ,jcol ]
        +hydcondy[irow ,jcol+1]*hydhead[irow ,jcol+1]))/(2.0*deltax^2)
    }
  }
# set the boundary conditions
for(jcol in 1:ncols){ #top and bottom rows
  if(boundarytop[jcol] == 0){hydhead[1,jcol]<-hydhead[2,jcol]} #no-flow at top
  if(boundarybottom[jcol] == 0){hydhead[nrows,jcol]<-hydhead[2,jcol]} #no-flow at bottom
  # otherwise values are fixed head
}
for(irow in 1:nrows){ #left and right columns
  if(boundaryleft[irow] == 0){hydhead[irow,1]<-hydhead[irow,2]} #no-flow at left
  if(boundaryright[irow] == 0){hydhead[irow,ncols]<-hydhead[irow,ncols-1]} #no-flow at
    right
  # otherwise values are fixed head
}
  for (irow in 2:(nrows-1)){
    for (jcol in 2:(ncols-1)){
      hydhead[irow,jcol] <- ( headnow[irow,jcol] +
        qr[irow,jcol] +
        amat[irow,jcol]*hydhead[irow-1,jcol ] +
        bmat[irow,jcol]*hydhead[irow+1,jcol ] +
        cmat[irow,jcol]*hydhead[irow ,jcol-1] +
        dmat[irow,jcol]*hydhead[irow ,jcol+1] )/
      (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol]+1)
    }
  }
# test for stopping iterations
percentdiff <- sum((hydhead-headold)^2)
if (percentdiff < tolerance){
  message("Exit iterations in velocity potential because tolerance met");
  message("Iterations =", iter);
  message("Current error : ",percentdiff);
  tolflag <- TRUE
}

```

```

    break
  }
  headold<-hydhead #update the current head vector
# if( iter %% 1000 == 0){message("Calculating in Potential Function ",iter," of ",maxiter,
  " iterations")}
}
if (tolflag == FALSE){
# message("Exit iterations in potential function at max iterations")
  message("Current error : ",percentdiff)
}
# perform an update
headnow <- hydhead
##### end of a time step #####
#print(headnow)
# check for plotnow
if(tstep%%iprint==0){plotnow(headnow,ncols,nrows,tstep,deltat)}
}
# output ending conditons
write(t(hydhead), file='2d-unsteady-confined.txt',ncolumns = ncols,sep=",")
# deallocate (disconnect) graphics device
dev.off() # disconect the pdf file

```

Example 2: Unconfined Rectangular Aquifer with Pumping – Unsteady Flow

Figure ?? is a plan view of an unconfined aquifer with a Well field located as shown, near the Red River. The head in the river is 1000 meters as is the head at the South boundary, the Green Swamp. The well field pumps 20,000 cubic meters per day. The net recharge to the aquifer is 0.001 meters/day. The hydraulic conductivity of the aquifer is 50 meters per day. The hydraulic conductivity of the swamp is about 500 meters per day. The specific yield of the aquifer is 0.35. The town of Hubbleville is at elevation 1020 meters. The managers of the Green Swamp Conservation area are concerned that pumping will significantly reduce groundwater discharge to the swamp and threaten wildlife habitat. The town claims that the well is located sufficiently close to the Red River so that induced recharge will contribute the significant portion of water that flows to the well and none will come from the swamp. Each cell in the sketch is 1000X1000 meters.

Use the unsteady, unconfined groundwater flow script to

1. Demonstrate that flow is towards the swamp with the well field inactive (zero pumping, but non-zero recharge).
2. Determine if the flow is still towards the swamp with the well field activated (non-zero pumping, non-zero recharge).
3. Assuming that water is eventually drawn from the swamp, determine how many years the well field can be operated, before water is drawn from the swamp into the well, or until the well de-waters (head drops below zero).

Answering these three questions is a matter of running our script three times, with two input files. The first input file would be the problem conditions with the well off, but with recharge on. Given that our prototype script does not have a recharge array, we have two choices — add one and refactor the requisite parts of the script, or observe that recharge is just negative pumping and simply change the pumping array. I chose that second approach, the first is left as an exercise. So the recharge for any cell in the correct problem units is

$$R_{i,j} = rate \times \Delta x \times \Delta y = 0.001 \text{ m/day} \times 1000 \text{ m} \times 1000 \text{ m} = 1000 \text{ m}^3/\text{day} \quad (262)$$




Figure 138. Hubbleville Aquifer System.

The first question is asking about a “pre-development” condition, so we need to run our script once until the solution stops changing (an equilibrium solution) then use the head distribution for this solution as the initial condition for the second run, when we turn on the pump. An important observation is that the aquifer bottom elevation is at 980 meters, so the hydraulic head in the swamp or at the river is only 20 meters based on the profile sketch. Thus the maximum head in the aquifer should be less than 1020 meters (otherwise the whole area is a swamp) and greater than zero. Once the computations are completed we can add back the datum (980 meters) to recover the various water elevations if we wish. For the example, I will just stipulate the desired head is between zero and 40 meters.

Listing 74 is the input file for the pre-development conditions. The file was created by running the script once with all the heads set to 20. Then at the end of the run (30 years), the output heads were pasted back into the input file, and it was rerun. This process was repeated several times until the program ran through all the time steps using only a single iteration and the error term essentially reported as a small constant value. At this condition, the system was deemed to be at equilibrium. The listing shows the result that was used to produce Figure 139.

Examination of Figure 139 shows that the contour lines are parallel to the swamp and the river (bottom and top) and slope from the middle ($Y = 5500$ meters) towards each boundary. The highest plotted water elevation is 34 meters, which would be a depth of about 6 meters beneath the town of Hubbleville. The high “ridge” is called a groundwater divide. All water South of the divide flows to the swamp, whereas all water North of the divide flows to the river.

Listing 74. Input file for the Hubleville problem – but as an unconfined aquifer – model is run to equilibrium to establish the pre-development conditions. Here we choose to use 30 years simulation time as sufficiently long and printed contour maps at 5 year intervals to assess if there is much change as the solution proceeds.

```

1000
1000
10
1
10
10
1e-8
1000
10950
1825
0 1000 2000 3000 4000 5000 6000 7000 8000 9000
0 1000 2000 3000 4000 5000 6000 7000 8000 9000
1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
20 20 20 20 20 20 20 20 20 20
21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491
23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453
29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826
32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564
34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182
33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271
31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183
27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346
20 20 20 20 20 20 20 20 20 20
500 500 500 500 500 500 500 500 500 500
500 500 500 500 500 500 500 500 500 500
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
500 500 500 500 500 500 500 500 500 500
500 500 500 500 500 500 500 500 500 500
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000

```

Now we will activate the pumping. The pumping rate is positive 20,000 cubic meters per day. The recharge rate is a negative 1,000 cubic meters per day, so net pumping is 19,000 cubic meters per day. We will change the contents of the single well field cell, rename the file, and rerun the script.

Figure 140 is the result of the change. The plot is promising because the slope of the water is still towards the swamp. However, if we examine the plot at 30 years we will discover we have dewatered the aquifer in the wellfield (at some point the wells would




Figure 139. Hubbleville Aquifer head distribution at equilibrium.

fail). So we can conclude that the swamp is safe, but the town of Hubbleville cannot realistically recover the amount of water they desire.

While we have the model available, we can try different pumping rates to see how much firm yield the town could expect. For the example, let's stipulate that the smallest allowable head in the well field is 5 meters. Our goal is now to determine how large a pumping rate we can use and keep the head at this location greater than 5 meters. The exercise is quite simple, we will change the pumping value and run the script and examine the output. If the head is too small, decrease the pumping, if the head is too large then increase the pumping and we stop when it is just right.⁷³

⁷³Goldilocks optimization!




Figure 140. Hubbleville Aquifer head distribution at 5 years after well field is active.

Listing 75 is the input file used for the post development case with a pumping rate set to 18,000 cubic meters per day. The simulation forecast that the head in the well field cell would be 6.55 meters after 60 years of pumping. The stopping criteria for each time step was still changing, hence we have not quite reached equilibrium – however it would be safe to conclude that a firm yield of 18,000 cubic meters per day is achievable and does not impact the swamp as was feared in the original problem statement.

Figure 141 is the contour plot of the head distribution under these conditions. The induced recharge from the river is evident by the contour that extends to the boundary at the top. The swamp is South of the saddle point in the water table plot and water south of this point flows towards the swamp (as desired).

Listing 75. Input file for the Hubleville problem – 60 year simulation. Pumping rate is 18000 cubic meters per day.

```

1000
1000
10
1
10
10
1e-8
1000
21900
2190
0 1000 2000 3000 4000 5000 6000 7000 8000 9000
0 1000 2000 3000 4000 5000 6000 7000 8000 9000
1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1
0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
20 20 20 20 20 20 20 20 20 20
21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491 21.33491
23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453 23.18453
29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826 29.33826
32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564 32.70564
34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182 34.12182
33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271 33.83271
31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183 31.79183
27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346 27.61346
20 20 20 20 20 20 20 20 20 20
500 500 500 500 500 500 500 500 500 500
500 500 500 500 500 500 500 500 500 500
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
500 500 500 500 500 500 500 500 500 500
500 500 500 500 500 500 500 500 500 500
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
50 50 50 50 50 50 50 50 50 50
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35 0.35
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000
-1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000 -1000

```

These three chapters introduced an important computational hydraulics problem – namely that of flow in a porous medium. The governing equation of motion is Darcy’s law. We computed the heads using a variety of methods, but mostly focused on Jacobi iteration because for these kinds of problems it is a stable solution technique. It is not particularly computationally efficient, but it works.

We choose not to even do explicit methods – not because they are bad, but because they are conditionally stable which can be an undesired problem.⁷⁴ We also focused

⁷⁴Explicit methods are easy to program, and if I were tasked with simulating an explosion (a chemical reaction that moves outward from ignition at or above the speed of sound) or some other process




Figure 141. Hubbleville Aquifer head distribution at 60 years after well field is active with pumping rate of 18,000 meters per day.

on code reuse so the reader can look back over the three chapters and see the evolution of the code as the problem got more complex.

Lastly, while not apparent, we have replicated much of the internal functionality of the professional code MODFLOW which we could compare our work to if we had to use the **R** script for something the professional code does not do easily. The scripts here are very crude – to professionalize them we would want to add the concept of “layers” so we could do 3D models, and want to generalize the file reading components of the tool. We would also want to move much of the computation engine components into

that is not well documented, I would start with an explicit approach because I could get code up and running quickly and identify other problems that might be encountered along the way.

their own prototype functions so the structure of the program is easier to understand and maintain. Nevertheless, if you have worked this far, you have seen what's under the hood of a groundwater flow code.⁷⁵

18.4 Exercises

- Figure 142 is a plan view of a confined aquifer system bounded by mountains (no-flow) and a river (constant head). The 10 meter thick aquifer is approximately horizontal. The material properties are reported as transmissivities ($T = K \times \text{thickness}$). The storage coefficient for the aquifer system is 0.05. Figure 143 is the same aquifer system with a recharge gallery and a well field.




Figure 142. Aquifer system bounded on three sides with mountains, and one side by a river. The aquifer is about 10 meters thick at the river.

- Develop or modify an **R** script to simulate the unsteady confined aquifer system shown. Build your tool so that recharge and pumping are read in as separate arrays. To handle the mountain intrusion into the plain from the East in the sketch you will either have to modify how your script

⁷⁵ Actually you have also seen the code one would use for linear heat flow computations – because the PDEs look the same, hence the solvers will work the same.




Figure 143. Aquifer system bounded on three sides with mountains, and one side by a river. The aquifer is about 10 meters thick at the river. A recharge gallery is shown where water is recharged into the aquifer. A discharge gallery (wellfield) is shown where water is removed from the aquifer..

implements boundary conditions, or treat the mountain range as a low-permeability (small T) inclusion into the aquifer.

- (b) Run your script and produce a contour plot of the head distribution in the pre-development case at equilibrium.
 - (c) Describe how you decided equilibrium is reached.
 - (d) Compare your equilibrium solution to the same problem conditions run using your steady-flow solver.
 - (e) Run the script with the post-development case, and produce contour plots of the head distribution at 5 and 10 years of operation.
 - (f) Estimate the fraction of recharge water that is captured by the well field, assuming the 10-year contour plot is an equilibrium condition.
2. Repeat Example 2 with the following changes:
- (a) Modify the script (or build your own) that reads recharge as a separate array (instead of the R-P hack that was employed in the example).
 - (b) Assume cloud seeding doubles the average recharge rate. Determine the

amount of water that can be pumped from the well field such that the head in the wellfield is greater than 5 meters after 60 years of pumping.

- (c) Produce a contour map of the head distribution at your solution.
- 3. (Advanced) Modify your scripts to allow for different stress periods – intervals of time where recharge and pumping remain constant but where different stress periods can have different recharge and pumping rates. Develop examples to illustrate the effect of the modification.
- 4. (Advanced) Modify the scripts to multiple layers to approximate a 3D aquifer system. You will need to modify the difference equations to simulate the vertical flow component between the layers.

19 Flow Nets

Example 3: 2D Stream Function using Jacobi Iteration In steady aquifer flow, the flow is irrotational (or at least can be modeled as such). There exists an orthogonal function called the stream function (it is the function that exists in the flow field when vorticity is zero). A really good explanation of stream functions (and streamlines) appears on pages 381–398 in Zheng and Bennett (1995).

This function can be used to generate plots of streamlines for the same system. The principal changes are the material properties representation and the boundary conditions.

The stream function Ψ as a partial differential equation is

$$0 = \frac{\partial}{\partial x} \left(\frac{1}{T_y} \frac{\partial \Psi}{\partial x} \right) + \frac{\partial}{\partial y} \left(\frac{1}{T_x} \frac{\partial \Psi}{\partial y} \right) \quad (263)$$

Observe how the material property is inverted and changes directional identity, otherwise the equation is structurally identical to the groundwater flow equation (for steady flow).

The difference equation is also almost the same

$$0 = \left[\frac{1}{\Delta x} \frac{1}{T_y} \frac{\Psi_{i-1,j} - \Psi_{i,j}}{\Delta x} + \frac{1}{\Delta y} \frac{1}{T_x} \frac{\Psi_{i,j-1} - \Psi_{i,j}}{\Delta y} \right] - \left[\frac{1}{\Delta x} \frac{1}{T_y} \frac{\Psi_{i,j} - \Psi_{i+1,j}}{\Delta x} + \frac{1}{\Delta y} \frac{1}{T_x} \frac{\Psi_{i,j} - \Psi_{i,j+1}}{\Delta y} \right] \quad (264)$$

The substitutions are

$$\begin{aligned} A_{i,j} &= \frac{1}{2\Delta x^2} (T_{y,(i-1,j)}^{-1} + T_{y,(i,j)}^{-1}) \\ B_{i,j} &= \frac{1}{2\Delta x^2} (T_{y,(i,j)}^{-1} + T_{y,(i+1,j)}^{-1}) \\ C_{i,j} &= \frac{1}{2\Delta y^2} (T_{x,(i,j-1)}^{-1} + T_{x,(i,j)}^{-1}) \\ D_{i,j} &= \frac{1}{2\Delta y^2} (T_{x,(i,j)}^{-1} + T_{x,(i,j+1)}^{-1}) \end{aligned} \quad (265)$$

Substitution into the difference equation yields

$$0 = A_{i,j}\Psi_{i-1,j} + B_{i,j}\Psi_{i+1,j} - (A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j})\Psi_{i,j} + C_{i,j}\Psi_{i,j-1} + D_{i,j}\Psi_{i,j+1} \quad (266)$$

As before we can explicitly write the cell equation for $h_{i,j}$ as

$$\Psi_{i,j} = \frac{[A_{i,j}\Psi_{i-1,j} + B_{i,j}\Psi_{i+1,j} + C_{i,j}\Psi_{i,j-1} + D_{i,j}\Psi_{i,j+1}]}{[A_{i,j} + B_{i,j} + C_{i,j} + D_{i,j}]} \quad (267)$$

So at this point we could literally use the same script, however boundary conditions also “invert.” A no-flow head-domain boundary is a constant value stream function-domain boundary. A constant value head-domain boundary is a zero-gradient stream function-domain boundary.

So we could use the same code, but probably are better off building a separate code (it can read the same input file), and use it to generate stream functions. The prior example code is modified to generate the stream function for its case (and plot the stream function contours), which if overlaid on the head contours produces a flow net.

The modified code literally changes the names of the head array to stream function, modifies how the material properties (A, B, C, D) are constructed, and modified how the boundary conditions are incorporated. Listing 76 is a listing that implements the method – notice how the no-flow boundary conditions are implemented.

In this example the value of the stream function is arbitrarily set to range from 0 to 100. One useful interpretation of stream function values is that their differences indicate the flow fraction (of total flow) that flows between the streamlines (contour lines of constant stream function value).

Listing 76. R code demonstrating an Stream Function Simulator for 2D Steady Flow. This code fragment implements the Jacobi iteration. A subsequent listing shows the contour plot syntax. In the example the two fragments are joined and run as a single source file

```
# 2D-streamfunction
# Implements Finite-Difference Stream Function using Jacobi Iteration
# Assumes no-flow boundary row=1, and nrow ==> constant stream function
# Assumes fixed head boundary col=1, and ncol ==> no-flux stream function -- stream
  function runs from 0 to 100
zz <- file("input1.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
nrows <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
ncols <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
  skipNul = FALSE))
streamfn <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
hydcondx <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
hydcondy <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
  FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
streamfn <-as.numeric(unlist(strsplit(streamfn,split=" ")))
hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
# convert the numeric vectors into matrices for easier indexing
streamfn <-matrix(streamfn,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((1.0/hydcondy[irow-1,jcol ]+1.0/hydcondy[irow ,jcol ])*deltaz)
      /(2.0*deltax^2)
    bmat[irow,jcol]<-((1.0/hydcondy[irow ,jcol ]+1.0/hydcondy[irow+1,jcol ])*deltaz)
      /(2.0*deltax^2)
    cmat[irow,jcol]<-((1.0/hydcondx[irow ,jcol-1]+1.0/hydcondx[irow ,jcol ])*deltaz)
      /(2.0*deltay^2)
    dmat[irow,jcol]<-((1.0/hydcondx[irow ,jcol ]+1.0/hydcondx[irow ,jcol+1])*deltaz)
      /(2.0*deltay^2)
  }
}
# begin the calculations
streamold<-streamfn # copy the head array, used to test for stopping
maxit <- 100 # set the maximum number of iterations (to prevent infinite loop)
for (iter in 1:maxit){
  # first and last row special == no flow boundaries in head, fixed value streamfunction
  for(jcol in 1:ncols){
    streamfn[1,jcol]=0.0
    streamfn[nrows,jcol]=100.0
  }
  for (irow in 2:(nrows-1)){
    # first and last columns special == fixed head boundaries, no-gradient stream function
    streamfn[irow,1]=streamfn[irow,2]
    streamfn[irow,nrows]=streamfn[irow,nrows-1]
    for (jcol in 2:(nrows-1)){
      streamfn[irow,jcol] =
        (amat[irow,jcol]*streamfn[irow-1,jcol ] +
         bmat[irow,jcol]*streamfn[irow+1,jcol ] +
         cmat[irow,jcol]*streamfn[irow ,jcol-1] +
         dmat[irow,jcol]*streamfn[irow ,jcol+1] )/
        (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
    }
  }
  # test for stopping iterations
  percentdiff <- sum((streamfn-streamold)^2)
  if (percentdiff < tolerance){
    message("Exit iterations because tolerance met")
    break}
  streamold<-streamfn #update the current head vector
}
}
```

Listing 77. R code demonstrating an Stream Function Simulator for 2D Steady Flow. This code fragment implements the contour plotting

```
# echo contents for debugging
# streamfn
# streamold
#####
### built position arrays for contour plotting ###
#####
distancex<-numeric(nrows)
distancey<-numeric(ncols)
distancex[1]<-50
distancey[1]<-50
for (irow in 2:nrows){distancex[irow]<-distancex[irow-1]+deltax}
for (jcol in 2:ncols){distancey[jcol]<-distancey[jcol-1]+deltay}
#####
### contour plot of head -- note axes are rotated ###
#####
contour(distancex,distancey,streamfn,
        plot.title = title(main = "Stream Function Model",
                           xlab = "Meters (Y axis) ==>>>",
                           ylab = "Meters (X axis) ==>>>"))
}
```

Figure 144 is the stream function result. Compare it to Figure 121 and it should be clear that the “lines” are at right angles to each other –that is the stream function is orthogonal to the head function (which is anticipated, because of the nature of the relationship between stream function and head functions; the orthogonality is a consequence of the flow satisfying the the Cauchy-Riemann conditions.).




Figure 144. Output from R script for stream function model..

To complete the example, and prepare for the next example, we will modify the script one more time to:

1. Read in the material property values, head values, etc. (no change).

2. Compute the head distribution.
3. Compute the stream function distribution. (merge the two scripts)
4. Plot the head and stream function on the same graph – use different colors.
Also rotate the plots so axes agree with the problem statement.

Figure 145 is the result of these modifications.

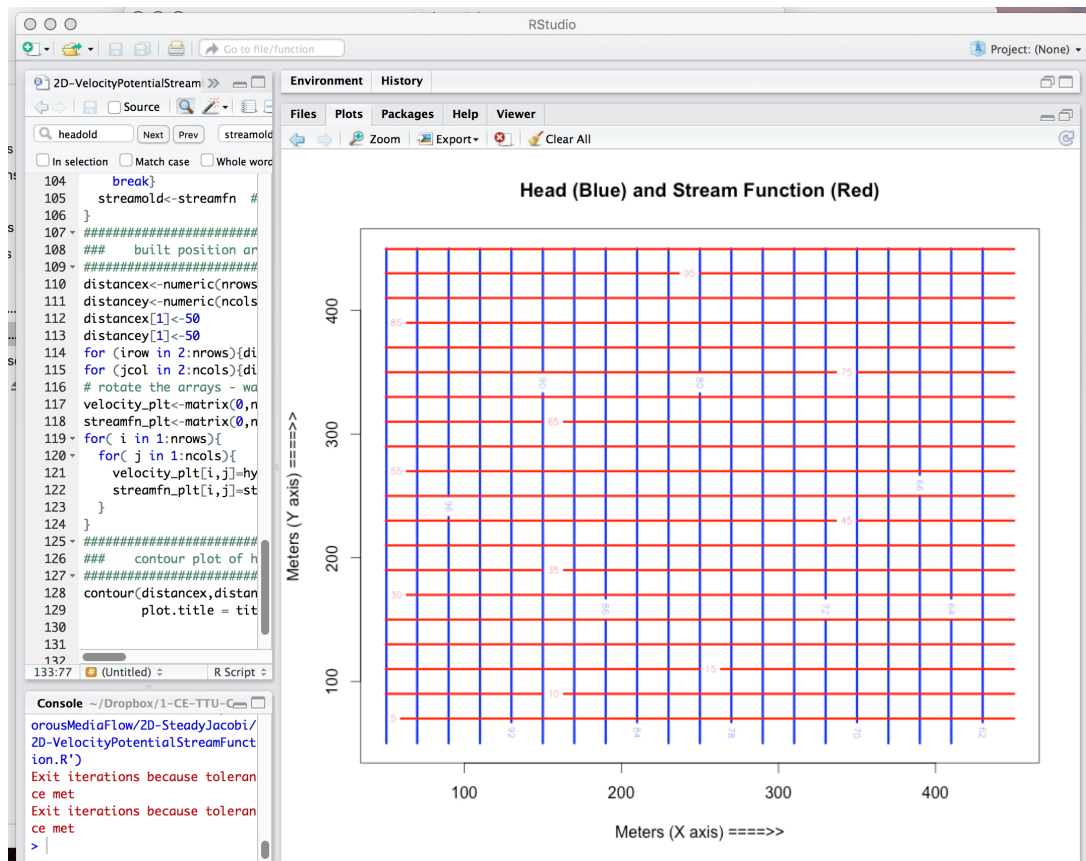


Figure 145. Output from R script for velocity-potential, stream function model, rotated and annotated to fit the original problem statement.

Listing 78 is a script fragment that implements the velocity-potential (here the same thing as the head equations) portion of the computations. The data file is the same, the only difference is the head values are copied to another vector (the stream functions) for use in the next fragment.

Listing 79 is a script fragment that implements the stream-function calculations. The A, B, C, D matrices are re-initialized and re-used. The stream function values are computed using the same computation engine (code is repeated – generally poor practice; done here to illustrate the re-use).

Listing 80 is the script fragment that rotates the results and plots the flow net.

Listing 78. R code demonstrating an Velocity Potential (Aquifer Head) Simulator for 2D Steady Flow. This code fragment implements the contour plotting

```
# 2D Velocity Potential Stream Function Model
# hydhead is velocity potential
# streamfn is stream function
zz <- file("input1.dat", "r") # Open a connection named zz to file named input.dat
# read the simulation conditons
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
nrows <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
ncols <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
hydhead <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondx <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondy <-(readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <-matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
# copy the hydhead array into streamfn for later use
streamfn <- hydhead
# here we perform the velocity potential calculations
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltaz)/(2.0*deltax
~2)
    bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltaz)/(2.0*deltax
~2)
    cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltaz)/(2.0*deltay
~2)
    dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltaz)/(2.0*deltay
~2)
  }
}
# veloity potential
headold<-hydhead # copy the head array, used to test for stopping
maxit <- 100 # set the maximum number of iterations (to prevent infinite loop)
for (iter in 1:maxit){
  # first and last row special == no flow boundaries
  for(jcol in 1:ncols){
    hydhead[1,jcol]=hydhead[2,jcol]
    hydhead[nrows,jcol]=hydhead[nrows-1,jcol]
  }
  for (irow in 2:(nrows-1)){
    for (jcol in 2:(ncols-1)){
      hydhead[irow,jcol] =
        (amat[irow,jcol]*hydhead[irow-1,jcol ] +
         bmat[irow,jcol]*hydhead[irow+1,jcol ] +
         cmat[irow,jcol]*hydhead[irow ,jcol-1] +
         dmat[irow,jcol]*hydhead[irow ,jcol+1] )/
        (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
    }
  }
  # test for stopping iterations
  percentdiff <- sum((hydhead-headold)^2)
  if (percentdiff < tolerance){
    message("Exit iterations because tolerance met")
    break}
  headold<-hydhead #update the current head vector
}
}
```

Listing 79. R code demonstrating an Stream Function Simulator for 2D Steady Flow. This code fragment implements the stream function

```
# built the streamfn transmissivity arrays -- notice reuse of the a,b,c,d arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((1.0/hydcondy[irow-1,jcol ]+1.0/hydcondy[irow ,jcol ])*deltaz)
      /(2.0*deltax^2)
    bmat[irow,jcol]<-((1.0/hydcondy[irow ,jcol ]+1.0/hydcondy[irow+1,jcol ])*deltaz)
      /(2.0*deltax^2)
    cmat[irow,jcol]<-((1.0/hydcondx[irow ,jcol-1]+1.0/hydcondx[irow ,jcol ])*deltaz)
      /(2.0*deltay^2)
    dmat[irow,jcol]<-((1.0/hydcondx[irow ,jcol ]+1.0/hydcondx[irow ,jcol+1])*deltaz)
      /(2.0*deltay^2)
  }
}
streamold<-streamfn # copy the head array, used to test for stopping
maxit <- 100 # set the maximum number of iterations (to prevent infinite loop)
for (iter in 1:maxit){
  # first and last row special == no flow boundaries in head, fixed value streamfunction
  for(jcol in 1:ncols){
    streamfn[1,jcol]=0.0
    streamfn[nrows,jcol]=100.0
  }
  for (irow in 2:(nrows-1)){
    # first and last columns special == fixed head boundaries, no-gradient stream function
    streamfn[irow,1]=streamfn[irow,2]
    streamfn[irow,nrows]=streamfn[irow,nrows-1]
    for (jcol in 2:(nrows-1)){
      streamfn[irow,jcol] =
        (amat[irow,jcol]*streamfn[irow-1,jcol ] +
         bmat[irow,jcol]*streamfn[irow+1,jcol ] +
         cmat[irow,jcol]*streamfn[irow ,jcol-1] +
         dmat[irow,jcol]*streamfn[irow ,jcol+1] )/
        (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
    }
  }
  # test for stopping iterations
  percentdiff <- sum((streamfn-streamold)^2)
  if (percentdiff < tolerance){
    message("Exit iterations because tolerance met")
    break}
  streamold<-streamfn #update the current head vector
}
}
```

Listing 80. R code demonstrating an Velocity Potential – Stream Function Simulator for 2D Steady Flow. This code fragment implements the contour plotting

```
#####
###    built position arrays for contour plotting    ###
#####
distancex<-numeric(nrows)
distancey<-numeric(ncols)
distancex[1]<-50
distancey[1]<-50
for (irow in 2:nrows){distancex[irow]<-distancex[irow-1]+deltax}
for (jcol in 2:ncols){distancey[jcol]<-distancey[jcol-1]+deltay}
# rotate the arrays - wasting memory.  Fix for homework.
velocity_plt<-matrix(0,nrows,ncols)
streamfn_plt<-matrix(0,nrows,ncols)
for( i in 1:nrows){
  for( j in 1:ncols){
    velocity_plt[i,j]=hydhead[j,i]
    streamfn_plt[i,j]=streamfn[j,i]
  }
}
#####
###    contour plot of head -- note axes are rotated    ###
#####
contour(distancex,distancey,velocity_plt,
        plot.title = title(main = "Head (Blue) and Stream Function (Red)",
                           xlab = "Meters (X axis) ==>>",
                           ylab = "Meters (Y axis) ==>>"),
        col="blue",lwd=3,xlim=c(50,450),ylim=c(50,450),nlevels=20)
contour(distancey,distancex,streamfn_plt,add=TRUE,col="red",lwd=3,nlevels=20)
}
```

Example 3: 2D Flow Net in a Confined Aquifer using Jacobi Iteration with Low Permeability Inclusion

Figure 146 is a schematic of a different situation that now only requires us to change the contents of the data file, and re-run the scripts unchanged. Some additional modifications have been added, mostly messages to the user because of anticipated long run times. The data file is changed a bit and two lines are added to help with the plotting – they represent the axis labels used in the contour plots. The boundary conditions are still directly coded into the algorithm, and that would be the next modification to the code - general boundary condition information.⁷⁶




Figure 146. Schematic of vertical slice in an aquifer with low permeability inclusion. Values are indicated on the schematic. Example illustrates how to use the scripts to generate flow nets for the vertical slice .

The following pages contain the code fragments (listings) for the velocity potential, the stream function, and the contour plotting. As above, these listings are combined into a single file (the fragmentation herein is to fit the printed page layout) and then run as a script.

Listing 81 is the listing for the velocity potential calculations.

Listing 82 is the listing for the stream function calculations.

Listing 83 is the is the listing for the plotting calculations.

Listing 84 is the input file for the problem. The file in this case is named `input2.dat`. In addition the generalized boundary conditons, the reader should consider making the program prompt the user for the file name, so that the program is somewhat deployable.

⁷⁶That modification is left for homework.

Listing 81. Velocity Potential Script.

```

# 2D Velocity Potential Stream Function Model
# hydhead is velocity potential; streamfm is stream function
rm(list=ls()) # deallocate (clear) memory
zz <- file("input2.dat", "r") # Open a connection named zz to file named input.dat to read
input conditions
deltax <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltay <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
deltaz <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
nrows <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
ncols <-as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
tolerance <- as.numeric(readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown",
skipNul = FALSE))
distancex <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
distancey <- (readLines(zz, n = 1, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydhead <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondx <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
hydcondy <- (readLines(zz, n = nrows, ok = TRUE, warn = TRUE,encoding = "unknown", skipNul =
FALSE))
close(zz)
# split the multiple column strings into numeric components for a vector
distancex <-as.numeric(unlist(strsplit(distancex,split=" ")))
distancey <-as.numeric(unlist(strsplit(distancey,split=" ")))
hydhead <-as.numeric(unlist(strsplit(hydhead,split=" ")))
hydcondx <-as.numeric(unlist(strsplit(hydcondx,split=" ")))
hydcondy <-as.numeric(unlist(strsplit(hydcondy,split=" ")))
# convert the numeric vectors into matrices for easier indexing
hydhead <-matrix(hydhead,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondx <-matrix(hydcondx,nrow=nrows,ncol=ncols,byrow = TRUE)
hydcondy <-matrix(hydcondy,nrow=nrows,ncol=ncols,byrow = TRUE)
# built the transmissivity arrays
amat<-matrix(0,nrows,ncols)
bmat<-matrix(0,nrows,ncols)
cmat<-matrix(0,nrows,ncols)
dmat<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amat[irow,jcol]<-((hydcondx[irow-1,jcol ]+hydcondx[irow ,jcol ])*deltaz)/(2.0*deltax
^2)
    bmat[irow,jcol]<-((hydcondx[irow ,jcol ]+hydcondx[irow+1,jcol ])*deltaz)/(2.0*deltax
^2)
    cmat[irow,jcol]<-((hydcondy[irow ,jcol-1]+hydcondy[irow ,jcol ])*deltaz)/(2.0*deltay
^2)
    dmat[irow,jcol]<-((hydcondy[irow ,jcol ]+hydcondy[irow ,jcol+1])*deltaz)/(2.0*deltay
^2)
  }
}
headold<-hydhead # copy the head array, used to test for stopping
tolflag<-FALSE
maxit <- 500000 # set the maximum number of iterations (to prevent infinite loop)
for (iter in 1:maxit){
  # first and last row special == no flow boundaries
  for(jcol in 1:ncols){
    hydhead[1,jcol]<-hydhead[2,jcol]
    hydhead[nrows,jcol]<-hydhead[nrows-1,jcol]
  }
  for (irow in 2:(nrows-1)){
    for (jcol in 2:(ncols-1)){
      hydhead[irow,jcol] <-
        (amat[irow,jcol]*hydhead[irow-1,jcol ] +
         bmat[irow,jcol]*hydhead[irow+1,jcol ] +
         cmat[irow,jcol]*hydhead[irow ,jcol-1] +
         dmat[irow,jcol]*hydhead[irow ,jcol+1] )/
        (amat[irow,jcol]+bmat[irow,jcol]+cmat[irow,jcol]+dmat[irow,jcol])
    }
  }
  # test for stopping iterations
  percentdiff <- sum((hydhead-headold)^2)
  if (percentdiff < tolerance){
    message("Exit iterations in velocity potential because tolerance met");
    message("Iterations =", iter);
    tolflag <- TRUE
    break}
  headold<-hydhead #update the current head vector
  if( iter %% 5000 == 0){message("Calculating in Potential Function ",iter," of ",maxit, "
iterations")}
}
if (tolflag == FALSE){message("Exit iterations in potential function at max iterations")}
}

```

Listing 82. Stream Function Script.

```

streamfn<-matrix(1.0,nrows,ncols)
for (i in 1:nrows){
  for(j in 1:ncols){
    streamfn[i,j]=as.numeric(i)/as.numeric(nrows)
  }
}
streamfn
# built the streamfn transmissivity arrays
amats<-matrix(0,nrows,ncols)
bmats<-matrix(0,nrows,ncols)
cmats<-matrix(0,nrows,ncols)
dmats<-matrix(0,nrows,ncols)
for(irow in 2:(nrows-1)){
  for(jcol in 2:(ncols-1)){
    amats[irow,jcol]<-((1.0/hydcondy[irow-1,jcol ]+1.0/hydcondy[irow ,jcol ])*deltaz)
      /(2.0*deltax^2)
    bmats[irow,jcol]<-((1.0/hydcondy[irow ,jcol ]+1.0/hydcondy[irow+1,jcol ])*deltaz)
      /(2.0*deltax^2)
    cmats[irow,jcol]<-((1.0/hydcondx[irow ,jcol-1]+1.0/hydcondx[irow ,jcol ])*deltaz)
      /(2.0*deltay^2)
    dmats[irow,jcol]<-((1.0/hydcondx[irow ,jcol ]+1.0/hydcondx[irow ,jcol+1])*deltaz)
      /(2.0*deltay^2)
  }
}
streamold<-streamfn # copy the head array, used to test for stopping
tolflag<-FALSE
maxit <- maxit/10 # set the maximum number of iterations (to prevent infinite loop)
for (iter in 1:maxit){
  # first and last row special == no flow boundaries in head, fixed value streamfunction
  for(jcol in 1:ncols){
    streamfn[1,jcol]=0.0
    streamfn[nrows,jcol]=1.0
  }
  for (irow in 2:(nrows-1)){
    # first and last columns special == fixed head boundaries, no-gradient stream function
    streamfn[irow,1]=streamfn[irow,2]
    streamfn[irow,ncols]=streamfn[irow,ncols-1]
    for (jcol in 2:(ncols-1)){
      streamfn[irow,jcol] =
        (amats[irow,jcol]*streamfn[irow-1,jcol ] +
         bmats[irow,jcol]*streamfn[irow+1,jcol ] +
         cmats[irow,jcol]*streamfn[irow ,jcol-1] +
         dmats[irow,jcol]*streamfn[irow ,jcol+1] )/
        (amats[irow,jcol]+bmats[irow,jcol]+cmats[irow,jcol]+dmats[irow,jcol])
    }
  }
  # test for stopping iterations
  percentdiff <- sum((streamfn-streamold)^2)
  if (percentdiff < tolerance)
  {
    message("Exit iterations in stream function because tolerance met");
    message("Iterations =", iter);
    tolflag <- TRUE;
    break
  }
  streamold<-streamfn #update the current head vector
  if( iter %% 5000 == 0){message("Calculating in Stream Function ",iter," of ",maxit, "
    iterations")
  }
}
if (tolflag == FALSE){message("Exit iterations in stream function at max iterations")}
}

```

Listing 83. Contour plotting script.

```

#####
### rotate arrays for contour plotting -- position values read in input ###
#####
velocity_plt<-matrix(0,ncols,nrows)
streamfn_plt<-matrix(0,ncols,nrows)
for ( i in 1:nrows){
  for( j in 1:ncols){
    velocity_plt[j,i]=hydhead[i,j]
    streamfn_plt[j,i]=streamfn[i,j]
  }
}
#####
### contour plot of head -- note axes are rotated ###
#####
contour(distancey,distancex,velocity_plt,
  plot.title = title(main = "Head (Blue) and Stream Function (Red)",
    xlab = "Meters (X axis) ==>>>",
    ylab = "Meters (Y axis) ==>>>"),
  col="blue",lwd=3,nlevels=20)
contour(distancey,distancex,streamfn_plt ,add=TRUE,col="red",lwd=3,nlevels=20)
}

```


[Modifications to handle generalized boundary conditions]

19.1 Unconfined Aquifer Flow

19.1.1 Finite-Difference Methods

19.2 Exercises

20 Heat and Mass Transport

20.1 Concentration

20.2 Tracer Hypothesis

20.3 Advection (Convection) Process

20.4 Diffusion Process

20.5 Dispersion Process

20.6 Reaction Processes

20.6.1 Linear Equilibrium Isotherms and Concept of Retardation

20.7 Advection, Dispersion, Retardation, Decay Mathematics

21 Analytical Solutions to ADR Equations

21.1 Impulse Solution

21.2 Ogata-Banks Solution

21.3 2D-Injection (Hunt) Solution

21.4 2D-, and 3D- Domenico Robbins Solutions

22 Numerical Solutions to ADR Equations

22.1 Explicit Methods

22.1.1 Centered Differences

22.1.2 Upwind Formulation(s)

22.2 Crank-Nicholson Methods

References

- Christian, B. and Griffiths, T. (2016). Algorithms to Live By: The Computer Science of Human Decisions Henry Holt and Co.. Kindle Edition.
- Cunge, J.A., Holly, F.M., Verwey, A. (1980). Practical Aspects of Computational River Hydraulics. Pittman Publishing Inc. , Boston, MA. pp. 7-50
- Zheng, C. and Bennett, G. D. (1995). *Applied Contaminant Transport Modeling*. Van Nostrand Reinhold.
- Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W. T. 1986, Numerical Recipes:*The Art of Scientific Computing*. Cambridge University Press, London. 818p.
- Gill, P.E., Murray, W, Wright, M. H., 1981. Practical Optimization. Academic Press, San Diego. 401p.
- Wikipedia discussion of Newton's method. http://en.wikipedia.org/wiki/Newton's_method.
- Wikipedia discussion of matrix inversion. http://en.wikipedia.org/wiki/Invertible_matrix.
- Chin, D. A. (2006). *Water-Resources Engineering*. Prentice Hall.
- Haman YM, Brameller A. (1971) Hybrid method for the solution of piping networks. Proc IEEE 1971;118(11):1607-12.