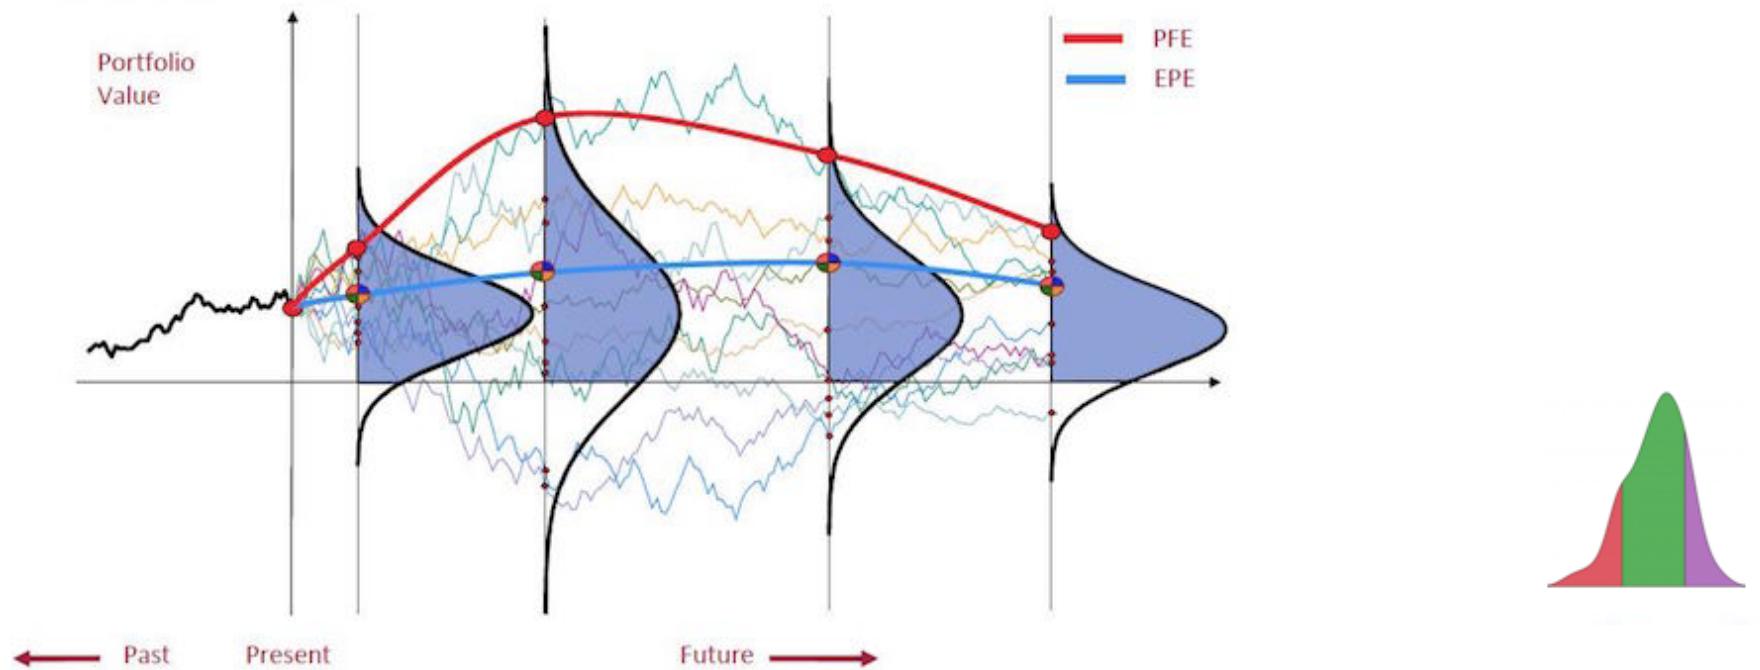


# Julia as an HPC alternative for Quant finance

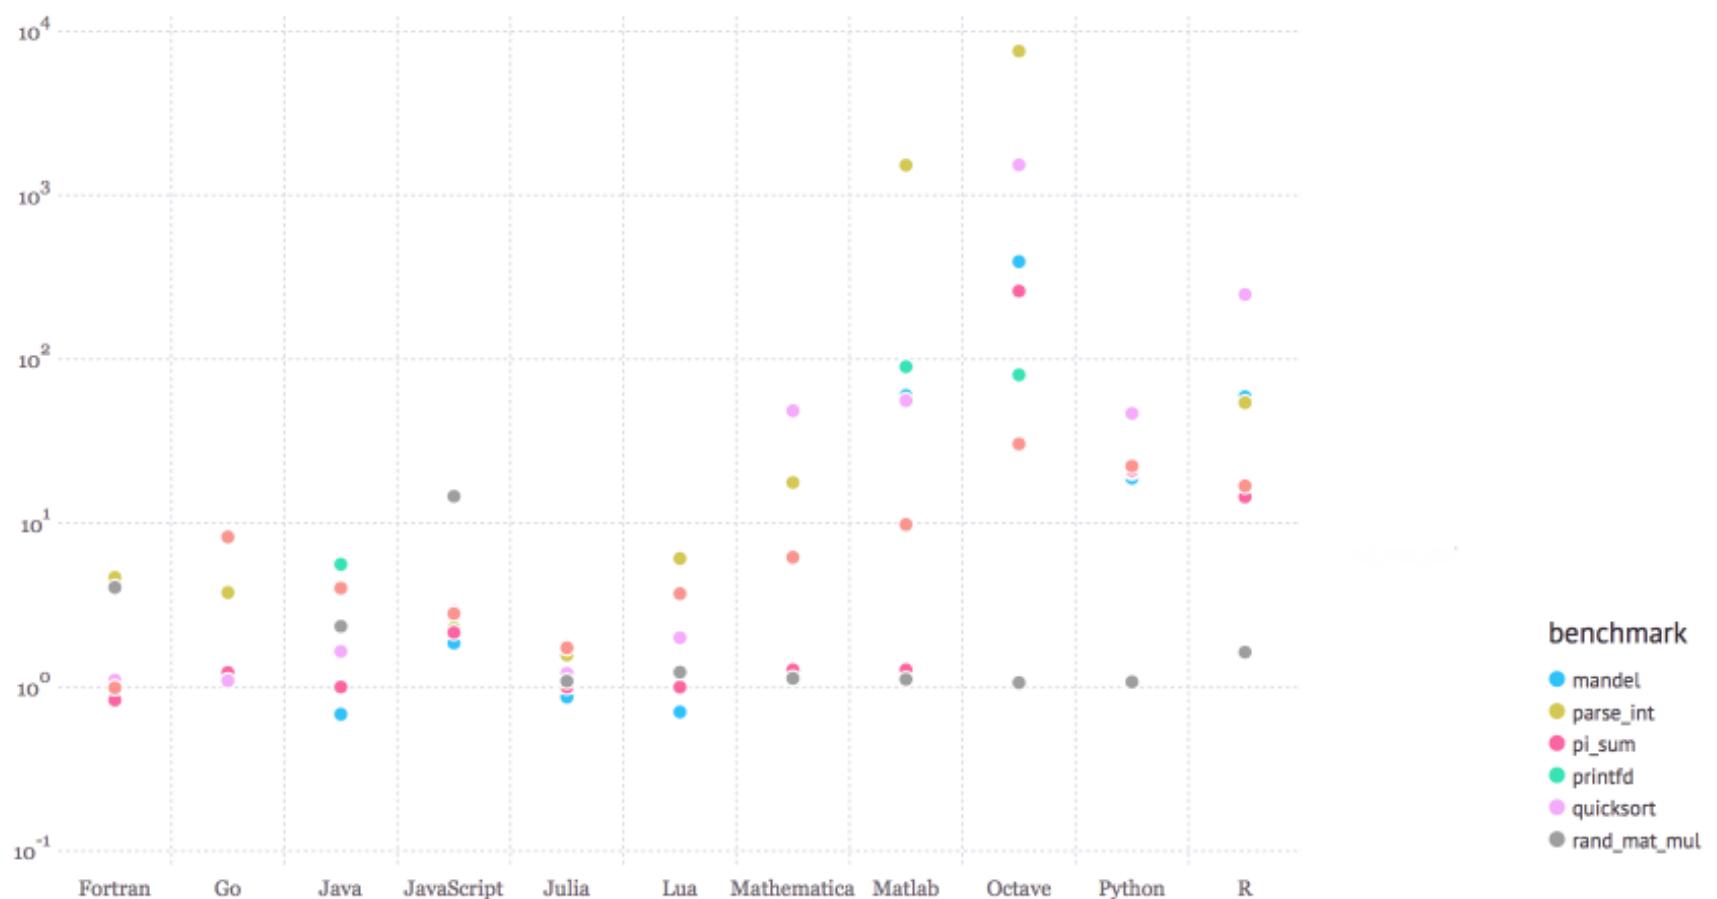


## Breakdown of the talk

- Main features of Julia
- Derivative pricing
- Parallel processing
- Julia Community Groups
- High Performance Computing
- JuliaQuant
- Time series analysis
- Heteroskedasticity

# What make Julia different ?

- Julia is written in Julia.
- Uses LLVM / JIT compilation, so it is fast !!!
- Code is uncluttered, basic functions are built in.
- Homoiconic design: provides runtime macros.
- Easy to execute on multicore and parallel processors.
- Can connect with modules / libraries written in other languages
- Can spawn tasks and interact with other programs



In [1]:

```
using PyPlot
T = 100;
S0 = 100;
dt = 0.01;
v = 0.2;
r = 0.05;
q = 0.0;
```

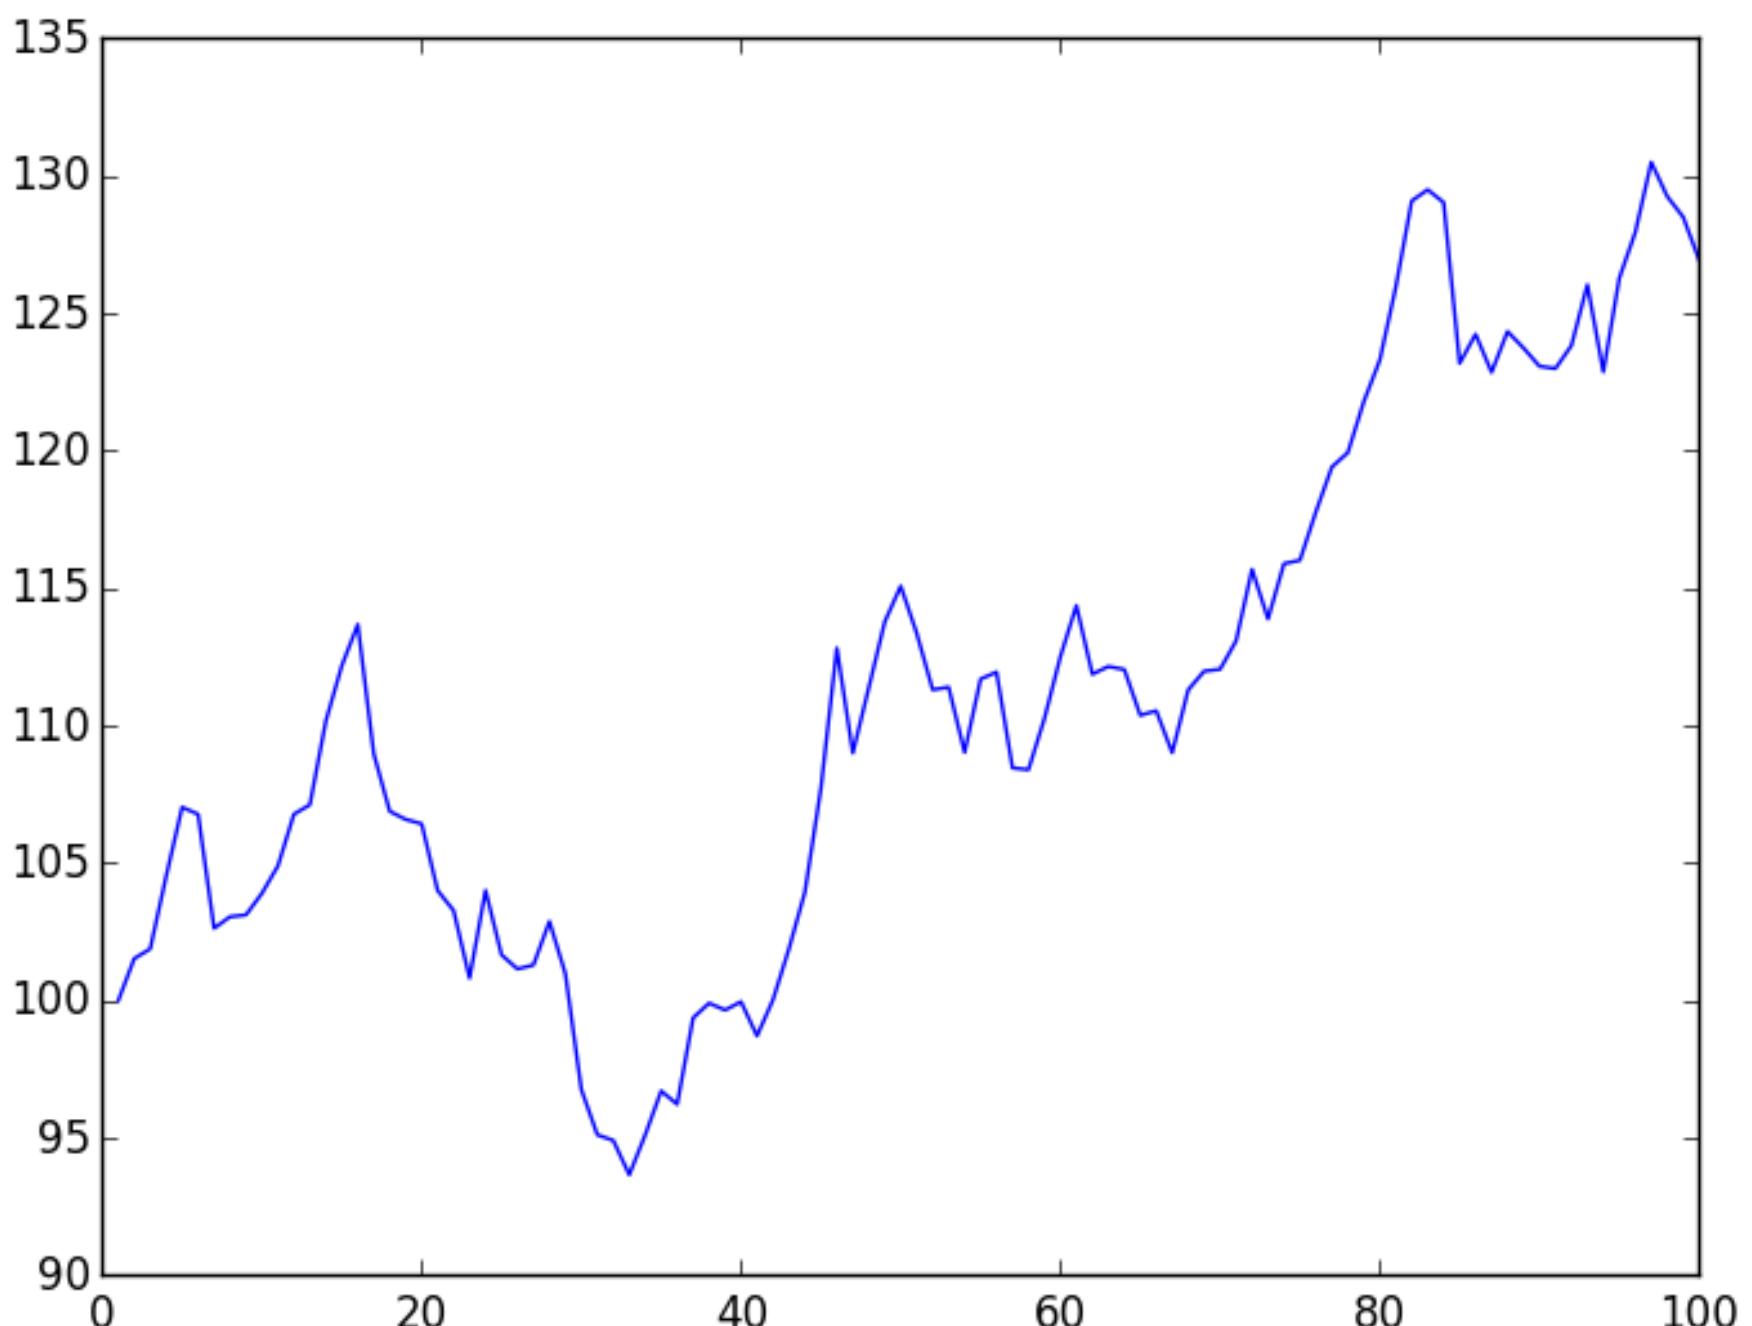
INFO: Loading help data...

In [6]:

```
S = zeros(Float64,T)
S[1] = S0;
iseed = ccall( (:clock, "libc"), Int32, () );
srand(iseed);

dW = randn(T)*sqrt(dt);
[ S[t] = S[t-1]*(1 + (r - q - 0.5*v*v)*dt + v*dW[t] + 0.5*v*v*dW[t]*dW[t]) for
t = 2:T ]

x = linspace(1,T);
plot(x,S)
```



Out[6]:

```
1-element Array{Any,1}:
PyObject <matplotlib.lines.Line2D object at 0x11c26e2d0>
```

In [7]:

```
function asianOpt(N::Int64, T::Int64; S0=100.0, K=100.0, r=0.05, q=0.0, v=0.1, tma=0.25)

# European Asian option.
# Euler and Milstein discretization for Black-Scholes.

@assert N > 0;
@assert T > 0;

dt = tma/T;
S = zeros(Float64,T);
A = zeros(Float64,N);

for n = 1:N
    S[1] = S0
    dW = randn(T)*sqrt(dt);
    for t = 2:T
        z0 = (r - q - 0.5*v*v)*S[t-1]*dt;
        z1 = v*S[t-1]*dW[t];
        z2 = 0.5*v*v*S[t-1]*dW[t]*dW[t];
        S[t] = S[t-1] + z0 + z1 + z2;
    end
    A[n] = mean(S);
end

# Define the payoff and calculate price

P = zeros(Float64,N);
[ P[n] = max(A[n] - K, 0) for n = 1:N ];
return exp(-r*tma)*mean(P);

end
```

Out[7]:

```
asianOpt (generic function with 1 method)
```

In [8]:

```
price = asianOpt(100000,100; K=102.0,v=0.2);
@printf "Option Price: %10.4f\n\n" price;
```

```
Option Price:      1.6638
```

In [9]:

```
@elapsed asianOpt(100000,100;K=102.0,v=0.2)
```

Out[9]:

```
0.374801638
```

# Comparison with Python, R, Matlab etc.

	C	Python (v3)	R	Octave	Julia	Java	Javascript
Timings	1	32.67	154.3	789.3	1.41	2.97	4.22
Price	1.681	1.671	1.646	1.632	1.680	1.663	1.664

In [10]:

```
# global CPU_CORES  (sysinfo.jl : line 22 )
#
ncores = ccall(:jl_cpu_cores, Int32, ())
```

Out[10]:

4

In [11]:

```
addprocs(3);
```

In [12]:

```
n = nprocs()
```

Out[12]:

4

In [13]:

```
;cat Exotics.jl

module Exotics

using Yeppp
yp = Yeppp

function asianOpt(N::Int64, T::Int64; S0=100.0, K=100.0, r=0.05, q=0.0, v=0.1,
tma=0.25)

# European Asian option.
# Euler and Milstein discretization for Black-Scholes.

@assert N > 0;
@assert T > 0;

dt = tma/T;
S = zeros(Float64,T);
A = zeros(Float64,N);

for n = 1:N
```

```

S[1] = S0
dW = randn(T)*sqrt(dt);
for t = 2:T
    z0 = (r - q - 0.5*v*v)*S[t-1]*dt;
    z1 = v*S[t-1]*dW[t];
    z2 = 0.5*v*v*S[t-1]*dW[t]*dW[t];
    S[t] = S[t-1] + z0 + z1 + z2;
end
A[n] = mean(S);
end

# Define the payoff and calculate price

P = zeros(Float64,N);
[ P[n] = max(A[n] - K, 0) for n = 1:N ];
return exp(-r*tma)*mean(P);

end

function asianOptYep(N::Int64, T::Int64; S0=100.0, K=100.0, r=0.05, q=0.0, v=0.1, tma=0.25)

# European Asian option.
# Euler and Milstein discretization for Black-Scholes.

@assert N > 0;
@assert T > 0;

dt = tma/T;
S = zeros(Float64,T);
A = zeros(Float64,N);
W = Array(Float64,N*T);

k = [1.0 + (r - q - 0.5*v*v)*dt, v, 0.5*v*v]
dW = randn(N*T)*sqrt(dt);
yp.evalpoly!(W,k,dW);

for n = 1:N
    S[1] = S0
    for t = 2:T
        i = (n-1)*T + t;
        S[t] = S[t-1]*W[i];
    end
    A[n] = mean(S);
end

# Define the payoff and calculate price

P = zeros(Float64,N);
[ P[n] = max(A[n] - K, 0) for n = 1:N ];
return exp(-r*tma)*mean(P);

end

end

```

In [14]:

```
import Exotics
@time Exotics.asianOpt(1000000,100; K=102.0,v=0.2)
```

elapsed time: 3.708321446 seconds (1818867656 bytes allocated, 41.92% gc time)

Out[14]:

1.675951151108044

In [15]:

```
@everywhere using Exotics
```

In [16]:

```
tic();
@everywhere price = Exotics.asianOpt(250000,100; K=102.0,v=0.2);
totprice=0.0;
for i = 1:n
    totprice += remotecall_fetch(i,()>price)
end
totprice = totprice/n;
toc()

@printf "Asian price is %7.4f\n" totprice
```

elapsed time: 1.952977148 seconds

Asian price is 1.6728

## Julia Community Groups:

- JuliaStats
- JuliaOpt
- JuliaWeb
- JuliaQuant
- JuliaFinMetrix
- QuantEcon
- JuliaParallel
- JuliaGPU

## JuliaStats group :

- StatsBase
- Distributions
- MultivariateStats
- KernelDensity / KernelEstimator
- Dataframes
- RDatasets
- GLM
- Lora / MCMC
- Clustering

---

## JuliaParallel group :

- ClusterManagers
- Elly, HDFS
- MPI
- Yeppp

In [17]:

```
using Yeppp
@time Exotics.asianOptYep(1000000,100; K=102.0,v=0.2)
```

elapsed time: 2.732427 seconds (2426614384 bytes allocated, 4.00% gc time)

Out[17]:

```
1.676402559571753
```

---

## JuliaGPU group :

- CUDA, CUDArt, CUBLAS, CUFFT
- OpenCL, CLBLAS, CLFFT

# Kepler GK110

## Architecture

- 7.1B Transistors
- 15 SMX units
- > 1 TFLOP FP64
- 1.5 MB L2 Cache
- 384-bit GDDR5
- PCI Express Gen3

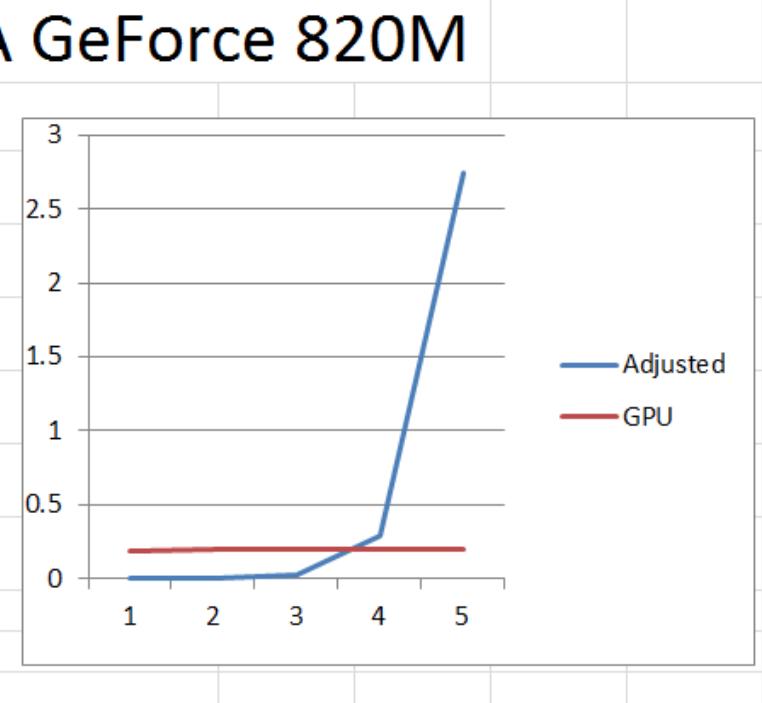


In [ ]:

```
;cat OCL/saxpy.jl
```

## OpenCL on Lenovo i7 / NVIDIA GeForce 820M

<i>Runs</i>	<i>Julia</i>	<i>Adjusted</i>	<i>GPU</i>
1000	0.0178	0.0002	0.192
10000	0.0201	0.0025	0.196
100000	0.0439	0.0263	0.195
1000000	0.3061	0.2885	0.197
10000000	2.7565	2.7389	0.196



In [18]:

```
;cat OCL/vaxbyz.jl
```

```

import OpenCL
cl = OpenCL

const vaxbyz_kernel =
__kernel void vaxbyz(
    __global const float *a,
    __global const float *b,
    __global const float *x,
    __global const float *y,
    __global const float *z,
    __global float *s)
{
    int gid = get_global_id(0);
    s[gid] = a[gid]*x[gid]*x[gid] + b[gid]*y[gid] + z[gid];
}
"

for n in (1, 1000, 50_000, 1_000_000, 5_000_000, 25_000_000, 50_000_000)

a = rand(Float32, n);
b = rand(Float32, n);
x = rand(Float32, n);
y = rand(Float32, n);
z = rand(Float32, n);
s = Array(Float32, n);

device, ctx, queue = cl.create_compute_context();

a_buff = cl.Buffer(Float32, ctx, (:r, :copy), hostbuf=a);
b_buff = cl.Buffer(Float32, ctx, (:r, :copy), hostbuf=b);
x_buff = cl.Buffer(Float32, ctx, (:r, :copy), hostbuf=x);
y_buff = cl.Buffer(Float32, ctx, (:r, :copy), hostbuf=y);
z_buff = cl.Buffer(Float32, ctx, (:r, :copy), hostbuf=z);
s_buff = cl.Buffer(Float32, ctx, :w, n)

p = cl.Program(ctx, source=vaxbyz_kernel) |> cl.build!
k = cl.Kernel(p, "vaxbyz");

cl.call(queue, k, size(x), nothing, a_buff, b_buff, x_buff, y_buff, z_buff, s_buff);
t0 = @elapsed s0 = cl.read(queue, s_buff);

s1 = zeros(n);
t1 = @elapsed begin
    for i = 1:length(x)
        s1[i] = a[i]*x[i]*x[i] + b[i]*y[i] + z[i]
    end
end
if n > 1
    @printf "%10d : %8.5f : %8.5f\n" n t0 t1;
else
    @printf "      Loops : GPU : Native\n";
end
end
println();

```

In [19]:

```
run(`./vaxbyz.sh`)
```

```
100000 : 0.21780 : 0.05390
500000 : 0.21689 : 0.31798
1000000 : 0.21892 : 0.52458
5000000 : 0.27383 : 2.73475
10000000 : 0.34168 : 5.90696
25000000 : 0.53296 : 15.33412
```

---

## JuliaQuant group :

- TimeSeries
- MarketData
- MarketTechnicals
- Quandl
- Ito
- GARCH
- FinancialAssets
- Grist
- TradeModels
- PortfolioModels

In [20]:

```
## FinancialSeries.jl api demo
using MarketData, PyPlot
```

In [21]:

```
#using FinancialSeries
include("./FS/FinancialSeries.jl")
```

In [22]:

```
# example stock from MarketData
AAPL.colnames
```

Out[22]:

```
12-element Array{UTF8String,1}:
 "Open"
 "High"
 "Low"
 "Close"
 "Volume"
 "Ex-Dividend"
 "Split Ratio"
 "Adj. Open"
 "Adj. High"
 "Adj. Low"
 "Adj. Close"
 "Adj. Volume"
```

In [23]:

```
# construct financial time series
Apple = TimeArray(AAPL.timestamp, AAPL.values, AAPL.colnames, FinancialSeries.
Stock(FinancialSeries.Ticker("AAPL")))
```

Out[23]:

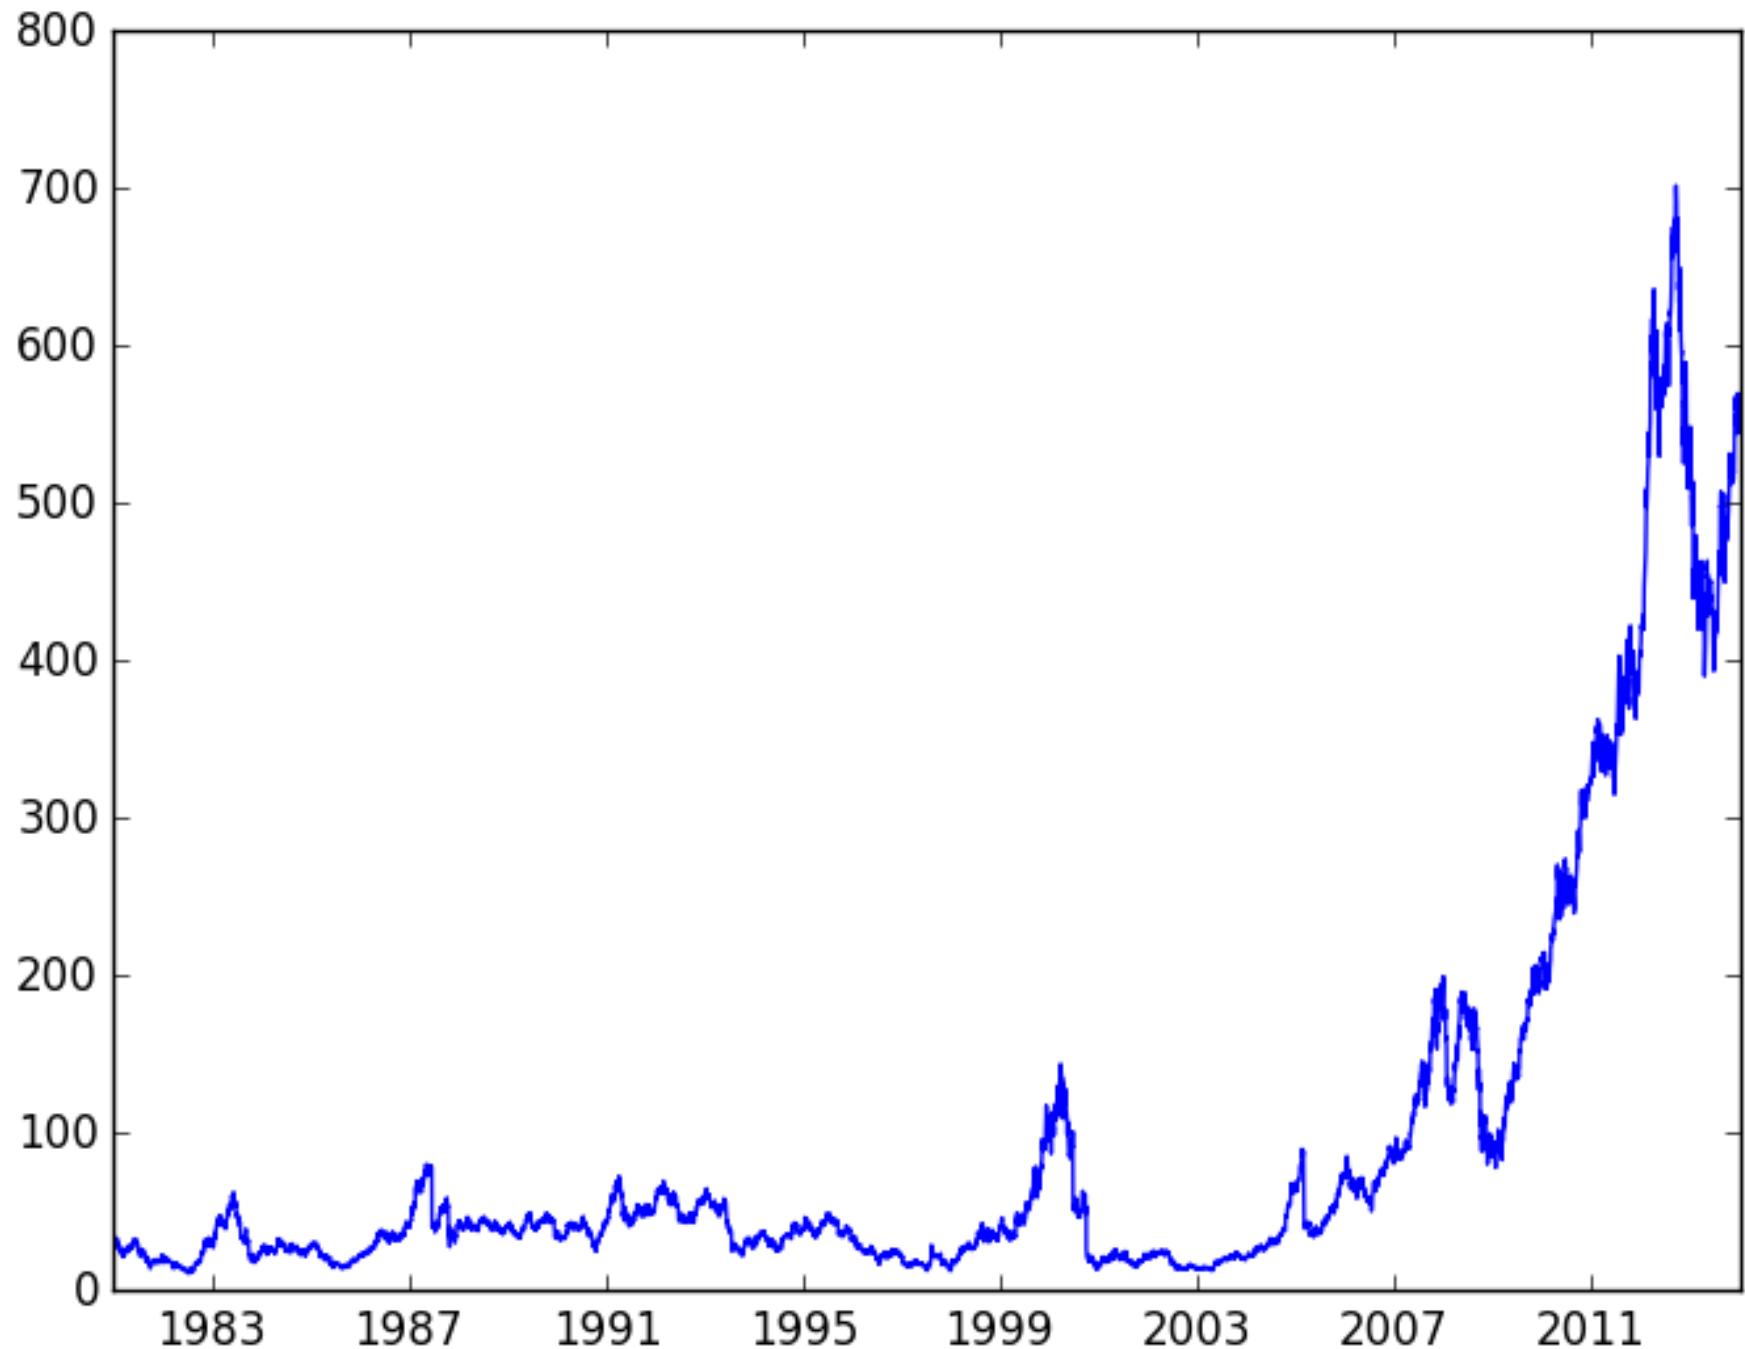
8336x12 FinancialTimeSeries for Stock, 1980-12-12 to 2013-12-31

```
ticker:          AAPL
currency:        USD
tick:            0.01
multiplier:      1.0
```

	Open	High	Low	Close	Volume	Ex-Dividend	Split R
ratio	Adj. Open	Adj. High	Adj. Low	Adj. Close	Adj. Volume		
1980-12-12   3.38	28.75	28.88	28.75	28.75	2093900	0.0	1
3.22	3.39	3.38	3.38		16751200		
1980-12-15   3.22	27.38	27.38	27.25	27.25	785200	0.0	1
2.98	3.22	3.2	3.2		6281600		
1980-12-16   2.98	25.38	25.38	25.25	25.25	472000	0.0	1
3.04	2.98	2.97	2.97		3776000		
1980-12-17   3.04	25.88	26.0	25.88	25.88	385900	0.0	1
:					3087200		
2013-12-26   564.74	568.1	569.5	563.38	563.9	7286000	0.0	1
560.48	566.13	560.05	560.56		7286000		
2013-12-27   560.48	563.82	564.41	559.5	560.09	8067300	0.0	1
554.16	561.07	556.19	556.78		8067300		
2013-12-30   554.16	557.46	560.09	552.32	554.52	9058200	0.0	1
550.89	556.78	549.05	551.24		9058200		
2013-12-31   550.89	554.17	561.28	554.0	561.02	7967300	0.0	1
					7967300		

In [24]:

```
ac = Apple["Close"].values;  
tm = timestamp(Apple);  
plot(tm,ac)
```

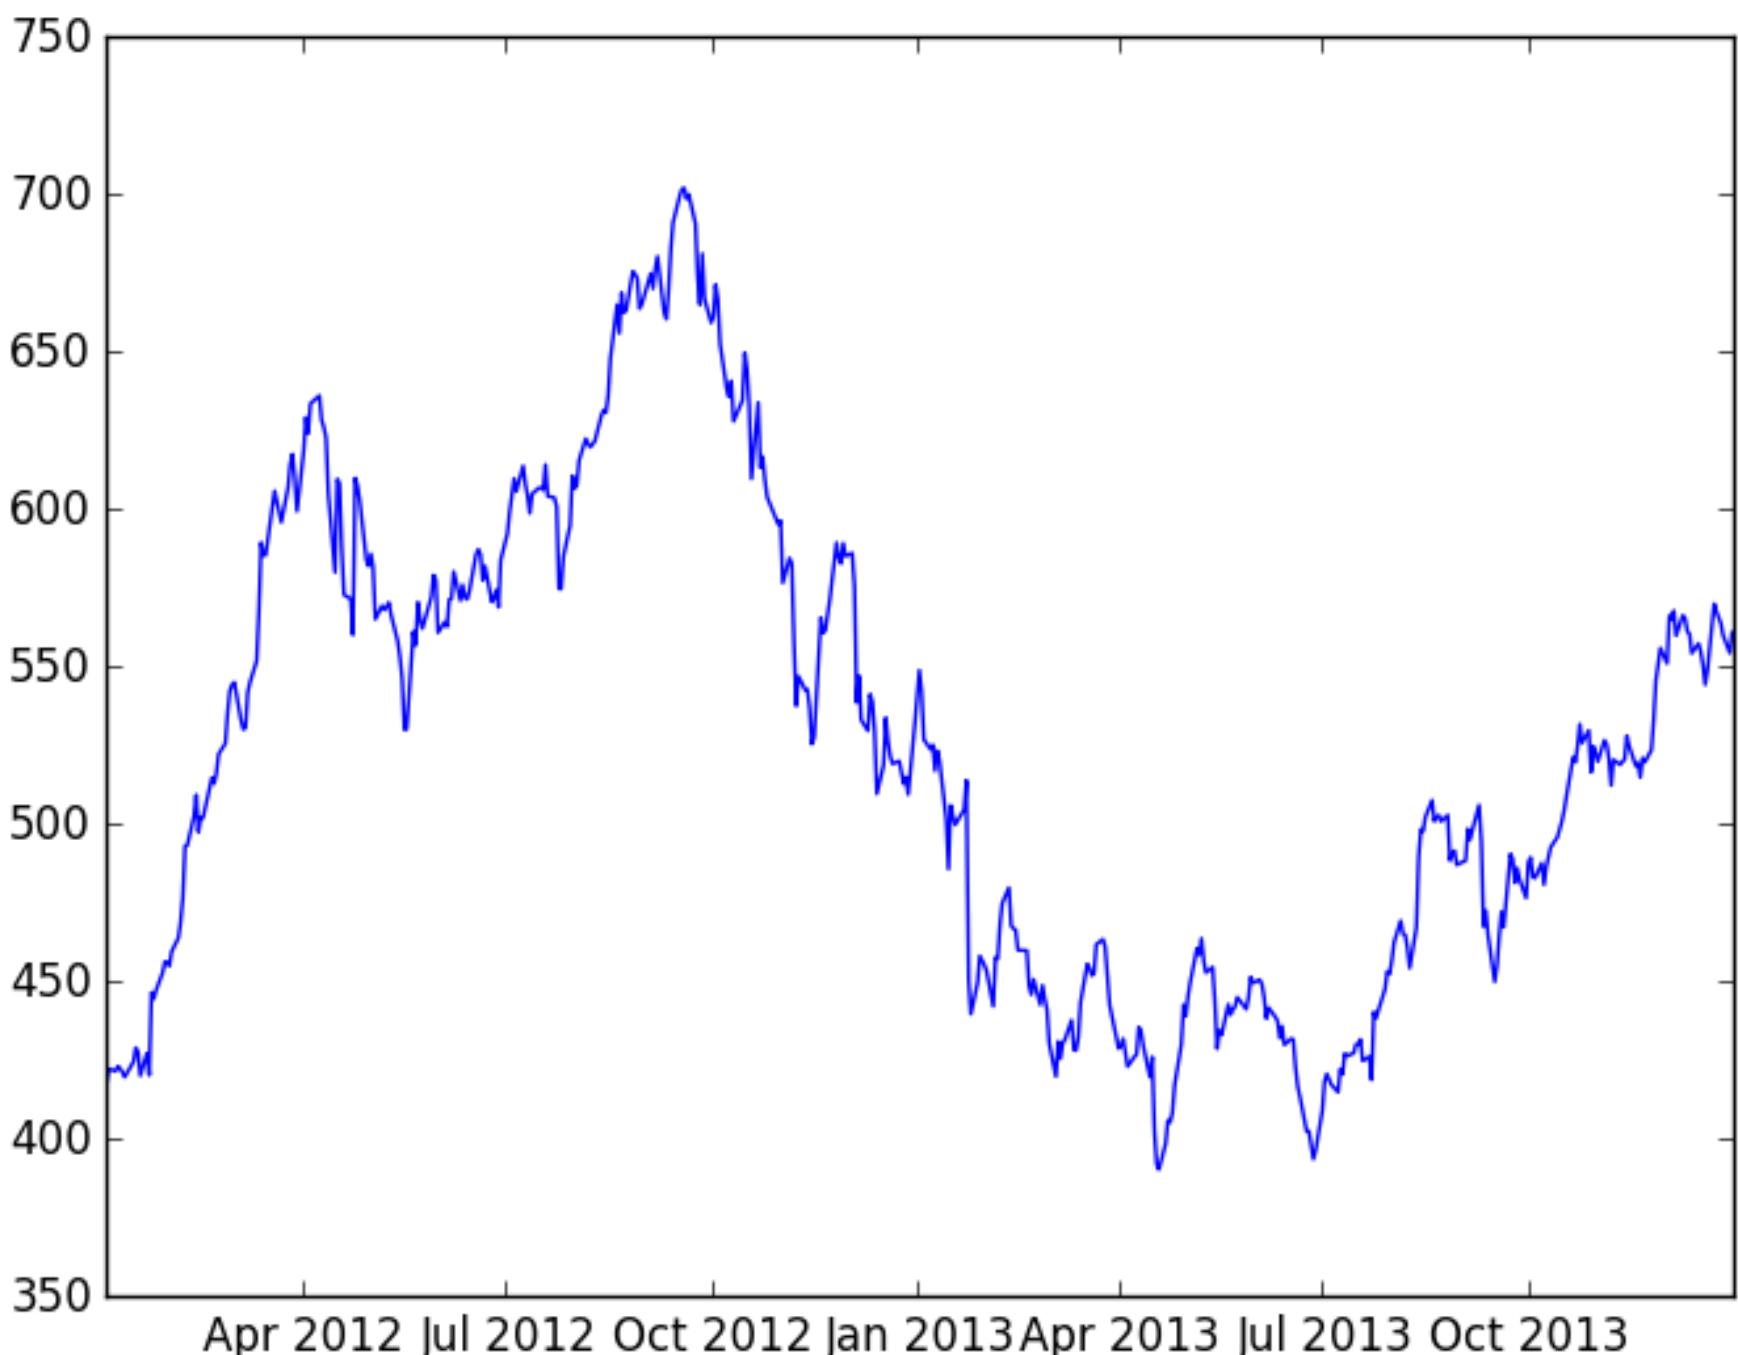


Out[24]:

```
1-element Array{Any,1}:  
PyObject <matplotlib.lines.Line2D object at 0x1aee8b250>
```

In [25]:

```
plot(tm[end-500:end], ac[end-500:end])
```

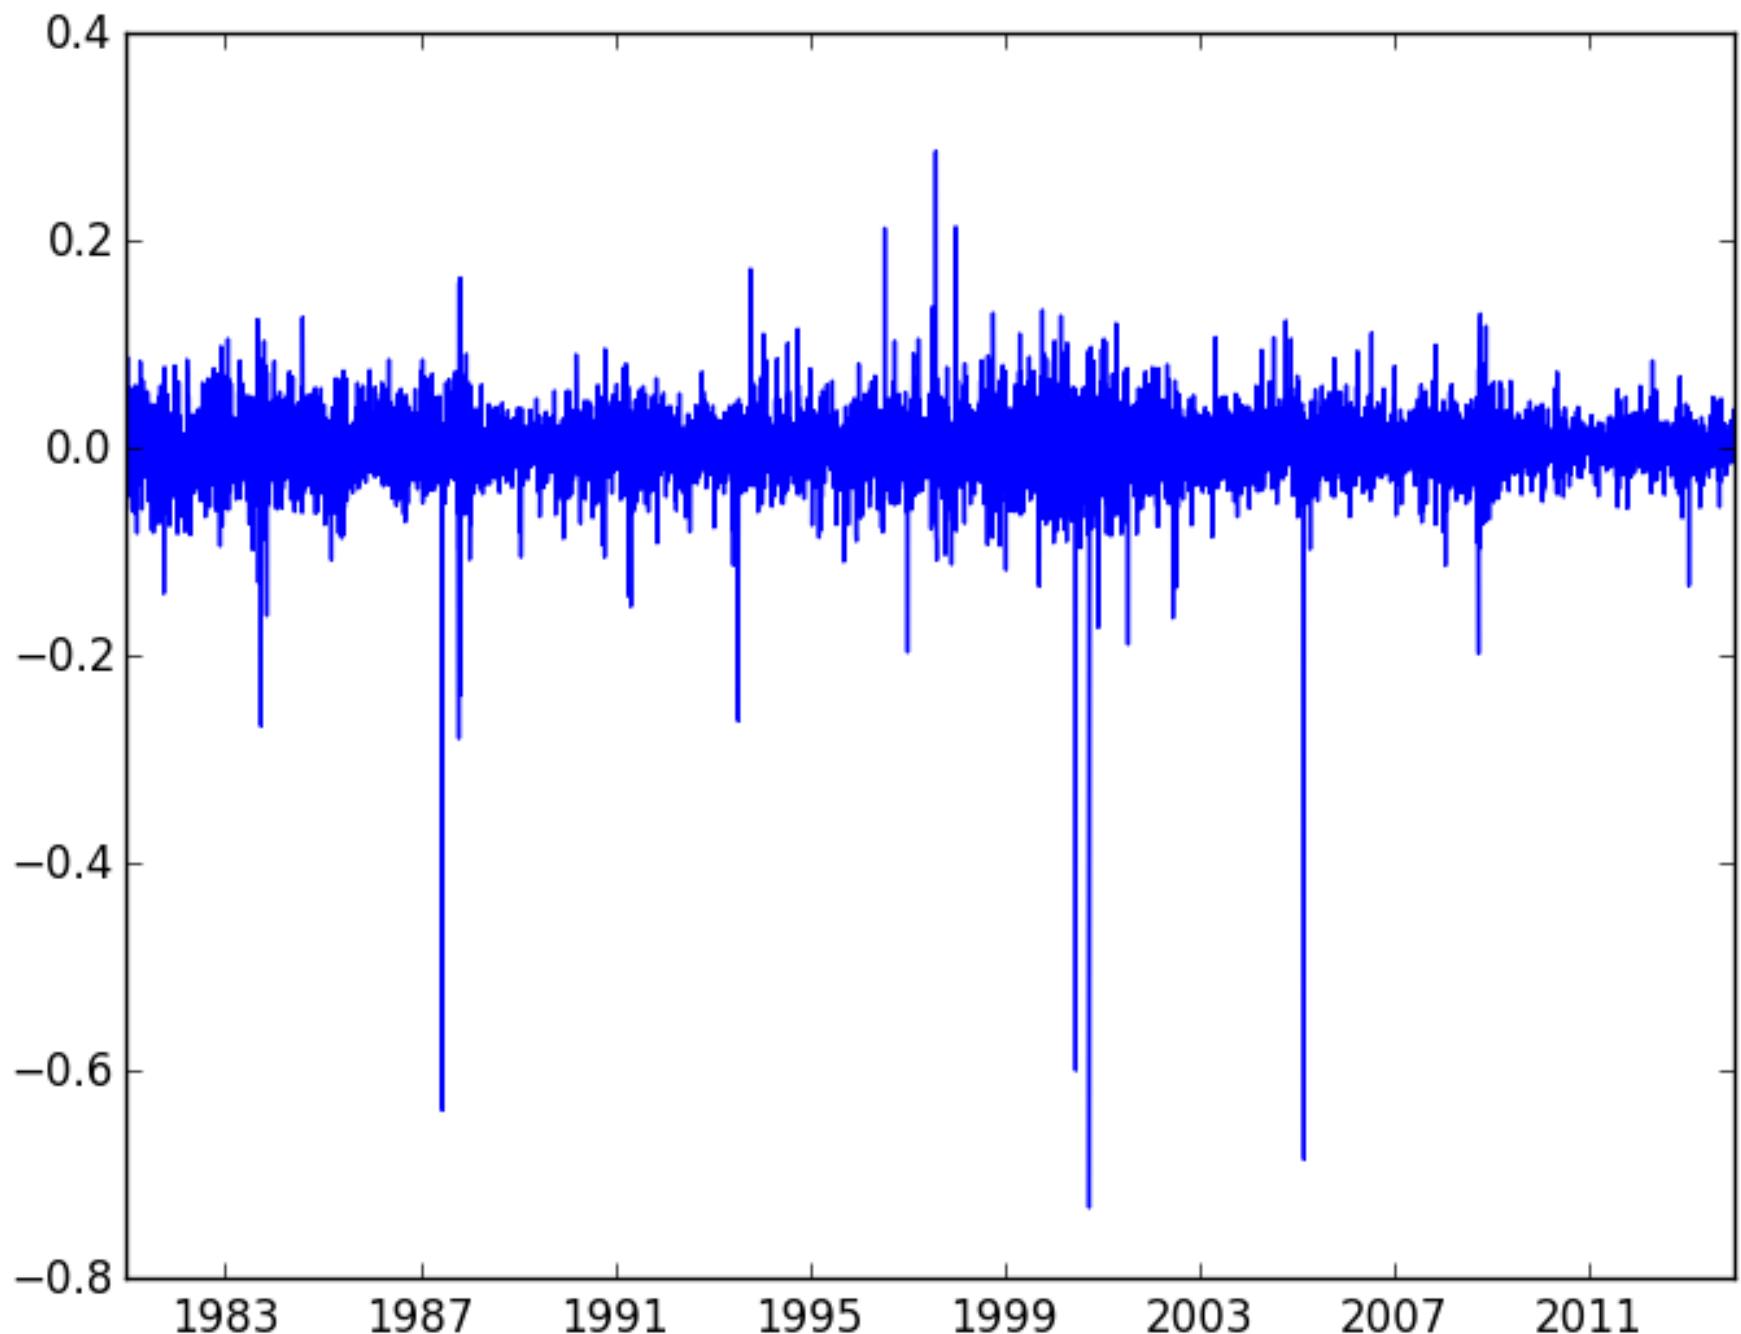


Out[25]:

```
1-element Array{Any,1}:
PyObject <matplotlib.lines.Line2D object at 0x1af19dad0>
```

In [26]:

```
pc = percentchange(Apple["Close"],method="log").values;
tm = timestamp(Apple)[2:end];
plot(tm,pc)
```



Out[26]:

```
1-element Array{Any,1}:
PyObject <matplotlib.lines.Line2D object at 0x1af4d9d10>
```

# Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) Process

Developed in 1982 to describe an approach to estimate volatility in financial markets.

The GARCH process is often preferred in financial modeling because it gives a more real-world context than other forms when predicting the prices and rates of financial instruments.

The general process for a GARCH model involves three steps:

- Estimate a best-fitting autoregressive model
- Compute autocorrelations of the error term
- Test for significance

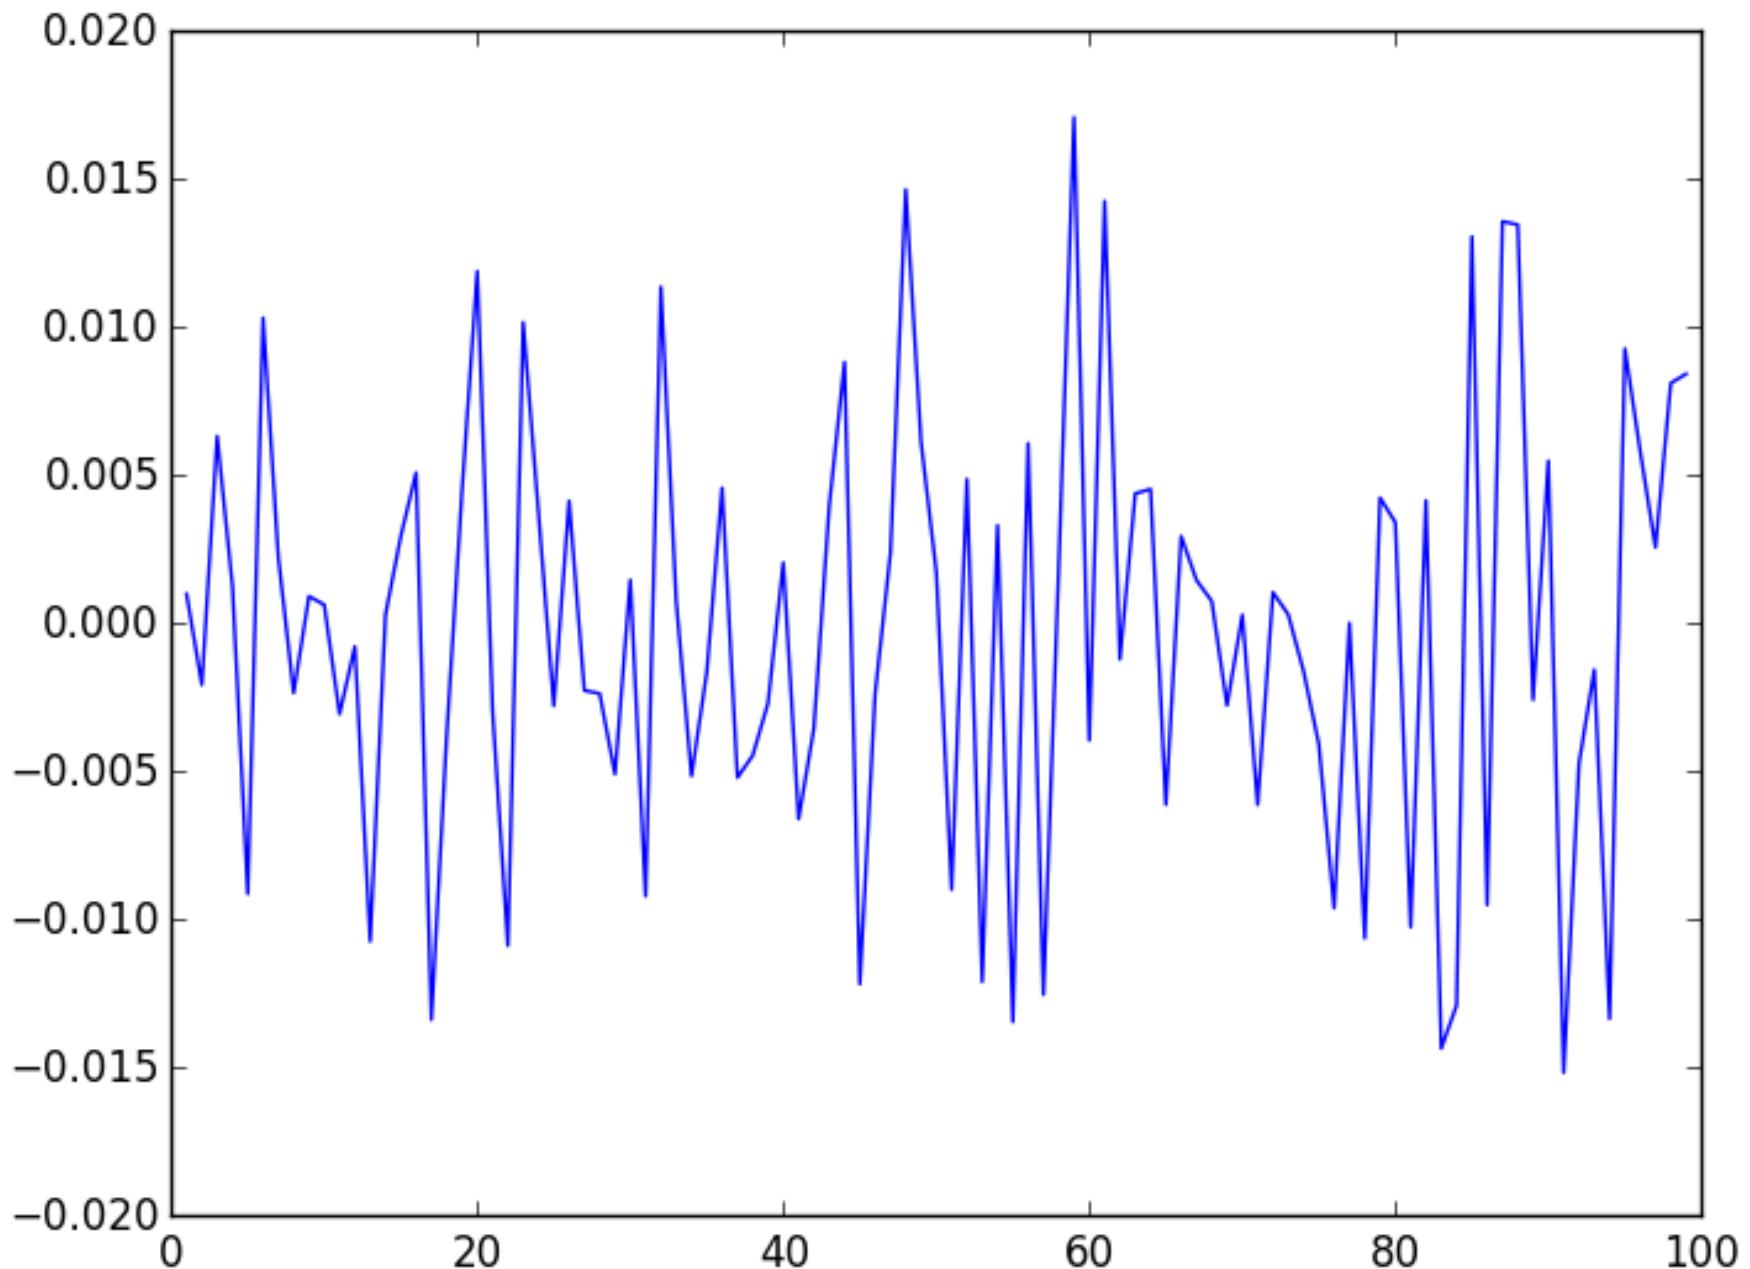
In [2]:

```
using Quandl  
quotes = quandl("YAHOO/INDEX_GSPC", format="DataFrame")
```

Out[2]:

In [3]:

```
using PyPlot
qac = diff(log(array(quotes[:Adjusted_Close])))
plot(1:length(qac),qac)
```



Out[3]:

```
1-element Array{Any,1}:
PyObject <matplotlib.lines.Line2D object at 0x11a0c3c90>
```

In [4]:

```
using GARCH  
garchFit(qac)
```

Out[4]:

```
Fitted garch model  
* Coefficient(s): omega alpha beta  
                  0.000000 0.014263 0.990000  
* Log Likelihood: 347.223248  
* Converged: false  
* Solver status: XTOL_REACHED  
  
* Standardised Residuals Tests:  
      Statistic p-Value  
Jarque-Bera Test  $\chi^2$  0.839373 0.657253  
  
* Error Analysis:  
      Estimate Std.Error t value Pr(>|t|)  
omega 0.000000 0.000003 0.005231 0.995826  
alpha 0.014263 0.012117 1.177084 0.239162  
beta 0.990000 0.079091 12.517164 0.000000
```

## In Summary

- Coding in Julia is simple
- It has (even now) impressive functionality
- This can be extended by using C/Fortran/Python/R/Java modules and libraries
- Julia has implemented data frames, similar to those in R and Python/Pandas
- Community groups exist and are active in statistics, finance, econometrics etc.
- JuliaQuant already covers many of the principal financial requirements