



# Optimizing APL Matrix Indexing for Application Programmers

Eugene Ying  
Software Development

Aug 8, 2011

# Introduction

In the business world where APL is being used to perform large scale data processing, data matrices tend to be very large, with millions of rows and thousands of columns of data. When a very large matrix is manipulated thousands of times inside a loop, optimizing matrix indexing is important in reducing the function's run time. There are many cases where an application program's run time can be significantly reduced when data indexing has been optimized.

This paper talks about several techniques that can be used to optimize APL matrix indexing. The CPU time shown in the examples were based on Dyalog AIX V12.1 64-bit Classic APL running on the IBM pSeries server.

# Indexing Optimization Techniques

The major topics to be discussed are:

Matrix Column Data Fragmentation

Progressive Indexing

Index Manipulation

Grade Up and Grade Down Indices

# Multidimensional Array Storage

According to Wikipedia,

“For storing multidimensional arrays in linear memory, row-major order is used in C; column-major order is used in Fortran and MATLAB.”

We notice that APL also stores array items in row-major order.

# Matrix Column Data Fragmentation

## Matrix

|     |     |     |     |
|-----|-----|-----|-----|
| 1;1 | 1;2 | 1;3 | 1;4 |
| 2;1 | 2;2 | 2;3 | 2;4 |
| 3;1 | 3;2 | 3;3 | 3;4 |
| 4;1 | 4;2 | 4;3 | 4;4 |

## Matrix[3;]

|     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1;1 | 1;2 | 1;3 | 1;4 | 2;1 | 2;2 | 2;3 | 2;4 | 3;1 | 3;2 | 3;3 | 3;4 | 4;1 | 4;2 | 4;3 | 4;4 |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|

## Matrix[:,3]

|     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 1;1 | 1;2 | 1;3 | 1;4 | 2;1 | 2;2 | 2;3 | 2;4 | 3;1 | 3;2 | 3;3 | 3;4 | 4;1 | 4;2 | 4;3 | 4;4 |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|

# A Wider Matrix

## Matrix

|     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|
| 1;1 | 1;2 | 1;3 | 1;4 | 1;5 | 1;6 | 1;7 | 1;8 |
| 2;1 | 2;2 | 2;3 | 2;4 | 2;5 | 2;6 | 2;7 | 2;8 |
| 3;1 | 3;2 | 3;3 | 3;4 | 3;5 | 3;6 | 3;7 | 3;8 |
| 4;1 | 4;2 | 4;3 | 4;4 | 4;5 | 4;6 | 4;7 | 4;8 |

## Matrix[3;]

|     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |     |            |            |            |            |            |            |            |            |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|------------|------------|------------|------------|------------|------------|------------|------------|-----|-----|-----|-----|-----|-----|-----|-----|
| 1;1 | 1;2 | 1;3 | 1;4 | 1;5 | 1;6 | 1;7 | 1;8 | 2;1 | 2;2 | 2;3 | 2;4 | 2;5 | 2;6 | 2;7 | 2;8 | <u>3;1</u> | <u>3;2</u> | <u>3;3</u> | <u>3;4</u> | <u>3;5</u> | <u>3;6</u> | <u>3;7</u> | <u>3;8</u> | 4;1 | 4;2 | 4;3 | 4;4 | 4;5 | 4;6 | 4;7 | 4;8 |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|------------|------------|------------|------------|------------|------------|------------|------------|-----|-----|-----|-----|-----|-----|-----|-----|

## Matrix[;3]

|     |     |            |     |     |     |     |     |     |     |            |     |     |     |     |     |     |     |            |     |     |     |     |     |     |     |            |     |     |     |     |     |
|-----|-----|------------|-----|-----|-----|-----|-----|-----|-----|------------|-----|-----|-----|-----|-----|-----|-----|------------|-----|-----|-----|-----|-----|-----|-----|------------|-----|-----|-----|-----|-----|
| 1;1 | 1;2 | <u>1;3</u> | 1;4 | 1;5 | 1;6 | 1;7 | 1;8 | 2;1 | 2;2 | <u>2;3</u> | 2;4 | 2;5 | 2;6 | 2;7 | 2;8 | 3;1 | 3;2 | <u>3;3</u> | 3;4 | 3;5 | 3;6 | 3;7 | 3;8 | 4;1 | 4;2 | <u>4;3</u> | 4;4 | 4;5 | 4;6 | 4;7 | 4;8 |
|-----|-----|------------|-----|-----|-----|-----|-----|-----|-----|------------|-----|-----|-----|-----|-----|-----|-----|------------|-----|-----|-----|-----|-----|-----|-----|------------|-----|-----|-----|-----|-----|

The wider the matrix, the more fragmented the data columns.

# Virtual Memory

| CACHE | RAM | Disk Swap                    |
|-------|-----|------------------------------|
| L1    |     | Seek Time (tracks)           |
| L2    |     | Rotational Latency (sectors) |
| L3    |     | Transfer Time                |

**For a very large matrix**

**An APL data row is concentrated in the cache, or a few pages of RAM, or perhaps a few sectors of disk.**

**An APL data column is scattered among a few locations in the cache, and many pages of RAM, and perhaps many sectors or tracks of disk.**

# AIX APL Run Time

```
N←10000
A←(N,N)ρ⊞AV      a 1000 by 1000 matrix
:For I :In ⍳10000
  J←N?N
  K←?N
  L←⊞AV[?256]
  A[J;K]←L      a Random items J of column K  9492 CPU ms
  A[K;J]←L      a Random items J of row K      725 CPU ms
:EndFor
```

**For a 1,000 by 1,000 character matrix, random row access is on the average more than 10 times faster than random column access.**



# Another Column Data Fragmentation Example

```
Y←1000 1000ρ0.1+ι1000
```

```
:For I :In ι1000
```

```
  {}+//Y      A Sum matrix columns first 18277 CPU ms
```

```
  {}+//Y      A Sum matrix rows first    8178 CPU ms
```

```
  {}+/,Y      A Sum vector              6445 CPU ms
```

```
:EndFor
```

# Platform & Version Dependency

Although these observations on CPU times are generally true on all platforms, the exact performance varies depending on the hardware configuration, especially the quantity and quality of cache relative to the size of the arrays being manipulated. The performance might vary in the future versions of APL.

# Matrix Organization Suggestion

When you design a large data matrix for an APL application, always ask this question.

“Will I be accessing the data columns more frequently than accessing the data rows?”

If the answer is yes, then you should consider redesigning your data matrix such that it is in the transposed format so that you will access consecutive memory more frequently.

If some of your virtual matrix data are on a disk, consecutive disk sectors can greatly speed up your program.

# Marketing Data Example

Assuming we have the data of a million customers or prospects, and each customer has 100 data attributes, the first 5 of which are customer ID, country code, # of employees, industry code, and revenue.

We construct the following large random data matrix for testing.

```
ID←ι1010000
```

```
CTY←((1000000ρι80),10000ρ85)[1010000?1010000]
```

```
EMP←(1010000ρ(5ρ1000),(10ρ500),(20ρ200),50ρ100)[1010000?1010000]
```

```
IND←(1010000ρι9999)[1010000?1010000]
```

```
REV←(10100000+1000000×ι1010000)[1010000?1010000]
```

```
DAT←100↑[2]ID,CTY,EMP,IND,[1.1]REV
```



# Marketing Data Selection Example

To select companies with more 1000 employees in the manufacturing industry in country 85, most APL programmers would use a one-liner Boolean logic to get the data indices.

```
MFG←2011 2012 2013 2014 2015
:For I :In ⍳100
    J←((DAT[;3]≥1000)^(DAT[;4]∈MFG)^(DAT[;2]=85))/⍳1↑ρDAT
    ⍎ 28733 ms CPU
:EndFor
```



# Marketing Data Selection Example

## Data Organized by Columns

```
MFG←2011 2012 2013 2014 2015
:For I :In 100
    J←((DAT[;3]≥1000)^(DAT[;4]∈MFG)^(DAT[;2]=85))/11↑ρDAT
    ⌞ 28,733 ms CPU
:EndFor
```

## Data Organized by Rows

```
:For I :In 100
    J←((DAT[3;]≥1000)^(DAT[4;]∈MFG)^(DAT[2;]=85))/11↑ρDAT
    ⌞ 12,227 ms CPU
:EndFor
```



# Progressive Indexing

## Data Organized by Rows

```
:For I :In 100
  J←(DAT[3;]≥1000)/1↑ρDAT      A   2,393 ms CPU
  J←(DAT[4;J]∈MFG)/J          A     612 ms CPU
  J←(DAT[2;J]=85)/J           A       5 ms CPU
:EndFor                        A   3,000 ms CPU Total
```

In multi-line progressive indexing,

Expression 2: DAT[4;J]∈MFG works with fewer items than  
expression 1: DAT[3;]≥1000,

Expression 3: DAT[2;J]=85 works with fewer items than  
expression 2: DAT[4;J]∈MFG.

# Progressive Indexing Optimization

Suppose we know that there are not too many country 85 records in the data matrix, and there are hardly any big companies in this country. A better filtering sequence would be:

```
:For I :In 1100
  J←(DAT[2;]=85)/1↑ρDAT      A      1782 ms CPU
  J←(DAT[3;J]≥1000)/J       A           9 ms CPU
  J←(DAT[4;J]∈MFG)/J       A           4 ms CPU
:ENDFOR                      A      1795 ms CPU total
```

# Marketing Data Selection Speed Comparisons

```
:For I :In 100      A Column Data
  J←((DAT[ ;3]≥1000)^(DAT[ ;4]∈MFG)^(DAT[ ;2]=85))/1↑ρDAT
:EndFor                                                    A 28,733 ms
```

```
:For I :In 100      A Row Data
  J←((DAT[3 ; ]≥1000)^(DAT[4 ; ]∈MFG)^(DAT[2 ; ]=85))/1↑ρDAT
                                                    A 12,227 ms
```

```
J←(DAT[3 ; ]≥1000)/1↑ρDAT      A 2,393 ms
J←(DAT[4 ;J]∈MFG)/J           A 612 ms
J←(DAT[2 ;J]=85)/J           A 5 ms
                                                    A 3,000 ms Total
```

```
J←(DAT[2 ; ]=85)/1↑ρDAT      A 1782 ms
J←(DAT[3 ;J]≥1000)/J         A 9 ms
J←(DAT[4 ;J]∈MFG)/J         A 4 ms
                                                    A 1,795 ms total
```

```
:EndFor
```

# Row Major Progressive Indexing

With proper arrangement of the Boolean statements such that most of the unwanted data are filtered out by the first and second statements, for very large matrices, progressive indexing can be many times faster than the single Boolean statement in performing data selection. In the previous example, we see the function is 16 times faster when progressive indexing is performed on row-major data.

# Index Manipulation

It is usually more efficient to manipulate a matrix inside the square brackets [ ] than outside the [ ].

```
Y←1000 100ρ□A
T←1000ρ1 0 0
:For I :In 1100000
  {} 4φY[;68+ι8]          A 6006 ms
  {} Y[;4φ68+ι8]         A 2110 ms

  {} 0 1↓Y[;ι7]          A 4448 ms
  {} Y[;1↓ι7]            A 2049 ms

  {} T/Y[;ι6]            A 2330 ms
  {} Y[T/ι1↑ρY;ι6]      A 982 ms

  {} Y[;1 3],Y[;2 4]     A 6359 ms
  {} Y[;1 3 2 4]         A 1561 ms
:EndFor
```

# Grade Up & Grade Down Index

It is usually more efficient to manipulate the grade up index or the grade down index than to manipulate the sorted matrix.

```
R←X ΔSUM Y;J;P
```

```
[1] X←X[⊠X[;Y];]
```

⌘ Sort X on columns Y

```
[2] P←{1=ρρω:1,(1↓ω)≠-1↓ω
```

⌘ Unique vector items

```
[3] 1,ν/(1↓[1]ω)≠-1↓[1]ω}X[;Y]
```

⌘ Unique matrix rows

```
[4] R←P/X
```

⌘ Rows with unique columns Y

```
[5] ...
```

can be rewritten as

```
R←X ΔSUM2 Y;I;J;P
```

```
[1] I←⊠X[;Y]
```

⌘ Sort X on column Y

```
[2] P←{1=ρρω:1,(1↓ω)≠-1↓ω
```

⌘ Unique vector items

```
[3] 1,ν/(1↓[1]ω)≠-1↓[1]ω}X[I;Y]
```

⌘ Unique matrix rows

```
[4] R←X[P/I;]
```

⌘ Rows with unique columns Y

```
[5] ...
```

# Speed and Space Comparisons

```
M←(10000ρι100),10000 100ρι1000000
:For I :In ι10000
  {}M ΔSUM 1      A 20647 ms
  {}M ΔSUM2 1     A 4542 ms
:EndFor
```

□WA

4044960

{ }DATA ΔSUM 1

WS FULL

ΔSUM[1] X←X[ΔX[;Y];]

^

□WA

121960

{ }DATA ΔSUM2 1

A Much less storage requirement

A NO WS FULL

# Array Dimensions

How many columns are there in matrix M?

$\rho M[1; ]$  should be coded as  $\sim 1 \uparrow \rho M$

How many rows are there in matrix M?

$\rho M[ ; 1 ]$  should be coded as  $1 \uparrow \rho M$

How many planes and columns are there in the 3-dimensional array A?

$\rho A[ ; 1 ; ]$  should be coded as  $(\rho A)[1 \ 3]$



# Conclusion

When you perform your data selection on a large matrix using carefully constructed progressive indexing instead of a simple APL one-liner Boolean logic, the run time can be reduced significantly. If the data to be progressively indexed are arranged in rows instead of in columns, the run time can be reduced even more.

In optimizing matrix indexing, we need to pay special attention to the ones that are in the innermost loop of a function. Optimizing the matrix indexing deep inside an intensive loop would give you the maximum benefits.