

# ECEN 615

## Methods of Electric Power Systems Analysis

### Lecture 11: Sparse Systems

---

Prof. Tom Overbye

Dept. of Electrical and Computer Engineering

Texas A&M University

[overbye@tamu.edu](mailto:overbye@tamu.edu)



TEXAS A&M  
UNIVERSITY

# Announcements

---



- Homework 3 is assigned today, due Thursday Oct 11
- Midterm exam is Oct 18 in class
  - Off campus students should work with Iyke to get their exam proctoring setup
  - Closed book, closed notes, but one 8.5 by 11 inch note sheet allowed

# Linear System Solution: Introduction



- A problem that occurs in many fields is the solution of linear systems  $\mathbf{Ax} = \mathbf{b}$  where  $\mathbf{A}$  is an  $n$  by  $n$  matrix with elements  $a_{ij}$ , and  $\mathbf{x}$  and  $\mathbf{b}$  are  $n$ -vectors with elements  $x_i$  and  $b_i$  respectively
- In power systems we are particularly interested in systems when  $n$  is relatively large and  $\mathbf{A}$  is sparse
  - How large is large is changing
- A matrix is sparse if a large percentage of its elements have zero values
- Goal is to understand the computational issues (including complexity) associated with the solution of these systems

# Introduction, cont.

---



- Sparse matrices arise in many areas, and can have domain specific structures
  - Symmetric matrices
  - Structurally symmetric matrices
  - Tridiagonal matrices
  - Banded matrices
- A good (and free) book on sparse matrices is available at [www-users.cs.umn.edu/~saad/IterMethBook\\_2ndEd.pdf](http://www-users.cs.umn.edu/~saad/IterMethBook_2ndEd.pdf)
- ECEN 615 is focused on problems in the electric power domain; it is not a general sparse matrix course
  - Much of the early sparse matrix work was done in power!

# Computational Complexity

---



- Computational complexity indicates how the number of numerical operations scales with the size of the problem
- Computational complexity is expressed using the “Big O” notation; assume a problem of size  $n$ 
  - Adding the number of elements in a vector is  $O(n)$
  - Adding two  $n$  by  $n$  full matrices is  $O(n^2)$
  - Multiplying two  $n$  by  $n$  full matrices is  $O(n^3)$
  - Inverting an  $n$  by  $n$  full matrix, or doing Gaussian elimination is  $O(n^3)$
  - Solving the traveling salesman problem by brute-force search is  $O(n!)$

# Computational Complexity

---



- Knowing the computational complexity of a problem can help to determine whether it can be solved (at least using a particular method)
  - Scaling factors do not affect the computation complexity
    - an algorithm that takes  $n^3/2$  operations has the same computational complexity of one that takes  $n^3/10$  operations (though obviously the second one is faster!)
- With  $O(n^3)$  factoring a full matrix becomes computationally intractable quickly!
  - A 100 by 100 matrix takes a million operations (give or take)
  - A 1000 by 1000 matrix takes a billion operations
  - A 10,000 by 10,000 matrix takes a trillion operations!

# Sparse Systems

---



- The material presented so far applies to any arbitrary linear system
- The next step is to see what happens when we apply triangular factorization to a sparse matrix
- For a sparse system, only nonzero elements need to be stored in the computer since no arithmetic operations are performed on the 0's
- The triangularization scheme is adapted to solve sparse systems in such a way as to preserve the sparsity as much as possible

# Sparse Matrix History

---



- A nice overview of sparse matrix history is by Iain Duff at <http://www.siam.org/meetings/la09/talks/duff.pdf>
- Sparse matrices developed simultaneously in several different disciplines in the early 1960's with power systems definitely one of the key players (Bill Tinney from BPA)
- Different disciplines claim credit since they didn't necessarily know what was going on in the others

# Sparse Matrix History

---



- In power systems a key N. Sato, W.F. Tinney, “Techniques for Exploiting the Sparsity of the Network Admittance Matrix,” Power App. and Syst., pp 944-950, December 1963
  - In the paper they are proposing solving systems with up to 1000 buses (nodes) in 32K of memory!
  - You’ll also note that in the discussion by El-Abiad, Watson, and Stagg they mention the creation of standard test systems with between 30 and 229 buses (this surely included the now famous 118 bus system)
  - The BPA authors talk “power flow” and the discussors talk “load flow.”
- Tinney and Walker present a much more detailed approach in their 1967 IEEE Proceedings paper titled “Direct Solutions of Sparse Network Equations by Optimally Order Triangular Factorization”

# Sparse Matrix Computational Order

---



- The computational order of factoring a sparse matrix, or doing a forward/backward substitution depends on the matrix structure
  - Full matrix is  $O(n^3)$
  - A diagonal matrix is  $O(n)$ ; that is, just invert each element
- For power system problems the classic paper is F. L. Alvarado, “Computational complexity in power systems,” *IEEE Transactions on Power Apparatus and Systems*, May/June 1976
  - $O(n^{1.4})$  for factoring,  $O(n^{1.2})$  for forward/backward
  - For a 100,000 by 100,000 matrix changes computation for factoring from 1 quadrillion to 10 million!

# Inverse of a Sparse Matrix

---



- The inverse of a sparse matrix is NOT in general a sparse matrix
- We never (or at least very, very, very seldom) explicitly invert a sparse matrix
  - Individual columns of the inverse of a sparse matrix can be obtained by solving  $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$  with  $\mathbf{b}$  set to all zeros except for a single nonzero in the position of the desired column
  - If a few desired elements of  $\mathbf{A}^{-1}$  are desired (such as the diagonal values) they can usually be computed quite efficiently using sparse vector methods (a topic we'll be considering soon)
- We can't invert a singular matrix (with sparse or not)

# Computer Architecture Impacts

---



- With modern computers the processor speed is many times faster than the time it takes to access data in main memory
  - Some instructions can be processed in parallel
- Caches are used to provide quicker access to more commonly used data
  - Caches are smaller than main memory
  - Different cache levels are used with the quicker caches, like L1, have faster speeds but smaller sizes; L1 might be 64K, whereas the slower L2 might be 1M
- Data structures can have a significant impact on sparse matrix computation

# ECEN 615 Sparsity Limitations

---



- Sparse matrices arise in many areas, and can have domain specific structures
  - Symmetric matrices
  - Structurally symmetric matrices
  - Tridiagonal matrices
  - Banded matrices
- ECEN 615 is focused on problems that arise in the electric power; it is not a general sparse matrix course
  - A variety of different sparse matrices are available at <https://sparse.tamu.edu/> (Prof. Tim Davis)
    - Some have more than 200 million rows and columns; some have more than six billion elements!

# Full Matrix versus Sparse Matrix Storage

---



- Full matrices are easily stored in arrays with just one variable needed to store each value since the value's row and column are implicitly available from its matrix position
- With sparse matrices two or three elements are needed to store each value
  - The zero values are not explicitly stored
  - The value itself, its row number and its column number
  - Storage can be reduced by storing all the elements in a particular row or column together
- Because large matrices are often quite sparse, the total storage is still substantially reduced

# Sparse Matrix Usage Can Determine the Optimal Storage

---



- How a sparse matrix is used can determine the best storage scheme to use
  - Row versus column access; does structure change
- Is the matrix essentially used only once? That is, its structure and values are assumed new each time used
- Is the matrix structure constant, with its values changed
  - This would be common in the N-R power flow, in which the structure doesn't change each iteration, but its values do
- Is the matrix structure and values constant, with just the **b** vector in  $\mathbf{Ax}=\mathbf{b}$  changing
  - Quite common in transient stability solutions

# Numerical Precision

---



- Required numerical precision determines type of variables used to represent numbers
  - Specified as number of bytes, and whether signed or not
- For Integers
  - One byte is either 0 to 255 or -128 to 127
  - Two bytes is either smallint (-32,768 to 32,767) or word (0 to 65,536)
  - Four bytes is either Integer (-2,147,483,648 to 2,147,483,647) or Cardinal (0 to 4,294,967,295)
    - This is usually sufficient for power system row/column numbers
  - Eight bytes (Int64) if four bytes is not enough

# Numerical Precision, cont.

---



- For floating point values using choice is between four bytes (single precision) or eight bytes (double precision); extended precision has ten bytes
  - Single precision allows for 6 to 7 significant digits
  - Double precision allows for 15 to 17 significant digits
  - Extended allows for about 18 significant digits
  - More bytes requires more storage
  - Computational impacts depend on the underlying device; on PCs there isn't much impact; GPUs can be 3 to 8 times slower for double precision
- For most power problems double precision is best

# General Sparse Matrix Storage

---



- A general approach for storing a sparse matrix would be using three vectors, each dimensioned to number of elements
  - AA: Stores the values, usually in power system analysis as double precision values (8 bytes)
  - JR: Stores the row number; for power problems usually as an integer (4 bytes)
  - JC: Stores the column number, again as an integer
- If unsorted then both row and column are needed
- New elements could easily be added, but costly to delete
- Unordered approach doesn't make for good computation since elements used next computationally aren't necessarily nearby
- Usually ordered, either by row or column

# Sparse Storage Example



- Assume 
$$\mathbf{A} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -4 & -3 & -2 & 10 \end{bmatrix}$$

- Then 
$$\mathbf{AA} = [5 \quad -4 \quad 4 \quad -3 \quad 3 \quad -2 \quad -4 \quad -3 \quad -2 \quad 10]$$
$$\mathbf{JR} = [1 \quad 1 \quad 2 \quad 2 \quad 3 \quad 3 \quad 4 \quad 4 \quad 4 \quad 4]$$
$$\mathbf{JC} = [1 \quad 4 \quad 2 \quad 4 \quad 3 \quad 4 \quad 1 \quad 2 \quad 3 \quad 4]$$

Note, this example is a symmetric matrix, but the technique is general

# Compressed Sparse Row Storage

---



- If elements are ordered (as was case for previous example) storage can be further reduced by noting we do not need to continually store each row number
- A common method for storing sparse matrices is known as the Compressed Sparse Row (CSR) format
  - Values are stored row by row
  - Has three vector arrays:
    - AA: Stores the values as before
    - JA: Stores the column index (done by JC in previous example)
    - IA: Stores the pointer to the index of the beginning of each row

# CSR Format Example



- Assume as before 
$$\mathbf{A} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -4 & -3 & -2 & 10 \end{bmatrix}$$

- Then 
$$\mathbf{AA} = [5 \quad -4 \quad 4 \quad -3 \quad 3 \quad -2 \quad -4 \quad -3 \quad -2 \quad 10]$$
$$\mathbf{JA} = [1 \quad 4 \quad 2 \quad 4 \quad 3 \quad 4 \quad 1 \quad 2 \quad 3 \quad 4]$$
$$\mathbf{IA} = [1 \quad 3 \quad 5 \quad 7]$$

# CSR Comments

---



- The CSR format reduces the storage requirements by taking advantage of needing only one element per row
- The CSR format has good advantages for computation when using cache since (as we shall see) during matrix operations we are often sequentially going through the vectors
- An alternative approach is Compressed Sparse Column (CSC), which is identical, except storing the values by column
- It is difficult to add values.
- We'll mostly use the linked list approach here, which makes matrix manipulation simpler

# Linked Lists: Classes and Objects

---

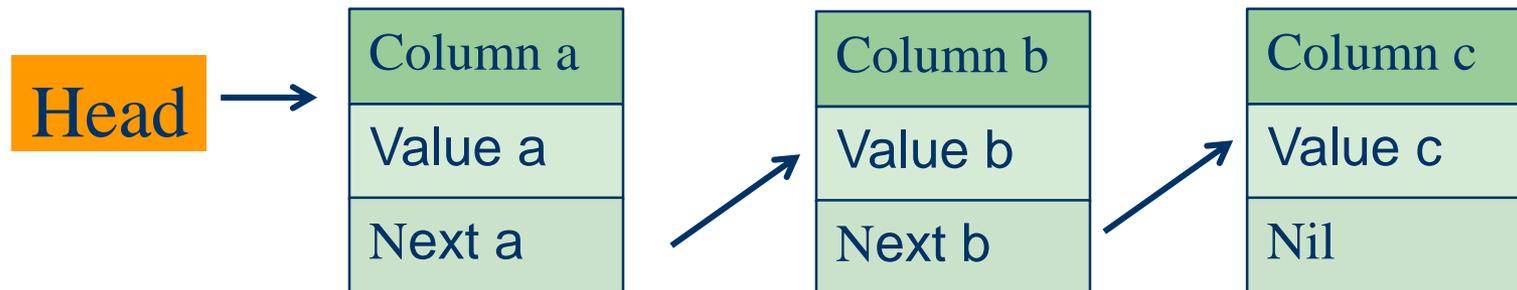


- In explaining the linked list approach it is helpful to use the concepts from object oriented programming (OOP) of classes and objects
  - Approach can also be used in non-OOP programming
- OOP can be thought of as a collection of objects interacting with each other
- Objects are instances of classes.
- Classes define the object fields and actions (methods)
- We'll define a class called sparse matrix element, with fields of value, column and next; each sparse matrix element is then an object of this class

# Linked Lists



- A linked list is just a group of objects that represent a sequence
  - We'll use linked lists to represent a row or column of a sparse matrix
- Each linked list has a head pointer that points to the first object in the list
  - For our sparse matrices the head pointer will be a vector of the rows or columns



# Sparse Matrix Storage with Linked Lists by Rows

---



- If we have an  $n$  by  $n$  matrix, setup a class called `TSparseElement` with fields `column`, `value` and `next`
- Setup an  $n$ -dimensional head pointer vector that points to the first element in each row
- Each nonzero corresponds to an object of class (type) `TSparseElement`
- We do not need to store the row number since it is given by the object's row
- For power system sparse matrices, which have nonzero diagonals, we also have a header pointer vector that points to the diagonal objects

# Linked Lists, cont.

---



- Linked lists can be singly linked, which means they just go in one direction (to the next element), or doubly linked, pointing to both the previous and next elements
  - Mostly we'll just need singularly linked lists
- With linked lists it is quite easy to add new elements to the list. This can be done in sorted order just by going down the list until the desired point is reached, then changing the next pointer for the previous element to the new element, and for the new element to the next element (for a singly linked list)

# On Board Example



- Draw the data structures for the matrix

$$\mathbf{A} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -4 & -3 & -2 & 10 \end{bmatrix}$$

# Example Pascal Code for Writing a Sparse Matrix



```
Procedure TSparMat.SMWriteMatlab(Var ft : Text; variableName : String; digits,rod : Integer;
                                ignoreZero : Boolean; local_MinValue : Double);

Var j : Integer;
    p1 : TMatEle;
Begin
For j := 1 to n Do Begin
    p1 := Row(j).Head;
    While p1 <> nil Do Begin
        If (not IgnoreZero) or (abs(p1.value) > local_MinValue) Then Begin
            If variableName <> " Then Writeln(ft,variableName+'(',j),',',(p1.col),')=' ,p1.value:digits:rod,');
            Else Writeln(ft,j:5,' ',p1.col:5,' ',p1.value:digits:rod);
        End;
        p1 := p1.next;
    End;
End;
End;
End;
```

# Sparse Working Row

---



- Before showing a sparse LU factorization it is useful to introduce the concept of a working row full vector
- This is useful because sometimes we need direct access to a particular value in a row
- The working row approach is to define a vector of dimension  $n$  and set all the values to zero
- We can then load a sparse row into the vector, with computation equal to the number of elements in the row
- We can then unload the sparse row from the vector by storing the new values in the linked list, and resetting the vector values we changed to zero

# Loading the Sparse Working Row



```
Procedure TSparMat.LoadSWRbyCol(rowJ : Integer; var SWR : PDVectorList);  
Var p1 : TMatEle;  
Begin  
  p1 := rowHead[rowJ];  
  While p1 <> nil Do Begin  
    SWR[p1.col] := p1.value;  
    p1 := p1.next;  
  End;  
End;
```

# Unloading the Sparse Working Row



```
Procedure TSParMat.UnLoadSWRbyCol(rowJ : Integer; var SWR :  
PDVectorList);  
Var p1 : TMatEle;  
Begin  
  p1 := rowHead[rowJ];  
  While p1 <> nil Do Begin  
    p1.value := SWR[p1.col];  
    SWR[p1.col] := 0;  
    p1 := p1.next;  
  End;  
End;
```

Note, there is no need to explicitly zero out all the elements each iteration since 1) most are still zero and 2) doing so would make it  $O(n^2)$ . The above code efficiently zeros out just the values that have changed.

# Doing an LU Factorization of a Sparse Matrix with Linked Lists



- Now we can show how to do an LU factorization of a sparse matrix stored using linked lists
- We will assume the head pointers are in the vector RowHead, and the diagonals in RowDiag
- Recall this was the approach for the full matrix

```
For i := 2 to n Do Begin // This is the row being processed
  For j := 1 to i-1 Do Begin // Rows subtracted from row i
    A[i,j] = A[i,j]/A[j,j] // This is the scaling
    For k := j+1 to n Do Begin // Go through each column in i
      A[i,k] = A[i,k] - A[i,j]*A[j,k]
    End;
  End;
End;
End;
```