Project Number: MXC-WO75

Profit and Loss Workflow Analysis

A Major Qualifying Project Report

submitted to the Faculty

of the

WORCESTER POLYTECHNIC INSTITUTE

in partial fulfillment of the

Degree of Bachelor of Science

By

Angelo Chandler
Thomas EP MacDonald

Date: December 14th, 2007

Approved:

Dean John Orr, ECE advisor
Professor Germano S. Iannacchionne, Physics advisor
Professor Arthur Gerstenfeld, Co-advisor
Professor Michael J. Ciaraldi, Co-advisor

1. Profit and Loss Workflow Analysis
2. Financial Engineering
3. Risk Engine Hardware Analysis
4. Econophysics

Abstract

This project is a summary of a gap analysis that was performed on Bank of America – London's profit and loss reporting process. The timeliness of the reconciliation process for the reporting of profit and loss led to an in-depth study and analysis of various computer systems and manual processes. This project is a summary of the gap analysis and an outline for a reconciliation program that will help the profit and loss reporting teams.

Executive Summary

This project is based around a profit and loss workflow gap analysis performed at Bank of America in London. The profit and loss workflow gap analysis examines the two systems that currently report the profit and loss for their traders. This project explores the current problems such as the time for reconciliation of the profit and loss reports and improper reporting of the profit and loss. Solutions for these problems are based on input directly from the teams in London and from an analysis of the reporting system.

Additionally, this project features an in-depth explanation of financial terms that were necessary in order to be able to communicate with the employees at Bank of America. Each of these has been elaborated on to promote understanding for not only the reader but also the authors of this project. Econophysics is covered in this project to promote understanding of physical sciences being applied to the financial world. This includes a review of topics that use the properties of thermodynamics to explain financial phenomena.

This project also explores an upgrade to a current computing cluster called the Risk Engine. This cluster of computers is in charge for calculating risk, the Actual profit and loss, and the general ledger. This project investigates future options for the Risk Engine by performing a power, performance, and space required analysis and examines the possibility of completely switching over hardware to such options as a quad-core Intel-based system or a Cell-based system.

Acknowledgements

Angelo Chandler would like to acknowledge:

I would like to acknowledge my mother for her continued support throughout this wonderful experience. She has supported me throughout my nine week experience here in the United Kingdom. Next, I would like to thank Abhi for his support and dedication to helping us throughout the 'ups and downs' of our project. He educated us on the basics of the financial world, opened up doors and networking abilities within Bank of America, and also, most important, provided comfort within our personal workplace. I would like to thank sincerely Ben Cooper, Priscilla Wong, Alex May, and all of the helpful Bank of America employees that contributed in their own unique way to our project experience. Their continued dedication to helping us in our project was the key element that allowed us to complete this project with such success. Thank you Bank of America!

Across the ocean and back home at Worcester Polytechnic Institute, I would like to thank Professor Iannacchione for being a great Physics MQP advisor throughout the project. Without you as my advisor, my physics abilities and my love for economics would not have been supported in my MQP as a physics undergraduate student. Professor Gerstenfeld and Professor Ciarldi, I thank you for supporting all of the Wall Street MQP students, in New York and in London, throughout the entire process. You granted us an opportunity that many students across the world could only dream of! I want to thank everyone back home at WPI for supporting us throughout these nine weeks away, because without your support I personally would have not completed this project with a smile on my face. Thank you to everyone involved and I hope everyone had a great B term while we were away!

Thomas MacDonald would like to acknowledge:

First and foremost, I would like to acknowledge my family for their continued support both emotionally and financially. Abhi, thank you for being a great boss, you allowed us to go in our own direction and guided us through the hardest times. Ben Cooper, thank you for all of your help. Your continued dedication to our project made us feel like we were making a difference and that we were truly helping you and your team. Priscilla Wong, thank you for all of your time as it helped us understand the profit and loss reporting system that much more. Alex May, thank you for being a part of the project and helping us with the Flash specific solutions for London.

Provost Orr, thank you for your continued dedication to the project. I was once lost on my requirements but I feel as if I've done enough work to make WPI and you proud. Professor Ciaraldi, thank you for your input on our project along the way. You were a joy to work with. Professor Iannacchione, thank you and I look forward to discussing my Physics MQP options with you as you appear to be very open-minded when it comes to that area. Professor Gerstenfeld, thank you for this opportunity and selecting me for the Wall Street – London project center. It was a great learning experience and I feel that I've become a well rounded student because of it. Thank you to all of my friends, I really missed you all here but it was comforting that you all continued on and stayed strong. As always, I save the best for last as I would like to thank my girlfriend, Natasha, who has driven me through the entire project to do my best.

Table of Contents

1: Introduction

1.1: Financial

This project was prepared by Abhishek Vashishth in order to have a profit and loss (P&L) workflow gap analysis performed. It started out with three different P&L systems, the Flash P&L, the Actual P&L, and the Explain P&L. However, it was soon obvious that the Explain P&L and the Actual P&L was the same exact thing except under a different name. This enabled us to specifically focus on the Flash and Actual P&L and the systems that input information to these calculators.

There are 11 different categories of P&L that are computed by Bank of America in the Flash and Actual P&L processes. There is the credit default swap credit spread, credit default swap credit spread convexity, corporate bond spread, government bond, interest rate, carry, credit curve shift, new business trading, brokerage, edit, and recovery. Credit default swap credit spread P&L is one of the most important profit and loss categories and is a summary of the sales surrounding credit derivatives. This comes from when the traders mark their curves at the end of the day. Credit default swap credit spread convexity is a second derivative and is only significant when there are major moves in the credit curve.

Corporate bond spread is a calculation of how large of a gap there is between their corporate bonds to an equal government bond. The government bond P&L is how much money is made from investing in a government's debt. Interest rate P&L comes from how much money is collected from interest. Carry P&L is how much money is made when a trade is kept from the beginning of the day to the end of the day. Credit curve shift is how a movement in the credit

yield curve affected their investment by either producing profit or loss. New business P&L includes any trades that are performed on that day and also includes assignments. Brokerage summarizes the profit and loss of acting as a firm that allows stocks and bonds to be sold for clients. Edit P&L is when a person had to go in physically and fix a trade because it was incorrectly booked. Recovery P&L comes from when an entity defaults and the company either has to pay or receives payment for the default.

With this knowledge, the various systems and the importance of each form of P&L had also been established. Problems with the carry, recovery, brokerage, interest rate, corporate bond spread, government bonds, credit curve shift, and credit curve spread convexity were considered to have small inaccuracies. This held true except for credit curve spread convexity that affects the P&L when there are large changes in the yield of the credit spread. This was identified as important information that Odin was not calculating and did not take into account.

We also established the importance of having multiple systems that have to agree. This provides multiple checkpoints to find errors, make corrections, and provide multiple sources for reconciliation. The Flash's purpose is to provide an estimate of the profit and loss for any given day and is distributed to the traders and front office at the beginning of the day. Front office is the office that does all of the trading while middle office performs the reconciliation and provides a checkpoint for the traders. The Actual P&L is calculated by Risk Engine and is reconciled before given to the traders.

In this project, we show that there are multiple areas that can be improved upon. These areas have been outlined by taking the input from people who actually need the P&L reports and then

the solution was found by brainstorming and taking in the available information. We then provide our ideas and our solutions with applicable information included to support our solutions.

1.2: Risk Engine Review

The Risk Engine is an important server cluster in Bank of America that computes risk, Actual P&L, and the general ledger for the entire corporation. This server cluster was reviewed to see if there are more power efficient options available that would be worthwhile to convert to. This is an area for saving money and server space for Bank of America.

Currently, the server cluster has 80 2.8GHz Dual Core Xeon Processors stored in 40 Hewlett-Packard DL360 G4p server chassis. This sets up a problem of whether or not a quad-core processor server or a Cell-based option would be worthwhile. This project focuses on a comparison of computational capabilities, electrical power consumption, and space usage. With these three points, a conclusion was drawn on the whether a full revamping was required. Once the conclusion was drawn on the benefits, this was weighed against the cost of the new hardware and if that hardware would pay for itself through power consumption.

In this project, it will be shown that the current system is sufficient as the hardware has already been purchased. Near future upgrades should be focused on the quad-core processor option while after mid-2008 the Cell-based option should be investigated further. This will be shown through the fact that even though the Cell will cut down power consumption and space required that the uncertainty in the difficulty of coding for its architecture is detrimental to its implementation and viability as an option.

1.3: Social Capitalism using Econophysics

In this project, it will show that econophysics is a growing research field that is being used to predict future world economics. Explanations of current and future research ideas have been expressed in this project. Studying Flash and Actual P&L systems, time was been found to research the basics of econophysics and the thermal physics skeleton used for basic concepts of Social Capitalism. This will show how various social classes are created and categorized using thermal physics analogies.

Thermal physics ideas represented in the project include the Boltzmann-Gibbs Law, thermal, and super-thermal physics. These characteristics and the mathematics involved are then applied to world economics using data that has been collected over the years. With this data analysis, accurate predictions can be made about past, current and future world economic problems and solutions. We have the ability, as a world, to unite and study our world economy, for future stabilization ideas. People do not realize that we all affect each other in some way, and choices on this Earth will decide if we continue to prosper here on this planet for years to come.

2: Background

2.1: Financial

2.1.1: Bonds

Bonds are a form of credit where a company or a government issues you debt in order to raise funds. The body that issues the bond makes a commitment to return the initial investment of the bond by a certain date. They also agree to pay interest at either fixed rates or floating rates on set dates. Thus, a bond is a loan where the traditional roles are reversed. Instead of having a bank or a company lending the money, you are providing the funds by buying into a company's debt. However, there are a few more aspects to a bond than there are to a loan. A bond includes the details of who was the issuer, the time frame of the bond, the coupon rate, and if there are redemption features.[1]

Issuer	**Examples** Corporation Municipality Government	**Domestic Examples** Japanese Government Bonds, U.S. Treasuries, UK Gilts
Maturity	Short, Medium, Long, Sinking Fund Provision	**International Examples** Eurobond Foreign Global
Coupon Rate	Fixed Income Floater Inverse Zero Coupon	
Redemption Features	Callable Convertible Puttable	

Table 1 - Features of Bonds[2]

[1] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. Credit Derivatives. Upper Sandle River, New Jersey 07458: Wharton School, 2006. p. 12

[2] Chako, George, Peter Hecht, Vincent Dessain, and Monika Stachowiak. "Note on Bond Valuation and Returns." Harvard Business School Note No. 205-008. Boston: Harvard Business School Publishing, 2004

This table demonstrates various features of the bonds with specific examples. As shown above, the issuer of a bond can vary from a corporation to a government. A bond will also have a maturity of when the bond can be cashed in. Coupons can also be collected from a bond if that is written in the original agreement. While the redemption features, allow a bond to be redeemed in different fashions. Each of these factors will affect the risk associated with the bond in varying degrees.

The issuer of the bond is one of the most important parts of bond's risk and its expected yield. When comparing a government to a corporation, it is easy to see that a corporation is a lot more likely to go bankrupt than a government. For this reasoning, government bonds are seen as safer and provide lower returns. In fact, when bonds are issued by the U.S. government these are seen to have the lowest chance of default. Often, U.S. Treasury bonds are commonly referred to as risk-free bonds. As shown in the table above, each of the bonds can have different names based on their country of origin.[3]

Eurobonds are a particularly interesting case of bonds. They are issued bonds in a currency that is not currency of the country where it is being originally issued. If a company sat in Australia and issued a bond in US dollars while not in the US that would be considered a Eurobond. Once a bond fully matures it is time for the issuer to make good on the commitment to repay its debt. They do this by paying back the principal. When dealing with government bonds, a short term bond is less than one year, a medium term is one to two years, and a long term has a term-of-maturity that is ten years or more. In the case of the United States, short-term bonds are Treasury Bills, middle-term bonds are Treasury Notes, and long-term bonds are called Treasury Bonds. On the other hand, a bond issued for the short term through a corporate source

[3] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp. 12-13

is often called commercial papers and have a maturity between 2 and 220 days. Long-term bonds issued by a corporation are called corporate bonds.[4]

Corporate bonds may also have what is called a sinking fund provision. This provision is a condition that mandates that the entity has to pay back the principal of a bond as time progresses. This would make it so that instead of getting one lump sum at the end that the bond holder would instead receive payments at fixed intervals. A zero-coupon bond is a bond that doesn't pay interest periodically and only pays back the principal when maturity is reached. Bonds can also have redemption features. When a bond is callable, that means that the issuer can recall the bond before it has fully matured. This will cause the issuer to pay more than they normally would. Puttable means that the bond holder can turn in the bond before it's fully matured. This will normally be redeemable for less. Convertible bonds give bond holders the option to convert their bonds into a set number of equity shares before or after the bond's maturity.[5]

An important distinction between bonds and stocks is that stocks do not promise dividends or returns. Bonds are a promise that you will be repaid the principal and interest while stocks provide ownership in a company.[6] In the United States, when you purchase a bond from the U.S. Treasury you only have to pay federal taxes on the income earned. When you purchase a municipal bond, you are exempt from paying federal taxes on the income you make from that. Also, a lot of municipalities will make tax exemptions for their bonds so that you don't have to pay taxes on the income from the municipal bond. These do come with a drawback because if

[4] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p. 13
[5] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p. 14
[6] "What is a Bond? - Education Center." Yahoo! 23 Oct. 2007 < http://finance.yahoo.com/education/bond/article/101194/What_is_a_Bond>.

you are offered complete tax-exemption these bonds will usually have a smaller return rate than a taxable bond.[7]

Risk is another aspect of bonds as it is with all forms of credit. When you buy a U.S. Treasury bond it is considered exceptionally safe. On the other hand, when you buy a municipal bond that is believed to be a bit riskier. Corporate bonds are seen as the riskiest because they're more vulnerable to various problems than a government. Two examples of these incidents are Chrysler and LTV Steel. Each of these bonds will come with a rating.[8]

Standard and Poor's Rating	Moody's Rating	Explanation	Grade
AAA	Aaa	Highest	Investment
AA	Aa	Very High	Investment
A	A	High	Investment
BBB	Baa	Middle	Investment
BB	Ba	Low	Non-investment
B	B	Very low	Non-investment
CCC	Caa	Great risk	Non-investment
CC	Ca	Greater risk	Non-investment
C	C	Default unavoidable	Non-investment

Table 2 - A list of Standard and Poor's Ratings and Moody's Ratings[9] [10]

Bonds are a specific kind of credit and with that a credit rating is given for the amount of risk that is likely to be incurred. There are agencies that rate how likely it is for a company to repay. Two of the major agencies are Standard and Poor's (S&P) and Moody's. These ratings are representative of the risk.[11] AAA (Aaa) bonds are highly likely to be paid back while CCC/CC/C or a Caa/Ca/C bonds aren't as likely to be paid back. On a side note, when investors look to

[7] "Types of Bonds - Education Center." Yahoo! 23 Oct. 2007 <http://finance.yahoo.com/education/bond/article/101197/Types_of_Bonds>.
[8] "Types of Bonds – Education Center" Yahoo!
[9] "Bond Ratings." Money-Zine. 24 Oct. 2007 <http://www.money-zine.com/Investing/Investing/Bond-Ratings/>.
[10] "Finance Corporate Debt Ratings Moody's and S&P." Blaha. 24 Oct. 2007 <http://www.blaha.net/Finance%20Corporate%20Debt%20Ratings.htm>.
[11] "Types of Bonds – Education Center" Yahoo!

invest across the world, they first have to look at the sovereign credit rating. This credit rating states how well a country can offer a safe investment environment.[12]

It is important to understand that a bond usually entails a legally binding document that is often at least one-hundred pages in length. The document that goes along with the bond and covers various aspects such as the parties that are involved, legal jurisdiction, and also include information on the cash flow and the amounts involved. An aspect of a bond that is also important to know is the grace period. This is the amount of time that is allowed to make payment before the bond defaults. This is usually the case when the borrower cannot support a large payment on the cash flow. This usually causes a default on all bonds and loans that body had.[13]

2.1.2: Stocks/Equity

While one can invest in the debt of a company through bonds, there is yet another way of investing in a company. A person can buy equity in a company. This equity is seen in the form of stocks. So while on one side of a company, there are a company's creditors, on the other side are the stockholders who have purchased equity in the company. These stockholders are all partial owners of the company. The percentage that a stockholder owns a company is determined by the percentage of stock that they own.[14]

When talking about actually purchasing the stock, it is referred to as buying shares. There are different types of stock that is available for purchase. Common stock is where you own a part of the company and have special rights to be able to vote for the company's board of directors.

---

[12] Heakal, Reem. "What is a Corporate Credit Rating?" Investopedia. 22 Oct. 2003. 23 Oct. 2007 <http://www.investopedia.com/articles/03/102203.asp>.
[13] Chaplin, Geoff. Credit Derivatives: Risk Management, Trading, & Investing. West Sussex, England: Wiley Finance, 2005. p.6.
[14] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p.227

The amount of votes you have is directly proportional to how much stock you have. On top of being able to vote for the board of directors, common stock shareholders also get to vote on various important events such as a stock split. Common stock shareholders are also the very end of the order for payout if the company defaults.[15]

Preferred stock is a type of stock where the shareholders still own a part of the company. However, they forfeit their right to vote but they do receive special treatment. There are instances where preferred stockholders do get voting rights. Preferred stock holders get paid a fixed dividend and will enjoy smaller price fluctuations because of the dividend paid. This dividend doesn't have to be paid necessarily if the company doesn't have the financial capacity to do so. Preferred stockholders, in the event of a default, will receive compensation before a common stock holder. There are four categories for preferred stock holders. Cumulative preferred stock enables their holders the right to be paid an accumulative back payment if the company doesn't have the financial capacity to do so for awhile. Non-cumulative does not receive this benefit. Participating preferred stocks may give out higher than normal dividend payments if the company produces a greater than expected profit. Convertible preferred stocks allow a preferred stockholder to convert their preferred stock into a certain amount of shares of common stock.[16]

It is important to also realize that not all common stocks are created equal. Some companies will issue different classes of common stock. This operates much like the different levels of bonds (senior and junior). These different levels of common stocks enable their stockholders to have different amounts of votes per share that they own. A company could have

[15] Investorguide Staff. "Stock Basics - Common and Preferred Stock." Investor Guide. 1 Nov. 2007 <http://www.investorguide.com/igu-article-818-stock-basics-common-and-preferred-stock.html>.
[16] Investorguide Staff. "Stock Basics - Common and Preferred Stock." Investor Guide.

a common stock that they rate higher and it gives five votes per share while they could have another common stock that only gives one vote per share.[17]

2.1.3: Credit

Credit is when a person or a company can use someone else's funds in order to sustain themselves financially. This is what makes credit an important economic concept. It allows people to buy what they want today even though they may not have the funds to repay it until sometime in the future. Credit carries with it the idea of interest. Interest can be paid back in regular intervals or at the end of the contract.[18]

When someone lends out money, they expect to be compensated for taking on the risk of the borrower not returning the money. This is where interest starts to come in. Since the creditor is lending out their money, they expect that they will be paid in return for parting with their money. This payment is called interest. An important aspect to interest; however, is that a time value is also associated with this. Time values are a way of compensating for inflation. Inflation lowers the value of money by increasing the amount of money required to purchase goods.[19]

Currency is a form of credit that is very commonly used. We may have United States dollars in our pocket but those pieces of paper that we have aren't really worth what their face value declares. That is to say, the physical piece of paper, the material that goes into it is worth less than its face value. These bills are backed by the United States government's credit. Every

[17] Investorguide Staff. "Stock Basics - Common and Preferred Stock." Investor Guide.
[18] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp. 9-10
[19] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p. 11

form of currency is backed by some sort of government's credit whether it's the US dollar, the British pound, or even the Euro. [20]

Loans are a type of credit where a lender gives the borrower something. This could range from money to property. The borrower then promises to return the capital or goods to the lender. On top of returning the money or property, they also have to pay either a fee or interest. A very common case of this is a mortgage. This is when a person takes out a loan and puts their house up as collateral. If the person then defaults on their loan, their house can be taken by the lender instead of paying the rest of the loan.[21] Loans also carry with them loan documentation, which are legally binding documents.[22]

Loans are a broader topic than bonds. They're very loosely defined. Often, bonds are things that are traded while loans are usually private agreements. On the other hand, there is actually a maturing market for secondary trading of loans. This is when the cash flows are sold to a third party. If this is to take place the loan has to be an assignable loan. Most loans are non-assignable. This trend seems to becoming less and less popular. Some non-assignable loans will become assignable during a time of default. This makes things a little more interesting when trying to deal with the default swap contracts.[23]

The inability of being able to pay for your credit is called defaulting. This can happen to people, corporations, and governments alike. People can over use their credit card. This form of debt is popular among college students and often leads to defaulting. There has been a steady increase from 1992 to 2003 for people filling bankruptcy. Corporations can easily go bankrupt

[20] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p. 11
[21] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp. 11-12.
[22] Chaplin, Geoff. p.3
[23] Chaplin, Geoff. p.3

and an early tell-tale sign is when the company stops paying interest on issued bonds. Governments are considered to be impervious to default for the most part as they can just print more money to pay off its debt. However, Russia defaulted on its loans in 1998, Turkey in 2001, and Argentina in 2002. That is when the International Monetary Fund will step in and try to formulate some sort of repayment plan.[24]

There are occasions that can happen to force an entity to default. These are called credit events. Examples of this are bankruptcy, failure to pay, a severe reduction of credit ratings, an event that happens after a merger leaving the new company weaker than before, government action, and market disruptions. Defaulting on a loan is likely to lead to bankruptcy. In the case of Argentina, however, they did not have to declare bankruptcy. On the other hand, corporations when they default will more than likely have to file for bankruptcy. When this happens, the company defaults on all of its loans and the company starts to be liquidated. This liquidation process is an attempt to salvage as much as possible from the company. All of the debt that that company once had is then ranked by seniority. The ones who hold the most seniority will be the first to be paid back.[25]

In the event of a company defaulting like what is stated above, all of the debt will then be broken up into senior and junior debt. Senior debt will be paid back before junior debt. Senior debt also comes with a lower potential for profit as it is lower risk than the junior debt holders. Debt holders have a priority for repayment over stock and equity holders. The company owners are given whatever is left over from the company. This starts to show why investing in bonds is

[24] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp. 14-17
[25] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp. 18-19

generally safer than investing in stocks. As a bond holder will get paid before a stock holder since the bond holder is investing in the company's debt not its equity.[26]

2.1.4: Credit Risk / Credit Default

Credit risk is the danger of not being paid back for the money that you lend. This danger is often called the default risk. In order to handle credit risk, it is usually broken down into pieces that are easier to deal with. There is a lot of debate on this topic on how to actually do this; however two generally accepted concepts are default probability and recovery rate. These two concepts when combined into a single form of measurement is frequently called credit spread.[27]

Default probability can be calculated in numerous ways, which includes but is not limited to the entire life of the loan or over specific time periods. Looking at a borrower's credit history is a very traditional way of establishing the probability of default. Credit ratings are assigned to everyone who uses credit in order to determine how likely they are to repay the loan and sometimes how much to charge extra. Someone with a low credit score would be charged much more and be kept on a shorter leash than something with a higher credit score. We can also run stress scenarios on the debt to see how a company might respond to changes such as a drop in interest rates or any other changes in its environment.[28]

The recovery rate is how much of loan can be recovered if the borrower defaults. This was covered in the credit section on the grounds of who gets paid first if a corporation defaults. During the liquidation process, everything will be sold and the lenders will be looking for as much compensation for their loan as quickly as possible. Also, there is another aspect that has to be looked into when evaluating a entity's credit. This facet is called credit exposure and it is also

[26] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p. 19
[27] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p. 20-21
[28] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p.22

known as exposure at default (EAD). This is a measure of how much the person or organization who lends the money would lose if the borrower were to default. A borrower could also be given a set threshold that he or she must not cross. The EAD measurements are usually for derivative instruments, which allow people to sell credit risk. By taking into consideration default probability, recovery rate, and credit exposure a formula can be founded for the expected loss.[29]

$$L_e = D * E_c * (1 - R_{y)}$$

Formula 1 - Calculation for Expected Loss[30]

L_e is the expected loss while D is the probability of default. E_c is credit exposure and R_r is the recovery rate. It would appear pretty difficult to put numbers to things that seem as arbitrary as these components. However, default rates are very easy to give values to. Recovery rates on the other hand are not so easy. Rating companies are becoming better in setting default rates. Very few high-rated companies are defaulting while the majority of defaulting companies were in low ratings as their ratings tanked. Companies that have a high rating are called investment grade while ones with low ratings are called speculative grade. This is used by investors who are wondering about the likelihood of the issuing entity defaulting.[31]

[29] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp.20-22
[30] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p.22
[31] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp.26-40

Original Rating	Probability for changing a rating at the end of the year (%)							
	AAA	AA	A	BBB	BB	B	CCC	Default
AAA	93.66	5.83	0.04	0.08	0.03	0.00	0.00	0.00
AA	0.66	91.72	6.94	0.49	0.06	0.09	0.02	0.01
A	0.07	2.25	91.76	5.19	0.49	0.20	0.01	0.04
BBB	0.03	0.25	4.83	89.26	4.44	0.81	0.16	0.22
BB	0.03	0.07	0.44	6.67	83.31	7.47	1.05	0.98
B	0.00	0.10	0.33	0.46	5.77	84.19	3.87	5.30
CCC	*0.16*	0.00	0.31	0.93	2.00	10.74	63.96	21.94
Default	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00

Table 3 - Probability of a bond-rating changing over one year (Standard and Poor's)[32]

It is clearly visible why bonds that are taken out over a longer period of time have larger yields. The chance that your bond shifts in rating increases as time goes on. This information was collected over a period from 1981-2000 and is originally from Standard & Poor's. Also, the probability that CCC becomes an AAA has been italicized in order to separate it as an irregularity. This is clearly evident as a B rated bond has a zero percent chance to become an AAA bond. That would imply that a CCC bond shouldn't have a chance of becoming a AAA bond. This is definitely something that an investor looking at bonds should understand before they take a step towards investing in a company.[33]

Original Rating	Probability for changing a rating at the end of the year (%)							
	Aaa	Aa	A	Baa	Ba	B	Caa	D
Aaa	93.38	5.94	0.64	0.00	0.02	0.00	0.00	0.02
Aa	1.61	90.53	7.46	0.26	0.09	0.01	0.00	0.04
A	0.07	2.28	92.35	4.63	0.45	0.12	0.01	0.09
Baa	0.05	0.26	5.51	88.48	4.76	0.71	0.08	0.15
Ba	0.02	0.05	0.42	5.16	86.91	5.91	0.24	1.29
B	0.00	0.04	0.13	0.54	6.35	84.22	1.91	6.81
Caa	0.00	0.00	0.00	0.62	2.05	4.08	69.19	24.06

Table 4 - Probability of a bond-rating changing over one year (Moody's)[34]

[32] "Default Model." Riskglossary. 23 Oct. 2007 <http://www.riskglossary.com/link/default_model.htm>.
[33] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p.31
[34] "The JP Morgan Guide to Credit Derivatives." Investing in Bonds. 24 Oct. 2007 <http://www.investinginbonds.com/assets/files/Intro_to_Credit_Derivatives.pdf>. p. 44

Moody's ratings can be comparable to Standard and Poor's to show that their predictions run closely together. We can see that both of the rating companies are very similar in their end results even if they do their ratings differently. This is very important information when considering what to invest in. Traders need to know and understand this in order to truly start evaluating risk. If the rating drops then it will bring down the total yield. If the rating drops significantly the profit may turn into a loss.

There are also other rating companies such as Fitch IBCA and Duff & Phelps. Fitch IBCA and Duff & Phelps both use a letter system that is like Standard and Poor's. They just give another opinion of how they think a company should be rated. These however aren't the most popular rating companies but they are still used. It also noteworthy that Moody's includes 1, 2, and 3 for their ratings from Aa to Ca to separate better stocks at each level. Standard and Poor's, Fitch IBCA, and Duff & Phelps use plus and minus signs on AA to CC to show which stocks are higher rated than others.[35]

2.1.5: Credit Derivatives

Credit derivatives are an over-the-counter financial contract that gets its basis in the idea of a credit event.[36] Credit events are when the protection buyer has the right to cash in on their credit derivative. It's like buying insurance for an investment and if that investment goes awry, the protection buyer receives compensation. [37] Credit events are a very important part of understanding credit derivatives as this is the foundation for how to collect the actual credit derivative. The International Swaps and Derivatives Association (ISDA) is the organization that

[35] Schurr, Stephen. "Bond Ratings." TheStreet. 16 Aug. 2000. 25 Oct. 2007 <http://www.thestreet.com/basics/gettingstarted/1036369.html>.
[36] Choudhry, Moorad. The Credit Default Swap Basis. New York: Bloomberg, 2006. p. 2.
[37] "Credit Event." Creditflux. 25 Oct. 2007 <http://www.creditflux.com/glossary/credit+event.htm>.

sets up definitions of what constitute a credit event.[38] These definitions are always under review and being updated.

A company can set their own requirements for what a credit event is exactly. So this subject matter stands on somewhat shaky ground. The ISDA's definitions are a guideline for which protection sellers can base their own requirements off of. Credit derivatives are a tool that is used to protect someone from credit exposure. Each can be engineered to fit a particular buyer's requests. Some of the benefits to using credit derivatives are customizability, profiting from a declining stock by selling short, more efficient pricing and it allows people to have access to elements in the market that an investor normally wouldn't deal with.[39]

Credit derivative instruments allow people to segregate and trade credit risk. This would be through the purchasing of protection or selling of protection. Credit derivatives have these features: what the protection is being issued for, applicable credit events and what constitutes them, and how things will be settled in case of a credit event. Credit derivatives fall into two categories. They are either funded or unfunded. Funded credit derivatives are usually seen as credit-linked notes. This is a note where the investor is the entity who sells the credit-protection. The issuer is the body that buys the actual note. If a credit event does not occur, then the investor is paid by the buyer of the note. The amount that is paid to the investor is called the redemption value of the note. Of course, if a credit event does occur then the investor will not receive the redemption value they will receive a lesser amount. This lesser amount depends on

[38] Kothari, Vinod. "ISDA Credit Event Definitions." Credit Derivatives. 25 Oct. 2007 <http://www.credit-deriv.com/isdadefinitions.htm>.
[39] Choudhry, Moorad. pp. 2-4

how the credit event affected the insurance. Also, how this settlement is done will based on the agreement of if it is to be a cash or physical settlement.[40]

When a credit-linked note is issued, the entity that purchases the protection is the seller of the note. The buyer will then transfer their credit risk through the seller of the protection. A note may be very similar to a credit default swap's set up. It can have an expiration date, pay some sort of rate (LIBOR based or not), or it could even just be a fixed payment plan. This is the part where the funded credit derivative ideology comes in. There is a periodic cash flow coming from the buyer of the protection. This form of protection will terminate early if there is a credit event on the seller of the protection's side. If this happens, then the buyer of the protection no longer has to pay coupons to the seller. The seller will have to pay the principal, which all of the periodic payments were based on.[41]

Unfunded credit derivatives are usually associated with credit default swaps. In this case, the protection buyer does not get an immediate payment from the protection seller. What happens is when a credit event happens, that is when the protection seller makes the payment. If the credit event never happens, then the protection seller never has to make payment. Both of these techniques are very important as they offer great flexibility when dealing with various trades. Unfunded credit derivatives embody the idea of a credit derivative as they segregate and relocate credit risk. This means that the value retained in these are a representation of the credit quality from the original source. On the other hand, if this was a bond the value would be determined by the risk and credit quality.[42]

[40] Choudhry, Moorad. pp. 2-6
[41] Chaplin, Geoff. p.143
[42] Choudhry, Moorad. p. 7

Also an interesting comparison to bonds, loans, and credit derivatives is that bonds and loans depend on credit quality, funding, the interest rate, and the currency that it's issued in. However, if you combine the bond with an interest-rate swap then it becomes an asset swap. Asset swaps eliminate the interest-rate risk, which in turn makes the asset swap only depend on credit and funding. To eliminate the funding aspect, we take the idea of an unfunded credit derivative, which was previously defined as not requiring funding at the start of the transaction. This allows us to produce a product where the protection seller's return only depends on the credit quality of the one who purchases the protection.[43]

Credit default swaps (CDS) are the most widespread credit derivative instrument that is traded. As previously established, CDS are characterized as being unfounded credit derivatives. CDS is also known as credit swap and default swap. With unfunded credit derivatives it is now understood that the protection buyer pays the protection seller a premium. This premium is set around what is called a basis point multiplier. A basis point (bp) is one-hundredth of one percent. This is a common financial phrase and a lot of other aspects in finance depend on having an understanding of basis point references. Credit default swaps are used by banks for countries and corporate credit. The protection buyer will also gain if the reference asset gets a ratings decrease or defaults.[44]

Contingent credit swaps require a credit event and another trigger in order to be activated. This extra trigger is usually another credit event that takes its indication from a change in equity prices, commodity prices, or even interest rates. This credit event could also be with respect to a reference entity. This makes the contingent credit swap cheaper as it needs two events in order to

[43] Choudhry, Moorad. pp. 7-8
[44] Choudhry, Moorad. pp. 7-8

be triggered. This also makes the contingent credit swap weaker than a regular credit swap. This is considered to be a good option when there is little likelihood that both triggers will occur.[45]

Dynamic credit swaps are used to deal with the hardships of managing risk in a portfolio. During a dynamic credit swap, the amount of credit exposure will shift over time and as the foundation market changes. Both of the entities involved are linked to the counterparty's credit exposure. An investor can then buy protection for this swap as a backup measure. This means that the investor will not have to worry about default unless the counterparty fails and the protection seller cannot come through with their obligation. With this special case, this means that the original investor is protected more effectively as he has two companies to fall back on for his protection. This is because it is very unlikely that both companies will default.[46]

Asset swaps are not truly credit derivatives; however they are still lumped into the same market. Their actual classification is as an interest-rate swap. Asset swaps main uses are as a way of altering the clash flow profile of a bond. They have a comparable nature to credit derivatives. Asset swaps are important to this market because they are set the price according to the London interbank offered rate (LIBOR).[47] LIBOR rates are used because general collateral (GC) rates are not readily accessible. The GC rates are known by the repurchase agreement (repo) trading desk and aren't made readily accessible to the public. The LIBOR rates, on the other hand, are simple to get. These rates are even very close to the GC rates. Usually the LIBOR rates are only a few basis points above the GC rates.[48]

[45] "The JP Morgan Guide to Credit Derivatives." Investing in Bonds.
[46] "The JP Morgan Guide to Credit Derivatives." Investing in Bonds.
[47] Choudhry, Moorad. p. 13
[48] Chaplin, Geoff. p.4

A better way of explaining asset swaps is we have a buyer of protection who has a fixed-rate bond that's on the riskier side. He wants insurance, so what he does is offer up the coupons of the bond to the seller of protection. The seller then agrees to pay the protection buyer LIBOR plus or minus the spread. This gives the investor a chance at possibly financing the CDS from the asset swap's income. The buyer will usually take the unwind cost of the interest-rate swap if the asset swap is finished.[49] Unwind is a term for a closure of an investment position. So it's as if they are taking the final amount.[50]

Credit derivatives also include credit options and total return swaps.[51] So on top of being funded and unfunded, a credit derivative can fall in the credit default swap, credit options, or total return swaps. Both of which are covered in later sections.

2.1.6: Total Return Swap

A total return is a collection of different forms of income, which include interest, capital gains, and dividends.[52] A total return swap is, therefore, a deal where a person will pay the recipient the total return on a reference asset. This deal depends on setting up a time span for this to occur. The payer takes on the role of being a buyer of protection while the recipient takes on the roll of the seller. What the payer gets in return for his payment is given a "floating leg".[53] A "floating leg" is basically a part of a deal that is not strongly linked to changes in interest rates.[54] This "floating leg" usually will involve being given a LIBOR and some sort of spread.

---

[49] Choudhry, Moorad. pp. 13-14
[50] "Unwind Financial Defintion of Unwind." Free Online Dictionary. 25 Oct. 2007 <http://financial-dictionary.thefreedictionary.com/unwind>.
[51] "The JP Morgan Guide to Credit Derivatives." Investing in Bonds.
[52] "Total Return." Investopedia. 29 Oct. 2007 <http://www.investopedia.com/terms/t/totalreturn.asp>.
[53] Chaplin, Geoff. p.163
[54] "Interest Rate Swaps Pricing." Derivative One. 29 Oct. 2007 <http://www.derivativesone.com/kb/interest_rate_swaps.aspx>.

The money coming from the reference asset will usually be paid to the receiver. This may sound like one is paying something for nothing, however, there is something else going on here.[55]

If the asset depreciates then the payer will have to pay a negative amount. This negative amount is where the protection comes in. The negative amount that the payer has to pay will ultimately mean that the receiver will actually owe the payer money. This can be very useful if the asset defaults. This will mean the receiver will have to pay a substantial amount in order to reconcile the debt. Risk starts at zero for the seller of protection; however it will become riskier if the reference asset change course from its increasing price curve. The dates that the cash flow is paid on will effectively help the seller of protection to decrease their risk.[56]

This type of transaction is called an off-balance transaction. This means that the transaction takes place off of the balance sheets. This is done because an organization will want to keep their debt to equity and leverage percentages low.[57] Banks, in particular, would use this type of transaction. This allows a bank that has gone past its balance sheet boundaries to go to another bank that has some room left on their balance sheet. This allows a bank to transfer the asset off of its balance sheet and move it to somewhere else. This is in spite of still having ownership of the asset. These deals are usually of a very large magnitude in the ballpark of billions of dollars and these deals would usually be very short-term.[58]

Total return swaps are starting to be used by hedge funds now. This total return swap allows a hedge fund to have cash flow coming in from something they don't have to specifically pay for themselves. These hedge funds are seen as undesirable by banks that are looking to buy

[55] Chaplin, Geoff. p.163
[56] Chaplin, Geoff. p.163
[57] "Off-Balance-Sheet Financing." Investopedia. 29 Oct. 2007 <http://www.investopedia.com/terms/o/obsf.asp>.
[58] Chaplin, Geoff. pp.163-164

protection. That is unless the hedge fund has a significant amount of collateral that it can offer up in case something happens. An investor could also use this in order to get ownership of something that he would normally not have access to. For example, an investor could ask a bank inside of that country buy the item and then it could give economic ownership to the investor through the use of this method.[59]

2.1.7: Collateralized Debt Obligations

Collateralized debt obligation (CDO) is an umbrella term that encompasses with it collateralized loan obligation (CLO) and collateralize bond obligation (CBO). With collateral debt obligations comes an idea of tranches.[60] A tranche is a level inside a deal and there can be multiple trenches offered for a CDO. Tranches are often used with bonds to offer different amounts of risk.[61] These collateralized debt obligations will usually pool together assets, which usually include loans and debt instruments. The investors will be the ones who have to sustain the credit risk of the collateral that is a part of the collateralized debt obligation. With this, an investor can select from a multitude of different tranches. These tranches will offer different maturities and credit risk. In this case, tranches are categorized as being senior, mezzanine, or equity depending on their amount of risk.[62]

If the collateralized debt obligation under performs or there is some sort of default then the senior tranches will have to be paid first. This is much like seniority of a bond. The people who hole senior tranches will be catered to first and then mezzanine. Once mezzanine is met then the equity level will be taken care of last. There are ratings that are associated with the senior and mezzanine levels to indicate priority also. For senior levels, the ratings go from A to AAA while

[59] Chaplin, Geoff. pp.163-164
[60] "Collaterialized Debt Obligation." Risk Glossary. 29 Oct. 2007 <http://www.riskglossary.com/link/collateralized_debt_obligation.htm>.
[61] "Tranches." Investopedia. 29 Oct. 2007 <http://www.investopedia.com/terms/t/tranches.asp>.
[62] "Collaterialized Debt Obligation." Risk Glossary.

for mezzanine the ratings go from B to BBB. This gives the investor an idea of the quality of the credit while also understanding where they come in the risk chain.[63]

The life of a collateralized debt obligation has different phases. These phrases are known as ramp-up, reinvestment, and sell-off. During ramp-up, the person who is in charge of the portfolio will invest the earnings from the sales involve with the CDO. The reinvestment period is where the manager will reinvest the cash flows of the CDO while buying and selling assets. The sell-off period includes when the collateral has matured and finally is sold. The investors will be paid at this point. Investors at the time of entering a CDO do not know what assets the CDO will be investing in. All they know is who is managing the profile and what sort of guidelines he will be using in order to choose worthy investments. Thus, this means that not only does an investor have to worry about the credit risks but they also have to worry about the person who is actually managing the portfolio. On top of all of that, they also have to pay for the portfolio manager.[64] This is where the idea of being a funded credit derivative comes into play. They have to pay a fee during the duration of the deal.

2.1.8: Bond Spreads, I-Spreads, Z-Spreads

A bond spreads' relative value is done by evaluating its yield spread in a comparison to another bond. When the comparison is made and the spread is determined, then that will show the yield of the bond. This is where the idea of interest rates per basis point comes in. This gives us the acronym IR01, which reads as interest rate per basis point. We are able to measure the spread of a swap by comparing the yield of the bond to the interest rate. This interest rate is evaluated by using a straight line and interpolating the results of the swap curve. For the most

[63] "Collaterialized Debt Obligation." Risk Glossary.
[64] "Collaterialized Debt Obligation." Risk Glossary.

part, investors will use asset-swap spread and Z-spread, in order to obtain a relative value. However, it's possible to use a government bond spread as these are considered risk-free bonds. When the risk is great, then the spread should also be greater. That is why investors are looking to have a narrower spread as that means that their investments are becoming less risky. However, the less risk there is in the investment, the less the return will be.[65]

I-spread is occasionally a tool for evaluating a cash bond with an equivalent CDS price. I-spread is effectively a spread over a swap. Spread over swaps can be easily calculated by taking the yield of the bond and subtracting the swap rate of a corresponding swap. Often investors will choose to use the asset-swap spread instead of the I-spread. It is however, very easy to relate the I-spread to the yield of a corporate bond. For this section, Y is the yield on the corporate bond, I is the I-spread, S is the swap spread, and T is the yield on a Treasury security.[66]

$$Y = I + S + T$$

Formula 2 - Equation for calculating a yield on a corporate bond using I-spread[67]

$$B_s = I + S$$

Formula 3 - Equation for calculating bond spread[68]

Z-spread is often used for relative analysis. This is because asset swaps use something called a bond's yield-to-maturity when it computes its spread. Using the yield-to-maturity approach makes it troublesome to actually use when calculating out a relative analysis. The Z-spread does things differently by actually using a zero-coupon yield curve. This allows for a much more realistic spread to be used in the comparison. Also, the zero-coupon curve is derived from the interest-rate swap curve when it is used for this method. This makes them very closely related. In fact, for short-term bonds and higher rated bonds the difference between the Z-spread

[65] Choudhry, Moorad. pp. 37-38
[66] Choudhry, Moorad. pp. 39-40
[67] Choudhry, Moorad. p. 39
[68] Choudhry, Moorad. p. 40

and asset swap spread is very small. If these differ greatly then that would lead to a conclusion that the bond has been priced incorrectly.[69]

Term	Definition
n	Number of interest periods until maturity
P	Bond price
C	Coupon
M	Redemption payment
Z	Z-spread
m	Frequency of coupon payments

Table 5 - Terms for equation for bond prices[70]

$$P = \sum_{i=1}^{n} \left[\frac{C_i + M_i}{\left(1 + (\frac{(Z + S_i + T_i)}{m}\right)^i} \right]$$

Formula 4 - Equation for bond price by understanding Z-spread[71]

2.1.9: Stock Options

Stock options are a way of giving you the right to buy or sell a stock at a predetermined price. This price is called the exercise price. However, this option is for a certain amount of time and eventually the option expires. A call option is when you have the right to buy shares of stock. On the other hand, if the option lets you sell the shares of stock then it is called a put option. It is important to know that when you buy this stock option, that you are not locked into purchasing or selling the stock. The stock option only gives you the right and doesn't make you obligated in any form to purchase or sell the stock at the strike price.[72]

[69] Choudhry, Moorad. pp. 41-45
[70] Choudhry, Moorad. p. 44
[71] Choudhry, Moorad. p. 44
[72] "NOVA Online | Trillion Dollar Bet | Site Map." PBS. 26 Oct. 2007 <http://www.pbs.org/wgbh/nova/stockmarket/formulaoptions.html#options>.

Effect: Increase	Call Value	Put Value	Effect: Decrease	Call value	Put Value
Underlying asset's value	**Increase**	Decrease	Underlying asset's value	Decrease	**Increase**
Strike price	Decrease	**Increase**	Strike price	**Increase**	Decrease
Variance of underlying asset	**Increase**	**Increase**	Variance of underlying asset	Decrease	Decrease
Time to expiry	**Increase**	**Increase**	Time to expiry	Decrease	Decrease
Interest Rates	**Increase**	Decrease	Interest Rates	Decrease	**Increase**
Dividends Paid	Decrease	**Increase**	Dividends Paid	**Increase**	Decrease

Table 6 - The effect on call and put value[73]

However, on a much larger scale, one can have a credit option. A credit option will allow one to apply call and put options on a form a credit. This credit could be a note, bond, or loan. This also includes an asset swap. So when dealing with this type of credit option it follows the same thought process. A person is given the right to purchase or sell the reference asset at a strike price. They are not required to purchase the item once the call or put option comes through. This gives one the ability to price credit options similarly as if they were stock options.[74]

Often, the buyers of credit options will be organizations such as banks. These organizations are looking to reduce their credit exposure. Since credit options are a part of credit derivatives, they share the same benefit of enjoying the off-balance-sheet aspect. This allows a company to keep its debt to equity ratio low while investing in credit derivatives. Credit swaps can also behave like credit options. That is to say, that credit swaps can actually be customized so that when a negative credit event occurs that payments will be triggered. In addition to these advantages, a portfolio that is seeing hard times and has to sell assets can create credit derivatives that will protect against the decline. This allows a portfolio to hang onto assets that have a mediocre credit quality but enjoy a smaller risk.[75]

[73] Damodaran, Aswath. "Option Pricing Theory and Models." Damodaran Online (Investment Valuation: Second Edition). New York University Stern. 26 Oct. 2007 <http://pages.stern.nyu.edu/~adamodar/pdfiles/valn2ed/ch5.pdf>.
[74] "The JP Morgan Guide to Credit Derivatives." Investing in Bonds.
[75] "The JP Morgan Guide to Credit Derivatives." Investing in Bonds.

When pricing an option, the option provider has to look at a few different things. The current value of the asset that the option is being taken out for is one of the first pieces of the option price. How much the value of the underlying asset varies will also determine the price of an option. The more that the value of the underlying asset varies the more the option will cost. This is because a person who buys an option cannot possibly lose more than they paid for the option. With that understanding, even if the variance is high, an investor can significantly gain ground if the price moves in the direction that they want. If there are dividends involved with the underlying asset then this will also affect the value of the option. If you want a call option, this will decrease as the dividends increase. If you want a put option, this will decease as the dividends increase[76].

The strike price of the option will also be important when determining the overall value of the option. When the strike price increases for a call, this will decrease the value of the call because it will not go into effect as soon as a lower strike price would. While for a put, the value of the put will increase as the strike price increases. As the date to expiry increases this will increase the value of the option. This is because this gives the market more time to shift whether it's in your favor or not. It is also important to compare the interest rates. When dealing with this portion of the option, the riskless interest rate comes into play. The exercise price for an option doesn't have to be paid until expiry so that means that the interest rate will have a direct effect on the value of the options.[77]

Let's establish an example to further exemplify what it means to put or call. When you price an option, you are expecting the market to do something that you can do in your favor.

[76] Damodaran, Aswath. "Option Pricing Theory and Models."
[77] Damodaran, Aswath. "Option Pricing Theory and Models."

When purchasing an option, you make a deal with the option provider. In the case of a call, you would expect the market price to increase. Let's say that company XYZ's stock is trading at $50 today. You expect that within the next year that it is going to increase much farther than that. You then go to an entity that sells options. This entity will then price out the option that you want to buy. For this case, we'll say it's a three-month option. You also have to set up a strike price at which you can exercise your right to buy the stock for a set price. For this exercise, we'll say that the strike price is $55.

This leads up to, if the market goes above $55, you are able to purchase the stock at $55 even though the market is above $55. That way you can turn around and sell it for cheaper. However, if the market never reaches $55 then your option is worthless. Also, you have to make up for the price of the option itself. Therefore, even if the market goes above $55, the stock will have to go even farther to set off the price of the option. For this case, the option will be for one hundred units of stock. The price of the option has been set at $165.

Figure 1 - Profit as a stock price increases from buying a call option.[78]

It can be easily seen from the figure above that as the price of the stock goes up the initial cost of the option will eventually break even. This cross happens at approximately $56.70. The point at which this cross occurs is called the break-even point. That is where if you bought the stock and then sold it, you would have a total profit of zero. The goal is then to hope that the price goes even further than it already has in order to obtain some profit. After the break-even point is when the exercising the option will lead to a profit.

For a put option, it's similar however different. For the put option, this will be applied to stock you already have. Let's assume that you have a hundred units of stock in XYZ. There is an expectation for this stock to fall, therefore you go and pick up a put option. This put option will enable you to sell your stock at a price that is predetermined even if the stock price drops tremendously. We use a similar set up, however we'll set the strike price at $45. There is still the assumption that the stock price is at $50, however now the general thought is that the stock price

[78] Compiled sources: Vashishth, Abhishek. Personal interview. 30 Oct. 2007. and "Long Call." Option Trading Tips. 2005. 31 Oct. 2007 <http://www.optiontradingtips.com/strategies/long-call-option.html>. and Dickson, John. "Rogue Trading?" Plus Magazine. Sept. 2001. 31 Oct. 2007 <http://plus.maths.org/issue16/features/derivatives/index.html>.

will go down. This option is taken out for a three-month period. We'll set the price of the option at $165 just like the call.

Figure 2 - Profit as a stock price increases from buying a put option.[79]

The figure above shows that if you have the option and your stock continues to fall that you will be able to sell your shares at the strike price. As the stock continues to fall, your profit range will increase. That is to say that if the stock price takes a nose dive then you'll be able to sell your stock at $45 and try to make a profit. This would be possible if you had purchased the stock before the dive. Perhaps, you purchased the stock when it was $35 and then it rose to $50. Therefore, you can either sell your stock at $50 before the drop or you can wait for the drop, buy the option before the drop, and have a guaranteed $45 because of the option that you purchased. This enables you to hang onto your stock even if you're worried that it might fall. An investor could combine this with a call option. This enables them to bet in both directions.

[79] Compiled Sources: "Long Put Option." Option Trading Tips. 2005. 31 Oct. 2007 <http://www.optiontradingtips.com/strategies/long-put-option.html>. and "Long Put / Put Purchase / Buy Put Option." Trade Soft. 31 Oct. 2007 <http://www.trader-soft.com/option-trading/option-strategies/put-long.html>.

This approach of combining a call and put option together is a form of hedging. There is a need to isolate from risk, thus one will simply bet in both directions in an attempt for each other to cancel out and leave only the profit. There are many ways to go about hedging as each will allow an investor to protect their investment and protect themselves from risk. The market is a very volatile entity but if precautions are taken then impacts of the market's volatility can be kept in check.[80]

2.1.10: Pricing Credit Derivatives

2.1.10.1: Reduced-Form Model Approach for Pricing Credit Derivatives[81]

The pricing of a credit default swap has two parts to it, which are the fee payments and default payment. Present value of a swap is the overall value of these two parts. There are a few ways of pricing credit default swaps. This is a pricing methodology using the reduced-form model approach. The variables that will be used in this part are T, t, $q(t)$, R, $u(t)$, $e(t)$, w, s, $v(t)$, π, and $A(t)$. This approach was originated by Duffie (1999) while the one we're using is specified by Hull and White (2000). The claim amount is also assumed to be face value plus any additional interest. There is also an assumption that this is being run for a one dollar principal.

[80] Dickson, John. "Rogue Trading?" Plus Magazine.
[81] This entire section is based off of Choudhry, Moorad. pp. 26-28

Term	Definition
T	The total time of the CDS in years.
q(t)	Risk-neutral probability at time *t*
R	The expected recovery rate as if it was in a risk-neutral world
u(t)	Present value of payments at the rate of $1 per year for payments between 0 and t.
e(t)	Present value of a periodically adding payment at time t equal to t-t* where t* is the payment date immediately before time t
v(t)	Present value of $1 received at time t
w	Total payment per year made by CDS buyer
s	Value of *w* that causes the CDS to be equal to zero
π	Risk-neutral probability of no credit event during the total time of the swap.
A(t)	Interest that is added periodically onto the reference obligation at time t and taken as a percentage of face value
t	Time this is being taken

Table 6 - List of definitions for the following equations

Thus, we know that π is the probably of the credit event not occurring. This gives us that π is one minus the probability of it happening. This gives us the first formula, which allows us to solve for π. This is done by integrating *q(t)* as shown below.

$$\pi = 1 - \int_0^T q(t)dt$$

Formula 5 - Equation for the probability of no credit event

$$w\int_0^T q(t)[u(t) + e(t)]dt + w\pi u(T)$$

Formula 6 - Expected present value

$$1 - R - A(t)R$$

Formula 7 - Risk-neutral expected payoff from CDS

$$\int_0^T [1 - R - A(t)R]q(t)v(t)dt$$

Formula 8 - Present value of the expected payoff from the CDS

$$\int_0^T [1 - R - A(t)R]q(t)v(t)dt - w\int_0^T q(t)[u(t) + e(t)]dt + w\pi u(T)$$

Formula 9 - Combined form of previous equations

$$s = \frac{\int_0^T [1 - R - A(t)R] q(t) v(t) dt}{w \int_0^T q(t)[u(t) + e(t)] dt + \pi u(T)}$$

Formula 10 - Rearranged form of above to solve for *s*

The above was found by first looking for π. This was necessary in order to continue onto the next step where we can calculate the expected present value. In order to understand the next formula, one has to realize that the present value if there is no default is *wu(T)* while if there is a default the present value is *w[u(t)]*. Therefore, it then follows that we are left with a way for solving for the expected present value. This model specifically makes the assumption that there is no counterparty risk of default. It will then follow that default probabilities, interest rates, and recovery rates do not depend on one another. As stated above, the claim amount in the case of default is assumed to be face value plus any additional interest that had been accumulated.

It is then clear to see that we need to know what the risk-neutral expected off is. This is done by the step in the next equation. This shows the recovery rate being multiplied by a number that can be easily determined as equal to or greater than one. This is because the interest should always be positive for this. Therefore, we see that as *A(t)* increases, the risk neutral expected payoff will shrink. Thus, we then try to obtain the expected payoff from the CDS by integrating from zero to *T*. This also combines the risk-neutral probability of *t*. This also has to include how much the dollar is worth at that time. This clearly gives an answer in dollars. That makes sense as this is an expected value. Now, we know that if we combine the present values together by subtracting the present value of the payment from the present value of the expected payoff, that we have all we need to know in order to find the default swap spread. This is done by figuring out what value of *w* makes the equation equal to zero. This ultimately gives us a compact form for being able to price CDS.

2.1.10.2: The Black-Scholes Model

Fischer Black and Myron Scholes were interested in finding a way to deal with the problem of options. This was because an option allowed you to be only on the upside of things, thus starting to control risk. They worked on other models and quickly weeded out anything that they saw as immeasurable. Eventually, they were left with what they had thought was actually important to pricing an option. These were the stock price, its volatility, the strike price, how long the contract is, the rate of interest, and how risky the option is. However, this left them with an immeasurable variable. This left them with risk. In order to get around this, they used a process called dynamic hedging, which allows someone to reduce risk where fluctuations in a portfolio cancel each other out.[82]

Before delving into the Black-Scholes model further, it is important to understand what assumptions were made in order to arrive at it. The first assumption is that the stock does not pay any dividends while the option is good. However, a lot of companies will pay dividends. General Electric is an example of a company that almost always pays a dividend. A way of getting around this limitation is to subtract the reduced value of a future dividend from the original stock price. Second, European terms are used for when an option can be exercised. The difference between European terms and American terms are European terms can only be exercised at the end of the option while American terms can be used at anytime. This isn't seen as a big problem because most investors will not exercise their call option until the very end of the option's life.[83]

They also assumed that markets are efficient. This is based on that an individual cannot consistently predict how the market is going to behave. This assumption is valid as economists

[82] "NOVA | Transcripts | Trillion Dollar Bet." PBS. 8 Feb. 2000. 26 Oct. 2007 <http://www.pbs.org/wgbh/nova/transcripts/2704stockmarket.html>.
[83] Rubash, Kevin. "Black - Scholes -- Option Pricing Models." Bradley University. 26 Oct. 2007 <http://bradley.bradley.edu/~arr/bsm/pg04.html>.

after 1929 wanted to find out if they could figure out how stock prices would move. They ended up picking stocks at random and throwing darts at the Wall Street Journal while blindfolded. When all was said and done, the random picks outperformed the top traders.[84] Black and Scholes turned to a different type of calculus to deal with the randomness. They used the foundations of Itô calculus in order to do their calculations in continuous time. As the Itô process itself is just a Markov process but in continuous time.[85]

Black and Scholes also assumed that no commissions were ever charged in their model. This restriction affects the model differently for different people. That is to say that floor traders pay a very small fee usually. So this assumption is minor to insignificant for them. On the other hand, an individual investor usually has to pay a much larger fee. This makes the model flimsy as the final product would be a distorted value. Also, they set interest rates to be known and unchanging. Even though there isn't a risk-free rate there are thirty-days-until-maturity U.S. Government Treasury Bills. However, if the interest rates are changing rapidly in the market then this assumption gets thrown out the window and breaks the model.[86]

The last assumption is that all returns are lognormally distributed. This says that underlying stocks have returns that are distributed normally. This finally gives the foundation for the Black-Scholes model. This model gives the ability to price a call option. There are two parts to the model. They are the expected benefit from obtaining a stock completely and the present value of paying the exercise price on the day of full maturity.[87]

---

[84] "NOVA | Transcripts | Trillion Dollar Bet." PBS. 8 Feb. 2000. 26 Oct. 2007 <http://www.pbs.org/wgbh/nova/transcripts/2704stockmarket.html>.
[85] Rubash, Kevin. "Black - Scholes -- Option Pricing Models."
[86] Rubash, Kevin. "Black - Scholes -- Option Pricing Models."
[87] Rubash, Kevin. "Black - Scholes -- Option Pricing Models."

Term	Definition
V	Value of the option
S	Current price of the asset
T	Time until expiry
K	Strike price / the price at which you have to right to buy the asset
R	Risk-free interest rate
N	Standard normal distribution
σ	Standard deviation of the stock returns
d_1	$\frac{\ln\left(\frac{S}{K}\right)+\left(r+\frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}$
d_2	$d_1-\sigma\sqrt{T}$

Table 7 - List of definitions for Black-Scholes model[88]

$$V_c = SN(d_1) - Ke\text{^}(-rT)N(d_2)$$

Formula 16 - Solves the call option price for the Black-Scholes model[89]

$$V_p = Ke^{(-rT)}N(-d_2) - SN(-d_1)$$

Formula 17 - Solves for the put option for the Black-Scholes model[90]

This enables a person to do is to effectively price options. Also, a key point on the Black-Scholes model is that it uses the standard normal distributions. So when the current stock price is much higher than the price of exercising the option, the normal distribution is very close to one. While if the stock price is a lot lower than the exercise price, the option will cost very little.[91]

When dealing with lognormal distributions, it is important to realize that the expected return and volatility both carry assumptions in order to make it work. The expected return is how much the investor of a return on average in small amount of time. The volatility of is a measure of the proportional stock price shifts uncertainty. The higher the risk, return, and the risk-free interest rate will generate a higher expected return for a stock. To figure out a ball-park figure of volatility one can actually use a formula. In this case, n is the number of observations, S_i is the

[88] "NOVA Online | Trillion Dollar Bet | the Formula That Shook the World." PBS. Nov. 2000. 26 Oct. 2007 <http://www.pbs.org/wgbh/nova/stockmarket/formula.html>.
[89] "NOVA Online | Trillion Dollar Bet | the Formula That Shook the World." PBS.
[90] Janusz, John. "Derivative Pricing Theories: Welcome." John's Finance Page. 26 Oct. 2007 <http://members.allstream.net/~johnjaz/bsm.htm>.
[91] "NOVA Online | Trillion Dollar Bet | the Formula That Shook the World." PBS.

stock price rated at *i*, **u** is the mean of the u_i *results*, and *T* is the length of time interval (rated in years).[92]

$$\sigma = \frac{\sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(u_i^2\right) - \frac{1}{n(n-1)}\left(\sum_{i=1}^{n} u_i\right)^2}}{\sqrt{T}}$$

Formula 18 - Equation for solving for volatility[93]

Shown in the formula above is the equation for being able to derive a value for volatility. This becomes extremely effective when evaluating equations that are derived from the Black-Scholes model. This also is able to put a value to something that appears to be hard to accurately describe.

2.1.10.3: Derivation of the Black-Scholes Partial Differential Equation[94]

As with all partial differential equations, it is important to understand the boundary conditions that are associated with the problem you're trying to solve. We know that the value of an option (*V*) is a direct consequence of the underlying assets value (*S*) and the current time period (*t*). When the option reaches its maturity and if the stock has risen, then the option at that point is worth $S - K$. *K* is the strike price at which the option can be exercised. This means when the option is able to be used. The profit would be $S - K$ in this case. This would be the value of the call at the time that it is exercised. It's important to note that if *S* does not reach *K* then the call option won't be used and its value would be equal to zero. Thus, we understand that when $S \leq K$ that *V* is equal to zero. Also it's important to note that there is also a boundary condition on *S*. When $S = K$, the implied boundary condition is that this function has to be continuous.

---

92 Janusz, John. "Derivative Pricing Theories: Welcome."

93 Janusz, John. "Derivative Pricing Theories: Welcome."

94 This section is based on Silverman, Dennis. "Solution of the Black-Scholes Equation Using the Green's Function of the Diffusion Equation." 20 May 1999. University of California, Irvine. 1 Nov. 2007 <http://www.physics.uci.edu/~silverma/bseqn/bs.pdf>.

We have to establish that we are looking for the net equity. This can be found using the equation below. Also, it is important to understand that *S* is value of the reference asset while *V* is the value of the option. Thus, we understand that the value of the asset is affected by the changes in the value of the asset and also the changes in the value of the option.

$$\boldsymbol{S - V\frac{\Delta S}{\Delta V}}$$

Formula 11 - Net equity for the derivation of the Black-Scholes model

The next step involves an understanding that we change *S* to ΔS while *V* will also change to ΔV. This is done to keep the equity the same. Thus we can use this to put everything into terms of deltas because we're looking at changes in the stock and the asset.

$$\boldsymbol{\Delta S - \Delta V\frac{\Delta S}{\Delta V} = 0}$$

Formula 12 - Change in equity due to changes in S and V

We know that for the Black-Scholes equation that *V* is a function of *S* and *t*. Thus, time and the actual value of the reference asset will directly affect the value of the call option or any option for that matter. For the next step, we understand that it is necessary to hedge the equity. Hedging allows us to protect ourselves against risk. For this case, we are looking for a neutral hedge where we protect ourselves in both directions equally. Thus, what happens is that we have to take into account that a fixed-interest rate comes into the equation. This fixed-interest rate is *r*.

$$\boldsymbol{\Delta S - \Delta V\frac{1}{\frac{\partial V}{\partial S}} = \left(S - V\frac{1}{\frac{\partial V}{\partial S}}\right)r\Delta t}$$

Formula 13 - Compensating for a fixed-interest rate

Then we take the Taylor expansion of ΔV in order to continue with the mathematical derivation. This gives us a specific term that needs some extra attention. $(\Delta S)^2$ comes with the assumption that the variance of this has to account for some random walk. Random walk comes

from the assumption that stocks follow a random path whether it be increasing or decreasing. This random walk also has to be proportional to Δt because the variance will be established over time.

$$\frac{\partial V}{\partial S}\Delta V + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}(\Delta S)^2 + \frac{\partial V}{dt}\Delta t$$

Formula # - Taylor Expansion of ΔS

$$(\Delta S)^2 = \sigma^2 S^2 \Delta t$$

Formula 14 - Accounting for the variance in $(\Delta S)^2$

The above formula includes σ, which is the volatility of the market. Thus, we can combine all of this established knowledge to come to the Black-Scholes partial differential equation for a European call option. We then end up being able to cancel out ΔS, dividing by Δt, and ultimately multiplying by $\partial V / \partial S$.

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0$$

Formula 15 - Black-Scholes PDE for a European call option[95]

2.1.10.4: The Greeks

	Call	Put
Delta	$\Delta = N(d_1)$	$\Delta = N(d_1) - 1$
Gamma	$\Gamma = \frac{\frac{1}{\sqrt{2\pi}} e^{-\frac{d_1^2}{2}}}{S\sigma\sqrt{T}}$	
Theta	$\theta = \frac{S\frac{1}{\sqrt{2\pi}} e^{-\frac{d_1^2}{2}}\sigma}{2\sqrt{T}} - rKe^{-rT}N(d_2)$	$\theta = \frac{S\frac{1}{\sqrt{2\pi}} e^{-\frac{d_1^2}{2}}\sigma}{2\sqrt{T}} + rKe^{-rT}N(-d_2)$
Vega	$\nu = S\sqrt{T}\frac{1}{\sqrt{2\pi}} e^{-\frac{d_1^2}{2}}$	
Rho	$\rho = KTe^{-rT}N(d_2)$	$\rho = -KTe^{(-rT)}N(-d_2)$

Table 8 - The Greeks for European Options[96]

[95] Compiled from: Lai, Yongzeng, and Jerome Spaingler. "Applications of Monte Carlo/Quasi-Monte Carlo Methods in Finance: Option Pricing." Smart Quant. 30 Oct. 2007 <http://www.smartquant.com/references/MonteCarlo/mc6.pdf>. and Cheng, Steve. "Black-Scholes PDE." Planet Math. 30 Oct. 2007 <http://planetmath.org/encyclopedia/BlackScholesPartialDifferentialEquation.html>. amd Silverman, Dennis. "Solution of the Black-Scholes Equation Using the Green's Function of the Diffusion Equation."

	Formula
Delta	$\Delta = \frac{dV}{dS}$
Gamma	$\Gamma = \frac{d^2V}{dS^2}$
Theta	$\theta = -\frac{dV}{dT}$
Vega	$\nu = \frac{dV}{d\sigma}$
Rho	$\rho = \frac{dV}{dr}$

Table 9 – The Greeks in Derivative Form[97]

Also coming out of the Black and Scholes model are some of the most important financial formulas. The main equations are named after Greek letters except for one. The main formulas are Delta, Gamma, Theta, Vega, and Rho. Delta is used for finding out how susceptible the option value is to minor changes in the share price. The Delta is also called the hedge ratio. Delta ranges from -100 to +100. The Delta can either be from 0 to 100 or from -100 to 0. For calls the range goes from 0 to 100 while for puts the range goes from -100 to 0.[98]

For calls that range from 0 to 100, this means that they are either far out of the money or deep in the money. A 0 when computing a call will mean that it is far out of the money. Being far out of the money, in this context refers, to the strike price of a call option being higher than the underlying futures price. While being in the money refers to a call option that has a strike price that is lower than the underlying futures price. For puts that range from -100 to 0, things are similar but still different. A value of 0 is still considered being far out of the money and -100 is considered being deeply in the money. However, in this case the situation turns as an out of the money put has a price lower than the underlying futures price. Being in the money means that the strike price is higher than the underlying futures price for puts. Being at the money is when the

[96] Janusz, John. "Derivative Pricing Theories: Welcome."
[97] Wystup, Uwe. "Vanilla Options - Chapter 1." Math Finance. 26 Oct. 2007 <http://www.mathfinance.de/FXRiskBook/chap-1.pdf>.
[98] Janusz, John. "Derivative Pricing Theories: Welcome."

option has a strike price which is close to the underlying futures price. For a put the rating is a -50 while for a call the rating is fifty. The underlying contract always has a delta of one hundred.[99]

Gamma is how sensitive the Delta is to small changes in share price. This function is also sometimes called the curvature of an option. As with any curvature, this gives the rate at which something increases or decreases. In this case, it is the rate at which an option gains or loses value when compared to the underlying contract. Theta is the option value's sensitivity as time progresses.[100] This measures the effect of time on the values of puts and calls. Sometimes it is even referred to as the time decay factor because it shows how an option loses value as time progresses. A short term at-the-money option will decay more rapidly than a long term at-the-money option. [101]

Vega is not a Greek letter; however it is included in this group because it measures the same sort of thing as the others. It is sometimes called kappa to keep with the Greek name and is sometimes represented as nu.[102] Vega like Gamma measures a sensitivity of a sensitivity. Vega is used to calculate the option value's sensitivity to small changes in the volatility. Options always increase in value when volatility increases. Consequently, for calls and puts Vega is always positive. Vega also decreases as the option nears expiration because the volatility is also decreasing. Out-of-the-money options will generate the largest Vega.[103]

Rho is the sensitivity of an option's value to small adjustments in interest rates. Futures options have a negative Rho. When interest rates rise, this will decrease the value of the option

---

[99] Janusz, John. "Derivative Pricing Theories: Welcome."
[100] "Theta." Risk Glossary. 26 Oct. 2007 <http://www.riskglossary.com/link/theta.htm>.
[101] Janusz, John. "Derivative Pricing Theories: Welcome."
[102] "Vega." Risk Glossary. 26 Oct. 2007 <http://www.riskglossary.com/link/vega.htm>.
[103] Janusz, John. "Derivative Pricing Theories: Welcome."

because it will cost more to be able to fund the option. Rho is typically not used to manage risk for futures or set up an options plan. Instead it is used for preparing strategies for options and managing risk. It helps manage the risk for things that can be affected by interest rates such as currencies, physical goods, and stocks.[104]

2.1.10.5: Binomial Theory for Options Pricing

This pricing theory is based on an idea that an asset price process can only go to one of two possible options. This involves construction of a tree of connected prices. Therefore, binomial theory sets up the idea that if you start with a current stock price it can either move up with a probability of *p* or it can move down with a probability of one minus *p*.[105] This method has been proven over time to be one of the best models. This is because it is very intuitive and flexible, which has lead it to becoming popular. It does not include some of the assumptions that the Black-Scholes model makes as it does not assume that markets are perfect and arbitrage opportunities. This approach can also be used as a shortcut to evaluating European options.[106]

This way of pricing options is usually applied to European options and American options. Both are these can be either with or without dividends. American calls that don't pay dividends will often act like European options and won't be exercised until expiry.[107] When dealing with a multiple step binomial tree that spans over multiple time periods, one has to start at the last time period and move backwards in time. They continue to move backwards in time until they reach

[104] Janusz, John. "Derivative Pricing Theories: Welcome."
[105] Damodaran, Aswath. "Option Pricing Theory and Models."
[106] Rubinstein, Mark. "Chapter 4. Binomial Option Pricing Model." In the Money. 1998. 26 Oct. 2007 <http://www.in-the-money.com/presentation/sld063.htm>.
[107] Janusz, John. "Binomial Pricing Theory" John's Finance Page. 26 Oct. 2007 <http://members.allstream.net/~johnjaz/bsm.htm>.

the current place in time. This means that one will start in the future and go period-by-period to travel backwards in time.[108]

It is also important to note that the value of an option using this method is not determined by the expected price but it is determined by the current price. Arbitrage is the cause of this.[109] Arbitrage is when someone purchases an asset and then sells the asset at the same time in order to take advantage of a difference in rate in the market place. This can also be called riskless profit.[110] The ones who actually take advantage of arbitrage are usually people with very low transaction costs and can do large amounts of transactions. This is because the discrepancies are usually small, which would in turn lead to little profit being made after the transactions have gone through.[111]

Although the binomial model gives an easy to understand background, it still requires a lot of information to be able to process an option price. So, to combat this, calculus comes into the thought process. The time iterations will become smaller. This will lead to the changes in prices becoming smaller. Then on top of this, we take the limit as time approaches zero. This will force the prices to become infinitesimally small. This then becomes a continue price spread. The Black-Scholes model is a special case of the binomial theory when the pricing becomes continuous. While if we say the prices actually stay large, this leads to a jump price process. There has been attempts made at using the jump price process, however any benefit that might be gained from using this model is usually drowned out by the parameters of the jump process. This is because the jump process parameters are noisy.[112]

---

[108] Damodaran, Aswath. "Option Pricing Theory and Models."
[109] Damodaran, Aswath. "Option Pricing Theory and Models."
[110] "Arbitrage." Investopedia. 30 Oct. 2007 <http://www.investopedia.com/terms/a/arbitrage.asp>.
[111] "Arbitrage." Risk Glossary. 30 Oct. 2007 <http://www.riskglossary.com/link/arbitrage.htm>.
[112] Damodaran, Aswath. "Option Pricing Theory and Models."

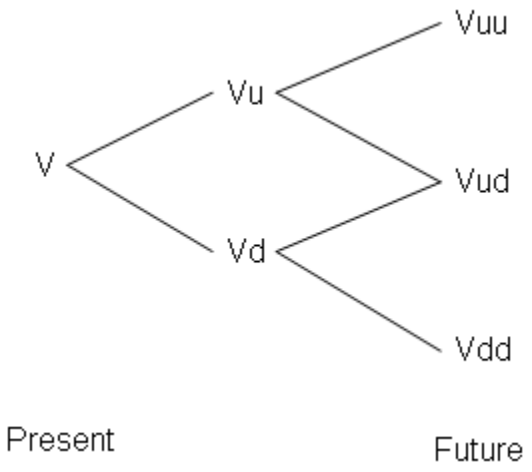

Figure 3 - Sample binomial theory tree[113]

The figure above is a simplified binomial theory tree. The time iterations are considered to be rather large so that there will be a significant change in price. V is the value of the asset and the addition of u or d signifies if it went up or down in value. V_{ud} isn't exclusive to moving upwards and then downwards. It can also be reached by going downwards and then upwards. Now we can define by how much gain or loss there is for going up or down in the tree. For simplicity, a move in the up direction will constitute a gain of 10% while a move in the down direction will constitute a loss of 10%.

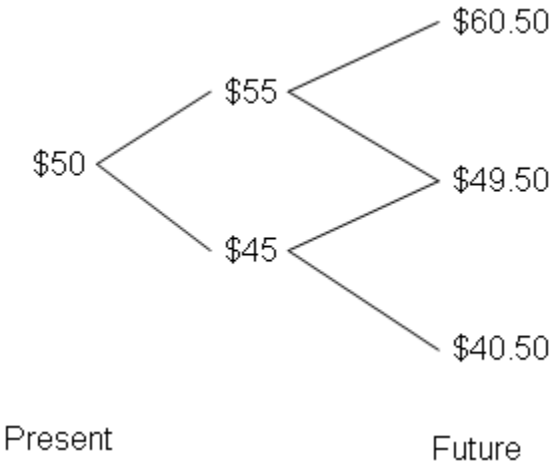

Figure 4 - Example of binomial theory

[113] Compiled from Damodaran, Aswath. "Option Pricing Theory and Models." and Janusz, John. "Binomial Pricing Theory." John's Finance Page. 30 Oct. 2007 <http://members.allstream.net/~johnjaz/bt.htm>.

The figure above shows that the price starts at $50. As such, the price can go up or down for the first time iteration. We have declared that first time iteration is a ten percent increase or decrease, so that gives us our two branches. Then based on how the market behaves the price will increase to one of the three possible outcomes, which is demonstrated in the next set of branches. This is a simple two step binomial tree. For this example, the strike price of the option will be $51. This makes it easy to see that the option only has value when the price goes up in this instance. This is because if the branch went up and then down the final price would be below the $51 threshold. However, if the price went up and then up again, this would up being higher than the threshold. This would give the option a price as an option below the strike price doesn't have any value.

In order to get the price of this option requires a little bit more work. To figure this out, we set up a value that we'll call O. Now O has subscripts that tell how it was calculated. O_{uu} is equal to $60.50 - $50. This number is positive so it stays. Then we're also interested in O_{ud}, which is equal to $49.50 - $50. This value is negative so the option value, O, will be equal to zero. Thus, we can solve for the probability of an upward movement. The probability for a downward movement will just be one minus *p*. Then with the probability of an upward movement, we can calculate the estimated value of the level before it.

$$p = \frac{e^{rT} - D}{U - D}$$

Formula 19 - Solving for the probability of an upward movement[114]

T in this case is the time between periods in years while *r* is the risk-free interest rate, which we will set to as 8%. For this formula to work, we have to use the multiplier version of the percentage. Therefore, our 10% increase will be 1.1 while our 10% decrease will be 0.9. *U* is the

[114] Janusz, John. "Binomial Pricing Theory."

amount of an increase, which is why we use the multiplier form instead of the percentage form. *D* is similarly the decrease caused by a downward movement. For this example, *T* is made up of two 6-months periods, which will make it 1 year. This will give us a probability of approximately 70% while the probability of decrease is approximately 30%.[115]

$$EV(V_u) = P_u(O_{uu}) + P_d(O_{ud})$$

Formula 20 - Solving for the expected value of the previous node[116]

We know that V_{ud} is 0 because this is under the strike price and therefore the option is worthless. So by applying the above formula, this will give us $6.21. We then have to take this expected value and discount the amount that we expect to actually give us the present value. We simply multiply this by e^{-rT} to give us our present value, which we calculate to be $5.96. The call option price will follow the same as in formula # but instead of O_{uu} and O_{ud} it will be O_u and O_d respectively. Also instead of being the expected value it will be the actual value of the option. This gives us a value of $4.20 for the price of the option. So we can now place the values of the option on the tree.[117]

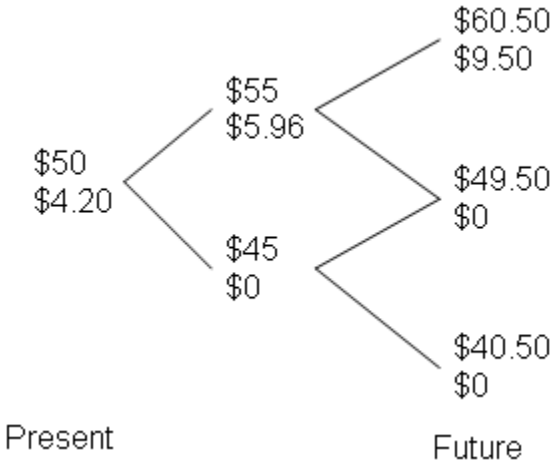

Figure 5 - Call option prices included with goal prices

[115] Janusz, John. "Binomial Pricing Theory."
[116] Janusz, John. "Binomial Pricing Theory."
[117] Janusz, John. "Binomial Pricing Theory."

As seen in the figure, the goal prices are set above by saying what the increase would be if there was one for each iteration. Then underneath those are the prices of the call if they were taken out at that point in time. This can be extended out infinitely to form the basis of the continuous price spread. The approach would still be the same; however it would involve calculus to deal with the continuous time span. This is just the basis for binomial theory as an entire project could be dedicated to just understanding the binomial theory in continuous time.

2.1.10.6: Monte Carlo

Monte Carlo simulations are a way of simulating random values that come out of a model.[118] It is also a way of being able to put a lot of random values into a model in order to get a probability density out of the model that's being used. This is one of the main ways of pricing derivatives.[119] Monte Carlo simulations are used often when it is tricky to derive a formula for pricing a derivative. An example of this is for pricing single variable European options that suffer from having a non-standard pay out form. Also another benefit to the Monte Carlo simulation is that when a model is developed, one doesn't have to change it completely. They can actually manipulate the variables in order to make the model suit one's needs.[120]

European options, following a similar trend as the binomial theory and Black-Scholes model, are easier to price. Monte Carlo simulations have been in use for European options for awhile now. On the other hand, American options are starting to be calculated by the Monte Carlo methodology also. Also, there are Asian options, which payoff depends on the history of the asset price. This is achieved by taking a time average.[121]

---

[118] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. p.227

[119] Janusz, John. "Monte Carlo Simulation." John's Finance Page. 30 Oct. 2007 <http://members.allstream.net/~johnjaz/mont.htm>.

[120] Fink, Peter. "Pricing Complex Options Using a Simple Monte Carlo Simulation." John's Finance Page. Comp. John Janusz. 30 Oct. 2007 <http://members.allstream.net/~johnjaz/montemain.htm>.

[121] Lai, Yongzeng, and Jerome Spaingler. "Applications of Monte Carlo/Quasi-Monte Carlo Methods in Finance: Option Pricing."

2.1.10.7: Cholesky Decomposition

Often with Monte Carlo simulations there is a need to correlate the random numbers that are generated. One of the most common ways is to use the Cholesky decomposition. This lends itself to the Cholesky matrix. This matrix is a triangular matrix. When given the correlation matrix R, it easy to calculate the Cholesky decomposition L.[122] It is important to understand that the matrix that precedes the Cholesky decomposition has to be symmetric and a positive definite matrix. The lower triangular matrix is defined to be the Cholesky triangle and it has positive diagonal elements.[123]

$$R = L * L^T$$

Formula 21 - The Cholesky decomposition obeys this relationship[124]

When given a correlation matrix, we are able to figure out the Cholesky decomposition. The L^T is the mathematical transpose of the Cholesky decomposition. This goes to show that if you have a correlation you can set up a Cholesky decomposition or if you have a distribution, you can solve for a correlation. Doing this in matrix form, allows you to see an overall correlation of two or more assets.[125] In order to go about this the proper way, to find out how to solve for $Rx=b$, one has to solve for $Ly = b$ for y. This is allows one to solve for the lower transpose. Then one has to solve for $L^Tx = y$ to be able to solve for x.[126]

It is possible to approach this situation differently and use a method of finding eigenvalues and eigenvectors. However, the Cholesky decomposition is a lot simpler and far easier to do. The results come out very similar.[127] Based on this, it is easy to see why people

---

[122] Chaplin, Geoff. p.43

[123] Krowne, Aaron. "Cholesky Decomposition." Planet Math. 4 June 2006. 1 Nov. 2007 <http://planetmath.org/?op=getobj&from=objects&id=1287>.

[124] Chaplin, Geoff. p.43

[125] Chaplin, Geoff. pp.43-44

[126] Bock, Rudolf K. "Cholesky Decomposition." 7 Apr. 1998. 1 Nov. 2007 <http://rkb.home.cern.ch/rkb/AN16pp/node33.html#SECTION00033000000000000000>.

[127] Team Latte. "All About the Cholesky Matrix." Risk Latte. 9 Aug. 2005. 1 Nov. 2007 <http://www.risklatte.com/portfolioEngineering/article20050809.php>.

would prefer to use the Cholesky decomposition over other methods. To further expand on this, here is a simple example of how to apply the Cholesky decomposition to a correlation. In this example, we will set up a general process to help understand what is happening during this process.

$$\begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix} = \begin{bmatrix} l_{11} & 0 \\ l_{21} & l_{22} \end{bmatrix} \begin{bmatrix} l_{11} & l_{12} \\ 0 & l_{22} \end{bmatrix}$$

Formula 22 - Correlation Matrix set equal to a Cholesky decomposition[128]

In the above equation, it is easy to see how to deal with the above situation. We follow the standard procedure for multiplying matrices. This is done by either following the formulas provided below or by performing the multiplication out by hand. The formulas provided below work for any size matrix.

$$L_{ii} = \sqrt{(R_{ii} - \sum_{k=1}^{i-1} l_{ik}^2)}$$

Formula 23 - Equation for solving for the diagonals[129]

$$Lji = \frac{R_{ji} - \sum_{k=1}^{i-1} l_{jk}\, l_{ik}}{l_{ii}}$$

Formula 24 - Equation for solving the other components[130]

A direct corollary that follows this is that if we know the probabilities of an event happening, we can solve for the correlation between multiple stocks. This is why Monte Carlo simulation can be a very useful tool because we can connect the probabilities of events occurring with a correlation. This is a very useful way of linking up Monte Carlo simulations with aspects of finance that can prove to be useful. Also, if there is no correlation, Monte Carlo simulations can also be used but instead it will just yield different results. These results will ultimately show for no correlation that fair spread percentage is larger than the correlated. If there is a perfect

[128] Compiled sources: Chaplin, Geoff. pp.44 and Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp.232-233
[129] Compiled sources: Bock, Rudolf K. "Cholesky Decomposition." and Krowne, Aaron. "Cholesky Decomposition." Planet Math.
[130] Compiled sources: Bock, Rudolf K. "Cholesky Decomposition." and Krowne, Aaron. "Cholesky Decomposition." Planet Math.

correlation between all of the reference assets, then the chances of none of them defaulting is greater than the zero correlation. Also under perfect correlation the chance of all of the companies defaulting is much greater than with 0 correlation.[131]

[131] Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. pp.234-242

2.2: Risk Engine Review

2.2.1: Multiple Core Processing

To understand multiple core processing, one has to first understand what a core is. When referring to a core in the computer sense, this is a complete central processing unit. A single core is a single processor that deals with all of the information. On the other hand, a multi-core processor has numerous execution cores that perform the calculations. A single core can only be pushed so far before it starts hitting physical limitations. One of these limitations is as the frequency is increased there is a rise in the number of clock cycles it takes to access memory. The i480 took 6-8 cycles to access memory while the Intel Itanium takes 224 clock cycles to access memory. A more pressing concern is the power density. As the frequency increases, the amount of power generated is increased. With smaller architectures and more transistors constantly being added, the power density continues to grow rapidly as the frequency increases. In 1993, the Intel Itanium 2 processor had approximately 3 million transistors while new Itaniums are currently pushing almost 1 billion transistors.[132]

Eventually, the heat created by the processor becomes expensive when compared to the increase in performance. Many people focus a lot of their efforts in cooling their processors down so that they can overclock their processors. Cooling comes with a price though. Eventually, it becomes uneconomical to cool the processor as the amount it costs to cool the processor will outweigh the benefit of the increase in performance. To address this, we can introduce more precise manufacturing processes to reduce the amount of heat generated.[133] As the transistors shrink down, it takes less power to keep the transistors on. This means that energy will be saved

[132] Koch, Geoff. "Discovering Multi-Core: Extending the Benefits of Moore's Law." Intel. July 2005. Michigan State University. 19 Nov. 2007 <http://www.intel.com/technology/magazine/computing/multi-core-0705.pdf>.
[133] "Planning Considerations for Multicore Processor Technology." Dell. May 2005. 7 Nov. 2007 <http://www.dell.com/downloads/global/power/ps2q05-20050103-Fruehe.pdf>.

because the transistor is dissipating less power to provide the same effect.[134] This isn't sustainable, however, at some point it has to end.

Another way around the problem of increasing heat from higher frequencies is to introduce multiple core processors. These processors will often run at a lower frequency but they enable a divide and conquer approach. The power density is then split across the chip and is also reduced because of the lower clock speeds. This includes a push towards a smaller manufacturing process too as this enables more processors to be put on the same chip. It also carries with a particular consequence that has to be addressed.

This consequence of switching to multiple core processors is the need to write the code in a parallelized form. That is because a program will only run on one core if it is not designed to take advantage of the additional cores. A basic example of parallelization of code and dealing with multiple cores is a factory. You can only have a single factory worker work so hard. He can only do things so fast with incentives like a bonus and a pay increase. Pay increases and bonuses are like frequency increases and more efficient processing with more transistors. Eventually, he reaches a physical limit of how hard he can work, even with the pay increase and bonus incentive. Now, we add a second worker. This second worker will take the single factory's work and split the work with him.

However, if the split isn't even, this hinders the entire efficiency of the process as one worker will still be doing a lot more work. This is where the code would have to be optimized to share processes. That is to say, if both workers get an equal work load they will be able to finish their overall task faster. Eventually, as we continue to add on additional workers to aid in the work, we would have to split the work even further to optimize the process. If the work load is

[134] Mutschler, Ann S. "Intel Creates 2nd 65nm Manufacturing Process." Electronic News. 20 Sept. 2005. 19 Nov. 2007 <http://www.edn.com/index.asp?layout=article&articleid=CA6258496>.

evenly distributed the work will get done exceedingly fast. If the work load is not evenly distributed, this will have a similar effect as the two worker situation. It will result in an area that can be improved through proper allocation of tasks. The distribution of the work is not the only area of concern. There could be problems with the workers. Each one might require certain amounts of resources to achieve their objective but those specific resources may be currently used by another worker. The workers could also be in the same workshop, the same building, or just the same company. This lends the concern of where the processors are placed.

If there are processors on the same board, then their communication is limited by the speed of the motherboard. We can make the processors closer together but this only works so far. Eventually a turn towards multiple processors that are inside of the same die is required. This gives us another reason for the multiple core processors. This is what enables the multi-core processor to do work more efficiently than if there are just multiple single-core processors on a board.[135] Servers have traditionally been systems that have extensively used multiple processors in order to deal with the workload that they experience.[136] Now, computer gaming has reached a new level of demand. Games like Crysis and Supreme Commander demand immense computing power in order to give optimal performance.

2.2.2: Parallelization of Code

There are multiple models used for the parallelization of code. Examples of these models are shared memory, threads, message passing, data parallel, and combinations of any of these. Each method has its own strong points and weak points. These models can be employed on any type of system. This is because they are not exclusive of other types of systems' architecture. That is to say that the shared memory model, for example, can be used on a system that is

[135] Jacobs, Jason. "Xeon 5150 Vs. E5320." Techware Labs. 29 Mar. 2007. 7 Nov. 2007 <http://www.techwarelabs.com/reviews/processors/xeoncomparison/>.
[136] Kyrnin, Mark. "Dual-Core Processors." About. 7 Nov. 2007 <http://compreviews.about.com/od/cpus/a/dualcore.htm>.

specifically designed as a distributed memory system. The drawback is that it may not be optimal to use that specific model on that system. It might also not make sense to use that model as it would make programming far more difficult for the system.[137]

The shared memory model is where you have tasks that read and write out of sync and share a common address space. This model excels in the aspect that it doesn't establish an ownership of certain data. This ultimately means that the development process can be cut down and made easier by enabling the programmers not to have to specifically identify the communication of data during events.[138] This model allows for easy scalability and enables portability across different memory architectures. It is a good step towards developing highly refined algorithms because of these reasons.[139] People who program in a symmetric multiprocessing environment will often employ this type of model.[140]

However, a major disadvantage to the shared memory model is that it will eventually become difficult to manage the data.[141] The difficulty of managing the data is because when a processor requests information that is not in the cache memory, it will stop any activity until the information is retrieved. This aspect of the shared memory model makes managing the data an absolute necessity. This is because if the data is managed properly these will be slowdowns that could be avoided if the system had what the resources it required. The data management problem becomes an issue as the program size increases and this weighs heaviest on a program using this

---

[137] "Introduction to Parallel Computing." Department of Computer Science and Information Systems. University of Limerick. 8 Nov. 2007 <http://www.csis.ul.ie/Modules/CS4838/tutorialsforweb/Introduction%20to%20Parallel%20Computing.htm>.
[138] "Introduction to Parallel Computing." Department of Computer Science and Information Systems. University of Limerick.
[139] Gibbons, Phillip B. "What Good are Shared-Memory Models?" Proceedings of the 1996 ICPP Workshop on Challenges for Parallel Processing (1996): 103-114. IEEEXplore. George C. Gordon Library, London. 19 Nov. 2007. Keyword: shared memory model.
[140] Scott, Michael L. "Shared Memory Models." Apr. 2002. University of Rochester. 8 Nov. 2007 <http://www.cs.rochester.edu/research/cashmere/2002-04_pmodels_talk.pdf>.
[141] "Introduction to Parallel Computing." Department of Computer Science and Information Systems. University of Limerick.

method.[142] Most programmers prefer to use message passing interface instead of using this model as there aren't any large parallel applications that have been written using this model.[143] The lack of massive shared memory model programs only hurts because there isn't a large enough example of how to write a program using this method properly.

A thread (also known as thread of execution) is when a process is written to have multiple execution paths. That is to say that, the process doesn't have to be run sequentially. The computer can give the illusion of the processes running in parallel while it quickly switches between threads. This illusion of processes running in parallel is achieved by having the first process run a thread through the processor. After that thread is completed, a thread from the second process can be run through the processor. Then these processes will alternate which thread is being run through the system. This creates the illusion of two processes being run simultaneously. Each process does run slower than if it was being run by itself because it is sharing the computation time with the onboard processor.[144]

The nature of the thread is why it lends itself to being used on parallel computing apparatuses. This is because threads can be split up among the processors and be executed almost simultaneously by the processors. Single-threaded applications are not written to take advantage of multiple processors as these applications will only be executed on one processor. An obstacle for writing multi-threaded code is that the threads have to be able to be executed independently of the other thread's instructions. This is because if a single thread is dependent on another thread will need the information from another thread. The problem arises if the threads run at the

[142] Buede, Dennis M. Engineering Design of Systems - Models and Methods. (pp. 281). John Wiley & Sons. Online version available at: http://www.knovel.com/knovel2/Toc.jsp?BookID=1410&VerticalID=0

[143] Scott, Michael L. "Shared Memory Models."

[144] O'Sullivan, Bryan. "What is a Thread?" Teideal Glic DeisbhéAlach. 9 June 2005. 21 Nov. 2007 <http://www.serpentine.com/blog/threads-faq/what-is-a-thread/>.

same time. This will cause the information to not be available and the dependent thread won't be able to complete its task until the other thread is finished. This would serve as a point of optimization for the code.[145]

The message passing model involves programs that are running in parallel with processes that cooperate with one another. The processes will send data to one another in the form of messages.[146] This model is considered to be portable, universal, and simple. The portability comes from the fact that it is the model that is used on most parallel processing apparatuses. The model's universality is based on the fact that it makes very few assumptions on the underlying hardware. By not making these assumptions, the program can be written to fit more systems than assuming that a computer will have certain specifications. This model's simplicity comes from the fact that it supports explicit control of memory references. [147] It is also code that can be read and understood by a human. This is because when using the message passing model, the code should read just as if one was reading the actual message passing model's outline.[148] The message passing model is used extensively with supercomputers.[149]

The data parallelism model is a peculiar entity. It's designed around concurrency. Concurrency is when there are several processes executed at the same time. This form of programming requires that the programmers decide how the data will be distributed to run the most efficiently. This type of programming works really well with single instruction, multiple

---

[145] Shamlin, David. "Threads Unraveled: a Parallel Processing Primer." SAS Institute Inc. 22 Nov. 2007 <http://www2.sas.com/proceedings/sugi29/217-29.pdf>.

[146] Gropp, William, Rusty Lusk, Debbie Swider, and Rajeev Thakur. "The Message Passing Model." Argonne National Laboratory. 22 Nov. 2007 <http://www-unix.mcs.anl.gov/mpi/tutorial/mpibasics/sld003.htm>.

[147] "The Message Passing Model." Hewlitt-Packard. Dec. 2003. 22 Nov. 2007 <http://docs.hp.com/en/B6060-96013/ch01s01.html>.

[148] Ercal, Fikrey. "2.2 Using Workstation Clusters." CSc 387. University of Missouri. 22 Nov. 2007 <http://web.umr.edu/~ercal/387/class.2-MPI.html>.

[149] Barney, Blaise. "Message Passing Overview." MHPCC. 3 July 1996. University of Hong Kong. 8 Nov. 2007 <http://www.hku.hk/cc/sp2/workshop/html/message_passing/message_passing.html>.

data (SIMD) environments.[150] This is because the data parallelism model is mainly focused on performing operations on data sets. These data sets will often be organized into a common structure such as an array. Different partitions of the same data structure will then be manipulated by a set of tasks. Each data partition will have the same tasks run on them as the goal of this architecture is to divide the work while still continuing through the process.[151]

2.2.3: Floating Point Arithmetic

Floating point is a way of representing real numbers using bits. However, it is important to understand what a fixed point number is first. This will establish an easy recognizable difference. Fixed point numbers put a decimal point somewhere in the middle of the digits. This could even be at the end in the case of integers. The important thing to understand about this is that the decimal point does not move. Fixed points only allow a certain amount of numbers to be represented inside of a small window. Floating points, on the other hand, allow the decimal point to move. This allows floating point numbers to represent a much larger range of numbers depending on the amount of precision that is allowed.[152]

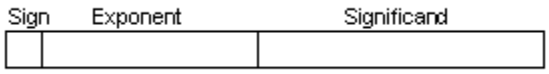

Figure 6 – The Floating-Point Format

Floating point numbers employ scientific notation to be able to shift in between large and small numbers. This allows the decimal point to float around and not get locked in like a fixed point representation. Floating point notation includes three major parts. These parts are the sign,

[150] Foster, Ian. "1.3 a Parallel Programming Model." Argonne National Laboratory. 1995. 8 Nov. 2007 <http://www-unix.mcs.anl.gov/dbpp/text/node9.html>.
[151] Barney, Blaise. "Introduction to Parallel Computing." Lawrence Livermore National Laboratory. 22 June 2007. 22 Nov. 2007 <http://www.llnl.gov/computing/tutorials/parallel_comp/>.
[152] Hollasch, Steve. "IEEE Standard 754 - Floating Point Numbers." 24 Feb. 2005. 23 Nov. 2007 <http://steve.hollasch.net/cgindex/coding/ieeefloat.html>.

the exponent, and the significand. A significand is the part of the floating number that contains the significant digits. The significand stores the J-bit and a binary fraction. The J-Bit is has an assumed value of 1 when the floating point number is normalized. The J-Bit is not actually stored inside of the physical digits. Normalized numbers in this context are non-zero, finite numbers that can be encoded to fit numbers from zero to infinity. However, with the physical limits of a format is limited by the amount of bits that are available for representation of the exponent. Denormalized numbers are numbers that are near zero and can only be represented by making the leading bits zero. This causes a loss of precision because of a reduction of significant digits by the leading zeros.[153]

The sign bit will denote that the number is either positive or negative. Positive numbers are represented by a 0 while negative numbers are a 1. The exponent block is able to represent both positive and negative numbers by applying a bias.[154] This bias is done by adding a specific number to the exponent. This bias has to be large enough to offset all of the negative numbers. This bias will depend on how many bits are available to represent the exponent. For a single precision number, the bias of the exponent will be 127 base 10.[155]

Figure 7 – Single precision floating number using the IEEE 754 standard

[153] "IEEE 754 Floating Point." Randelshofer Online. 23 Nov. 2007 <http://www.randelshofer.ch/fhw/gri/float.html>.
[154] Hollasch, Steve. "IEEE Standard 754 - Floating Point Numbers." 24 Feb. 2005.
[155] "IEEE 754 Floating Point." Randelshofer Online. 23 Nov. 2007

Figure 8 – Double precision floating number using the IEEE 754 standard

Figure 2 and figure 3 demonstrate the format that a floating number using the IEEE 754 standard would have to follow. The s is for the sign bit while the Es are for the exponent. The Ss are for the significand bits. The difference in precision is immediately obvious by looking at figure 2 and figure 3. The double precision floating number has the ability to store more significant digits than the single precision. Double precision has 29 more bits to represent the number with and it can also represent much larger numbers as the exponent range is increased by three bits. Even with all of these digits, there are numbers that cannot be represented in base 2. One of these numbers is 0.1. This number would require an infinite number of bits to be represented correctly.

We realize that not only because we're compressing an infinite amount of real numbers into a finite space but also because we are unable to properly represent a number in base two, that we should anticipate a rounding error. So just like a human, a computer will round a computed answer so that the number will be able to fit inside of the finite representation that our standards and technology allow. Error is not limited to the representation of a number but it also can come from performing a mathematical operation on the floating numbers. When performing a floating point operation it might be acceptable to round both floating point numbers and perform the operation on them. Another option is to add the numbers first and then round the

solution. If the numbers are significantly spaced apart by many orders of magnitude, these two methods should provide very similar results. [156]

If the numbers are close together, then a guard digit's value increases in importance. A guard digit's job is to save information that would disappear from the register normally. This extra information is used to mitigate the error where it is important to notice a small difference between two very similar numbers.[157]

2.2.4: Cell Processor

The Cell Broadband Engine is the product of a joint venture among IBM, Sony, and Toshiba. The Cell processor is a response to the need of increased efficiency and computing power with physical limits starting to weigh on conventional methods. The current Cell processor is made up of a POWER processing element (PPE), eight synergistic processing elements (SPEs), an element interconnect bus (EIB), a 32 KB data L1 cache, a 32 KB instruction L1 cache, and a 512 KB L2 cache.[158] In the future, it is possible that we'll see Cell processors with more or less SPEs depending on the application that they're being used for.[159] The current development roadmap is striving for a 1 TFLOP version of the Cell processor that will have 2 PPEs and 32 SPEs.[160] Toshiba and Sony have plans to use the Cell processors in their electronics such as televisions.[161] Sony currently uses the Cell processor in their Sony Playstation 3 ™.

---

156 Goldberg, David. "What Every Computer Scientist Should Know About Floating-Point Arithmetic." Numerical Computation Guide. Mar. 1991. 12 Nov. 2007 <http://docs.sun.com/source/806-3568/ncg_goldberg.html>.

157 Parlett, Beresford N. The Symmetric Eigenvalue Problem. SIAM, 1998. Google Book Search. 13 Nov. 2007 <http://books.google.com/books?id=uaYXlkHVPpEC&pg=PA31&lpg=PA31&dq=guard+digit&source=web&ots=06PkdJ-EeD&sig=dpRosxIj1U0HitNkMgl6sJ0gae0#PPA31,M1>.

158 Chen, Thomas, Ram Raghavan, Jason Dale, and Eiji Iwata. "Cell Broadband Engine Architecture and Its First Implementation." IBM. 29 Nov. 2005. 14 Nov. 2007 <http://www-128.ibm.com/developerworks/power/library/pa-cellperf/>.

159 Page, M. "IBM's CELL Processor: Preview to Greatness?" PCStats. 15 May 2005. 14 Nov. 2007 <http://www.pcstats.com/articleview.cfm?articleid=1727&page=3>.

160 Mercury Systems. Telephone interview. 15 Dec. 2007.

161 "Sony Global - Press Release - IBM, Sony, Sony Computer Entertainment Inc. and Toshiba Unveil Cell Processor." Sony. 29 Nov. 2004. 14 Nov. 2007 <http://www.sony.net/SonyInfo/News/Press_Archive/200411/04-1129E/>.

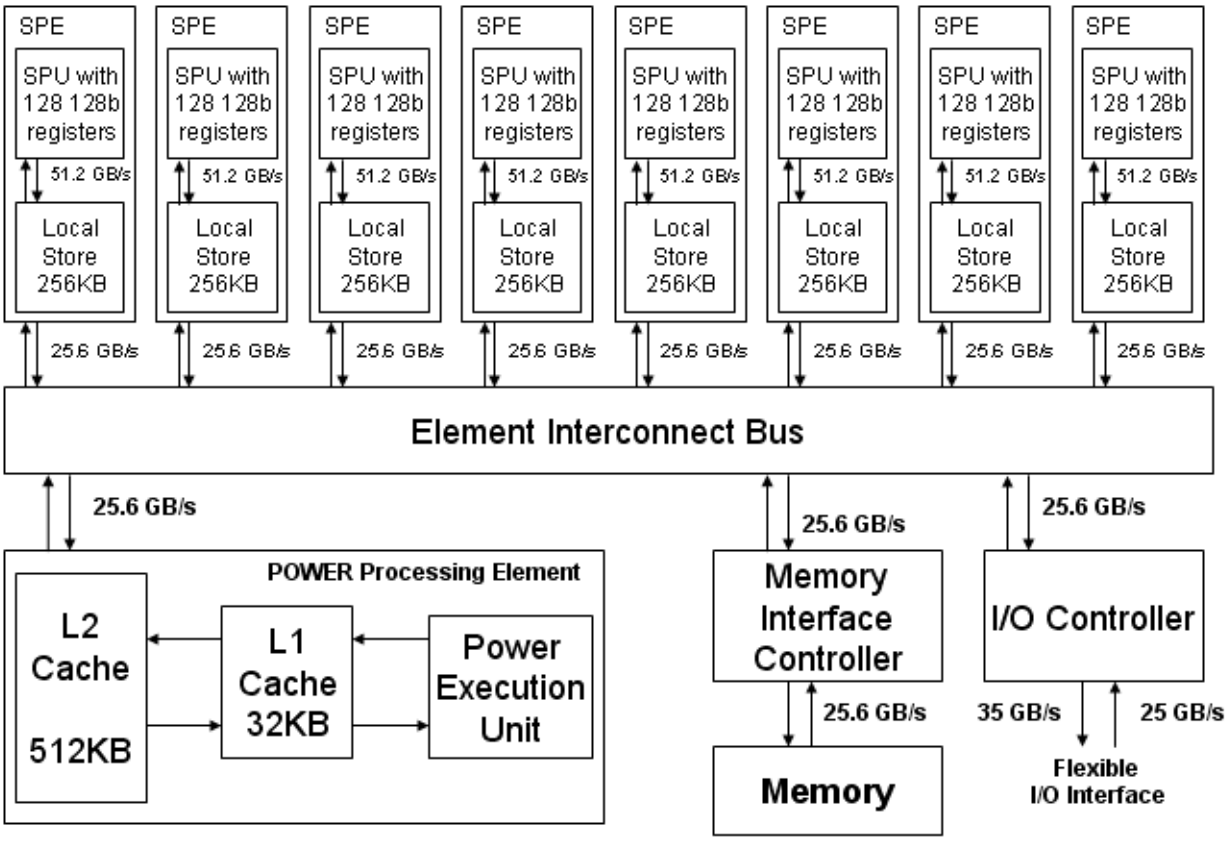

Figure 9 – The overall Cell architecture diagram[162]

The POWER processing element (PPE) is a POWER processing unit (PPU) and a 512KB L2 cache. The purpose of this processor is to run the operating system and to also direct the SPEs. The POWER processing unit also supports dual-threading.[163] The PPU is split up into three distinct sections. There is the instruction unit (IU) that is in charge of retrieving the instructions, figuring out what they mean, separating them, distributing them, and to confirming completion of the instructions. The next piece of the PPU is called the fixed-point execution unit (XU). This portion of the PPU is accountable for the fixed-point instructions. This also includes

[162] Compiled from:
Gschwind, Michael, H P. Hofstee, Brian Flachs, Martin Hopkins, Yukio Watanaba, and Takeshi Yamazaki. "Synergistic Processing in Cell's Multicore Architecture." Penn State. 20 Nov. 2007 <http://citeseer.ist.psu.edu/cache/papers/cs2/624/http:zSzzSzwww.research.ibm.comzSzpeoplezSzmzSzmikegzSzpaperszSz2006_ieeemicro.pdf/gschwind06synergistic.pdf>.
Easton, John, Ingo Meents, Olaf Stephan, Horst Zisgen, and Sei Kato. "Porting Financial Markets Applications to the Cell Broadband Engine Architecture." IBM. June 2007. 20 Nov. 2007 <http://dl.alphaworks.ibm.com/technologies/cellsw/cellFMwhitepaper.pdf>.

[163] Chen, Thomas, Ram Raghavan, Jason Dale, and Eiji Iwata. "Cell Broadband Engine Architecture and Its First Implementation."

any loading or storing instructions. The last component is referred to as a vector scalar unit (VSU). The vector scalar unit will pick up the rest of the work. This specifically is dealing with vector and floating-point instructions.[164]

The synergistic processing elements (SPEs) are the bread and butter of the Cell. The SPEs are split into two distinct pieces. There is the synergistic processing unit (SPU) and the memory flow controller (MFC). The SPU contains a 256KB of storage that IBM likes to call local store. The SPU also has a vaunted 128 by 128 bit set of registers. That means there are 128 registers that can store up to 128 bits each. The SPUs are based around being able to do single instruction, multiple data calculations. The SPEs have been optimized to do calculations in an extremely efficient manner. The Cell can then be expanded in the future based on the need for processing and how limited space is. This is done by adding more SPEs. However, if a task does not require a lot of computer power then SPEs could be taken away from the Cell.[165]

A special part of the SPE is the local store. These local stores were designed to reduce hardware complexity while offering a similar benefit as cache. Data will be loaded from the local stores to the registers of the SPEs. The local store is designed to be able to transfer data to the register at a rate of 16 bytes per cycle. The significance of the 16 byes per cycle is that it is enough to fill up a complete register (128 bits). The problem that arises, because of this, is that there is a sharing of transmission resources. That is to say that there has to be a write process to and from memory while also moving the data to and from the registers. To minimize this problem, the local memory is accessed for 1024 bits at a time during one cycle. By doing this,

[164] Kahle, J A., M N. Day, H P. Hofstee, C R. Johns, T R. Maeurer, and D Shippy. "Introduction to the Cell Multiprocessor." IBM. 7 Sept. 2005. 16 Nov. 2007 <http://www.research.ibm.com/journal/rd/494/kahle.html>.
[165] Chen, Thomas, Ram Raghavan, Jason Dale, and Eiji Iwata. "Cell Broadband Engine Architecture and Its First Implementation."

the data is moved quickly but also interferes less with the SPE's overall operation.[166] The SPEs are specifically designed to compute things in an SIMD manner.[167]

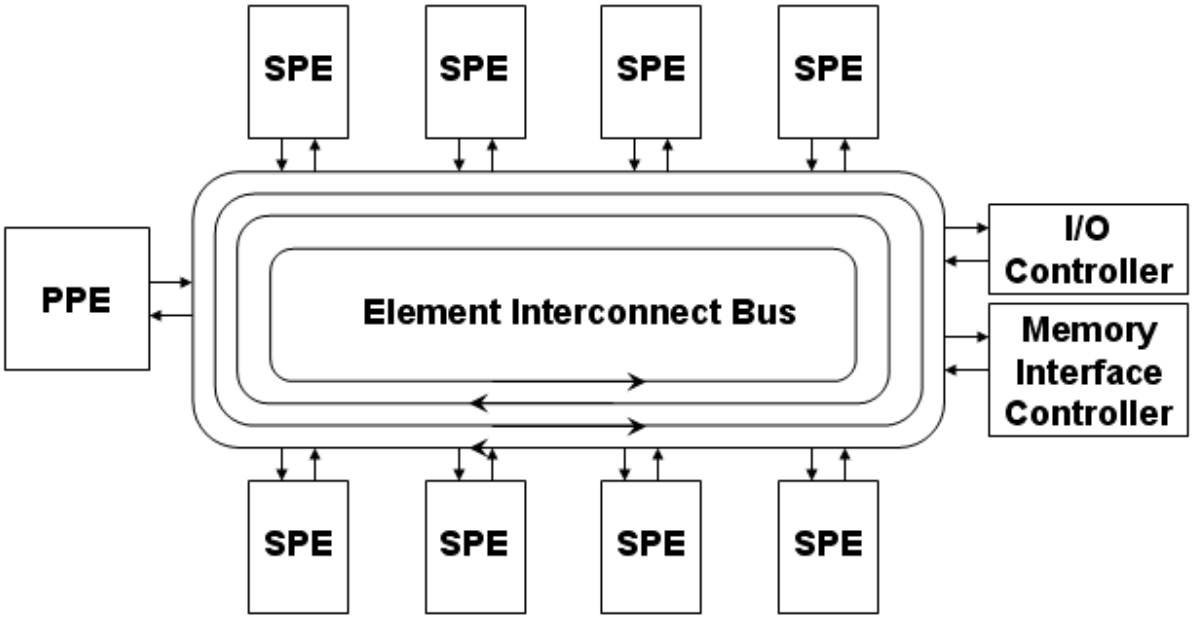

Figure 10 – Element Interconnect Bus Overview[168]

The element interconnect bus (EIB) is what allows all of the sections to talk to one another. It connects the PPE, the SPEs, the off-chip memory, and the external inputs and outputs. This bus is composed of four rings. Two of these rings run clockwise and the other two run counter clockwise. They allow simultaneous data transfers as long as the paths do not overlap one another. Controlling this bus is an arbiter that listens to the requests that are being made of the bus. It then decides which and when the requests will be fulfilled. The theoretical bandwidth of the bus is 204.8 GB/s. However, in practice it depends on where information is traveling to and from on the Cell processor. The theoretical bandwidth is achieved by summing up the

166 Blachford, Nicholas. "Cell Archiecture Explained Version 2 - Part 1: Inside the Cell." Blachford. 2005. 21 Nov. 2007 <http://www.blachford.info/computer/Cell/Cell1_v2.html>.
167 Stokes, Jon. "Introducing the IBM/Sony/Toshiba Cell Processor - Part I: the SIMD Processing Units." Ars Technica. 7 Feb. 2005. 14 Nov. 2007 <http://arstechnica.com/articles/paedia/cpu/cell-1.ars>.
168 Keable, Crispin. "17th Machine Evaluation Workshop." STFC Computational Science and Engineering Department. 6 Dec. 2006. 21 Nov. 2007 <http://www.cse.scitech.ac.uk/disco/mew17/talks/Keable_IBM_MEW17.pdf>.

individual 25.6 GB/s SPEs.[169] The theoretical bandwidth also includes the knowledge that the EIB runs at half the clock speed.[170] Being able to understand how the transfers are occurring will allow a programmer to manipulate the transfers in order to get the most out of the Cell processor.

The Cell has a lot of potential. Just reviewing the plans for future Cell iterations is enough to excite the mind with all the possibilities. During 2008, another iteration of the Cell will be released that has been enhanced for double precision calculations. Currently the Cell has a phenomenal ability to perform approximately 210 GFLOPs for single precision and approximately 14.6 GFLOPs for double precision. However, when the new Cell is released that has been optimized for double precision; the peak performance will jump to 100 GFLOPs for double floating point operations.[171]

	Cell (8 SPE)	X1(MSP)	X1E(MSP)	Opteron	Xeon	Xeon 5365 Quad Core
Frequency	3.2 GHz	800 MHz	1.13 GHz	2.2 GHz	2.3 GHz	3 GHz
GFLOPs (single)	204.8	25.6	36	8.8	4.6	24
GFLOPs (double)	14.6	12.8	18	4.4	4.6	24
	Pentium 4 3.0C	Pentium 4	Itanium2	Itanium2	Itanium2	Itanium2 Dual Core
Frequency	3 GHz	3.2 GHz	900 MHz	1.4 MHz	1.4 GHz	1.6 GHz
GFLOPs (Single)	5.8	6.4	3.6	5.6	6.4	12.8
GFLOPS (Double)	5.8	6.4	3.6	5.6	6.4	12.8

Table 10 – Comparison of an 8 SPE Cell versus common processors[172]

[169] Chen, Thomas, Ram Raghavan, Jason Dale, and Eiji Iwata. "Cell Broadband Engine Architecture and Its First Implementation."

[170] "Cell Architecture." University of North Texas. 21 Nov. 2007 <http://csrl.unt.edu/~kavi/CSCE6610/cell.ppt>.

[171] Mercury Systems. Telephone interview. 15 Dec. 2007.

[172] Compiled from:
Williams, Samuel, John Shalf, Leonid Oliker, Parry Husbands, Shoaib Kamil, and Katherine Yelick. "The Potential of the Cell Processor for Scientific Computing." Lawrence Berkley National Laboratory. 2005. University of California. 22 Nov. 2007 <http://repositories.cdlib.org/cgi/viewcontent.cgi?article=4262&context=lbnl>.
Williams, Samuel, John Shalf, Leonid Oliker, Shoaib Kamil, Parry Husbands, and Katherine Yelick. "Scientific Computing Kernels on the Cell Processor." 2007. Berkeley. 22 Nov. 2007 <http://bebop.cs.berkeley.edu/pubs/williams2007-scicomp-cell.pdf>.
Sun, C. "Intel Pentium 4 3.0 GHz 800MHZ FSB Processor Review." PCStats. 11 June 2003. 22 Nov. 2007 <http://www.pcstats.com/articleview.cfm?articleid=1403&page=4>.
Chen, Thomas, Ram Raghavan, Jason Dale, and Eiji Iwata. "Cell Broadband Engine Architecture and Its First Implementation." IBM.
"Pentium 4 Xeon Cluster 2.2/2.4/2.8 GHz." Top500 Supercomputing Sites. 22 Nov. 2007 <http://top500.org/system/6159>.
"SGI Altix ICE 8200, Xeon Quad Core 3.0 GHz." Top500 Supercomputing Sites. 22 Nov. 2007 <http://www.top500.org/system/8966>.
Kindratenko, Volodymyr. "High-Performance Computing on FPGAs." National Center for Computational Sciences. 7 Sept. 2007. Univesity of Illinois At Urbana-Champaign. 23 Nov. 2007 <http://www.ccs.ornl.gov/workshops/FallCreek07/presentations/kindratenko.pdf>.
"Sony/Toshiba/IBM (STI) CELL Processor." Spring 2007. University of Tennessee. 23 Nov. 2007 <http://www.cs.utk.edu/~dongarra/WEB-PAGES/SPRING-2007/Lect04b.pdf>.

Table 10 shows a comparison of GFLOPs using the LINPACK benchmark. The first element that jumps out is that the GFLOPs of the Cell with 8 SPEs is rated at 14.6 GFLOPs for double precision. The single precision capability of the Cell is 14 times greater than it's double precision computing capability. It is also important to note that Table 10 represents the peak GFLOPs of the processor not the average. The Cell that is quoted above is strictly the combined power of the SPEs and does not include the PPE. The PPE is capable of performing 6.4 GFLOPs of double precision calculations and 25.6 GFLOPs of single precision computations.[173] The release of the next iteration of Cell will increase the double precision GFLOPs raw capabilities to 100 GFLOPs. Also, with the advent of more precise manufacturing processes, more SPEs will be added to the Cell processors.[174]

A way of getting around the slower double precision capabilities of the Cell currently is by using a mixture of single precision and double precision calculations. This is done by using parallel matrix multiplication techniques. Approaching a problem like this has been able to yield up to 100 GFLOPs. However, this complicates the programming process by constantly deconstructing double precision calculations into single precision calculations. Then the single precision calculations have to be promoted back to double precision calculations through a comparison process with LU decomposition.[175] IBM's interest in including an upgraded double precision element to the Cell design stems not only from outperforming the competition but also simplifying the amount of work a programmer has to do in order to optimize their code for the Cell. Also, having this raw statistic increase will look favorably on its marketing power by showing an overwhelming computational power when compared to the competition. Even the

[173] "Sony/Toshiba/IBM (STI) CELL Processor." Spring 2007. University of Tennessee. 23 Nov. 2007

[174] Mercury Systems. Telephone interview. 15 Dec. 2007.

[175] Kurzak, Jakub. "New Approaches to Numerical Linear Algebra on the CELL Processor." Innovative Computing Laboratory. 25 Oct. 2006. University of Tennessee. 23 Nov. 2007 <http://www.cs.utk.edu/~dongarra/cell2006/cell-slides/03-Jakub-Kurzak.pdf>.

quad-core Xeon X5365 will be approximately four times slower when the new Cell is released in early 2008.

The Cell drives the Playstation 3™ and is used mainly for computing physics. However, programmers are now releasing their own code for the Cell based on the Playstation 3™ hardware. One of these is programs is the Folding@Home application. Folding@Home is a distributed computing program that pools resources of inactive Playstation 3™s.[176] IBM and Mercury systems are also pursuing the financial applications of the Cell processor. Both have written algorithms that perform Monte Carlo simulations, Mersenne Twister random number generation, and pricing using the Black-Scholes model.[177]

[176] "Folding@Home on the PS3." Standford University. 22 Nov. 2007 <http://folding.stanford.edu/English/FAQ-PS3>.

[177] Compiled from:
Bicer, Murat, Jonas Larsson, and Matt Sexton. "European Option Pricing with Monte Carlo Simulation on the Cell BE Processor." Mercury Computer Systems - White Paper.
Easton, John, Ingo Meents, Olaf Stephan, Horst Zisgen, and Sei Kato. "Porting Financial Markets Applications to the Cell Broadband Engine Architecture." IBM.

2.3: Econophysics Background

The study of thermal energy and how it can be used throughout research in econophysics continues to grow. Researching privately thermal physics and economics, it has been brought to my attention that there are various ways the American population have been divided and 'characterized' by wealth and income standards using thermal terminology. In the past decade the quote, "The RICH are getting richer, and the POOR are getting poorer" has been held over our world. Using the research of Christian Silva and Victor M. Yakovenko, Professors at the University of Maryland, combined with various other econophysicists' papers, I have created an example of how basic thermal physics studies are currently and in the future should be used to study the world economy using interdisciplinary studies between both physics and economics.

In greater detail, econophysics applies theories from uncertainty/stochastic elements, and nonlinear dynamics developed by physicists to solve problems in economics. These theories are taken to denote the principles of applying statistical mechanics to economic analysis. One facet of econophysics that can be taken into account, is the research based around mechanical and thermal equilibriums. For instance, the study of thermal and super-thermal characterizations that has been researched, the Boltzmann-Gibbs law has been used for solving economical mechanics directed towards income and wealth, and the application of the Pareto power law towards the wealthier portion of our world.[178] A few interesting topics not discussed in our paper are the study of Brownian motion for stock-market prices, stock predictions using Gaussian probability for the distribution of stock-price fluctuations, and the Carnot Process of Economic Growth and Wealth Distribution. These methods are often taken from statistical physics. Statistical physics used in econophysics is known as the application of economical, statistical, and chaotic models

[178] Philip Ball (2006). "Econophysics: Culture Crash". Nature 441: 686-688. Retrieved on 2007-08-27.

used throughout financial physics around the world to express economical data that has been collected. [179]

To a physicist, the most interesting thing about economics is the domination by fluctuations in quantities of economic interest. There is a strong belief that we can relate the study of econophysics to the 'butterfly effect'. This effect will show how one single small adjustment in our world economy plays a huge, and at most times, unpredictable part in our society. This single decision impacts everyone on this Earth. Due to big economic shocks that have affected the economy of our world, the possibility of an economic "meltdown" is one that we must be taken seriously. Big changes in big money affect not only people with large amounts of wealth, but also those who have very little of it—those on the boundaries of today's society. [180]

Finding ideas that solve worldwide economic problems can potentially help in making progress on unsolved physics problems also. A good example that has been looked at is 'turbulence'. If we take a bucket of water and disturb a surface, energy is now added to the system on a big scale. This energy then dissipates over a progressively smaller scale. This is an unsolved physics problem that is currently being looked at by university physicists. Many empirical facts have been made about this experiment, but little can be said about understanding the entirety of this theory. The world economy is comparable to this example of turbulence. One can add information on a big scale to an economic system; hence news of who wins a presidential elections, future technology advances, and any information on smaller and smaller

[179] Jean-Philippe Bouchaud, Marc Potters (2003). *Theory of Financial Risk and Derivative Pricing*. Cambridge University Press.
[180] Jean-Philippe Bouchaud, Marc Potters (2003). *Theory of Financial Risk and Derivative Pricing*. Cambridge University Press.

scales. The way that you handle this 'turbulence' and how it affects social equilibrium is to approach this turbulence problem with econophysics research.

2.3.1: Economic Phenomena

Econophysics is based concurrently around the idea of economic phenomena. For example, it represents how many particles interact between each other, taking into account the properties of human beings. Obviously human beings and particles are significantly different from each other, but even though humans and particles are different they both yield systematic interactions and chaotic behaviors. In reality, it's like a group of water molecules. The world economy is a complex system of individual members (usually countries) that interact with each other. World-wide financial data is provided and supplies opportunities for extensive statistical analysis of our present-day society. This world-wide analysis can lead to a better understanding of the behavior of economical complex systems. [181]

In early 1996 during the study of business firms, physicists and economists found the probabilities associated with the growth rate for a firm. It depicted a description using a single mathematical function for firms of all types and sizes (from sales of $100,000 to $1 trillion). They found that the width of the curve showing the probability distribution follows a "Pareto power law", in which the width is proportional to the firm size raised to a power of approximately 1/6. Today, a Boston University-MIT physics team led by Youngki Lee collaborating with a Harvard economist team led by David Canning has found the same universal patterns and power laws for the fluctuation in growth rates of the gross domestic products (GDP)

[181] http://www.nytimes.com/2005/12/11/magazine/11ideas1-15.html , Econophysics, Shea, Christopher, Published Dec, 11 2005. NY Times 11/20/07 2:28pm

of 152 countries since the 1950's! There are many other fields and tools from physics that have been used in econophysics with success.[182]

2.3.2: Basics of Income, Money, and Wealth

We'll start out with the basics. The theory of statistical mechanics from econophysics can be used to represent basics of money, income and wealth worldwide. Statistical mechanics is in correlation with statistical physics to describe the probability distributions of money, income, and wealth. The equilibrium of statistical mechanics is based on the Boltzmann-Gibbs law. Boltzmann-Gibbs states the probability distribution function (PDF) of energy:

$$\boldsymbol{E\ is\ P(E) = Ce - \frac{E}{T}}$$

where 'T' is the temperature and 'C' is a normalizing constant. The main ingredient in the derivation of the Boltzmann-Gibbs law is the conservation of energy. When two economic agents (such as companies) make a transaction, some amount of money is transferred from one agent to another, but the sum of their money before and after the transactions is the same:

$$\boldsymbol{m_1 + m_2 = m_1' + m_2'}$$

The equilibrium probability distribution function of money in a closed system of agents has the Boltzmann-Gibbs form:

$$\boldsymbol{P(m) = Ce - \frac{m}{T}}$$

where 'T' is the effective "money temperature" equal to the average amount of money per agent. This exponential distribution is observed in computer simulations, as shown in figure 11. In addition, a lot of statistical data is available for the distribution of income, 'r' for revenue. Using Boltzmann-Gibbs, figure 12 shows the probability distribution function of the individual income

[182] http://www.aip.org/pnu/1998/split/pnu395-1.htm , Number 395 Phillip F. Shcewe and Ben Stein. Oct 7, 1998 American Institute of Physics. 11/20/07 2:32pm

in USA and how it's well fitted by the exponential function P(r) = Ce−r/T. The standard plot of probability distribution function inevitably puts an upper limit on the horizontal axis at $120,000/year in figure 12. [183]

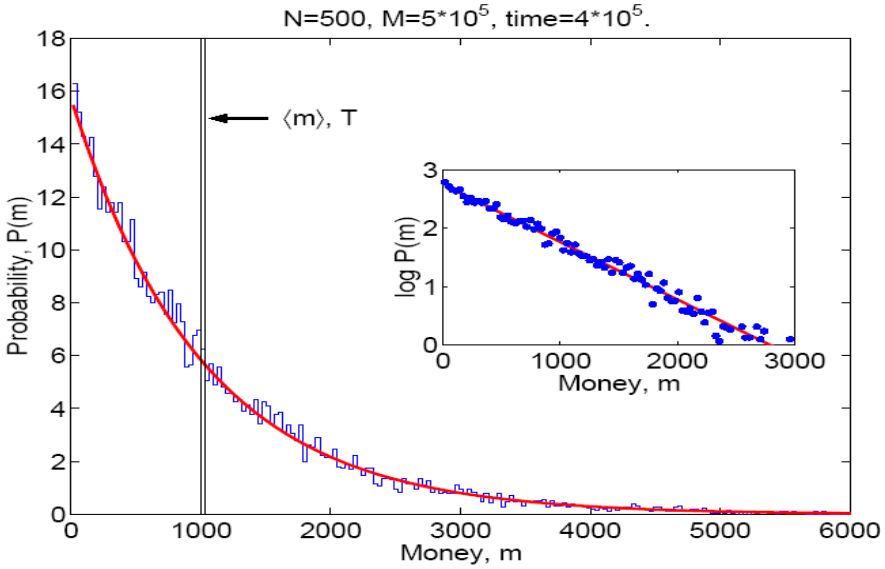

Figure 11- Probability Distribution of Money in Computer Simulation
A. A. Agulescu and V. M. Yakovenko, Euro. Phys. J. B
17, 723 (2000) [cond-mat/0001432].

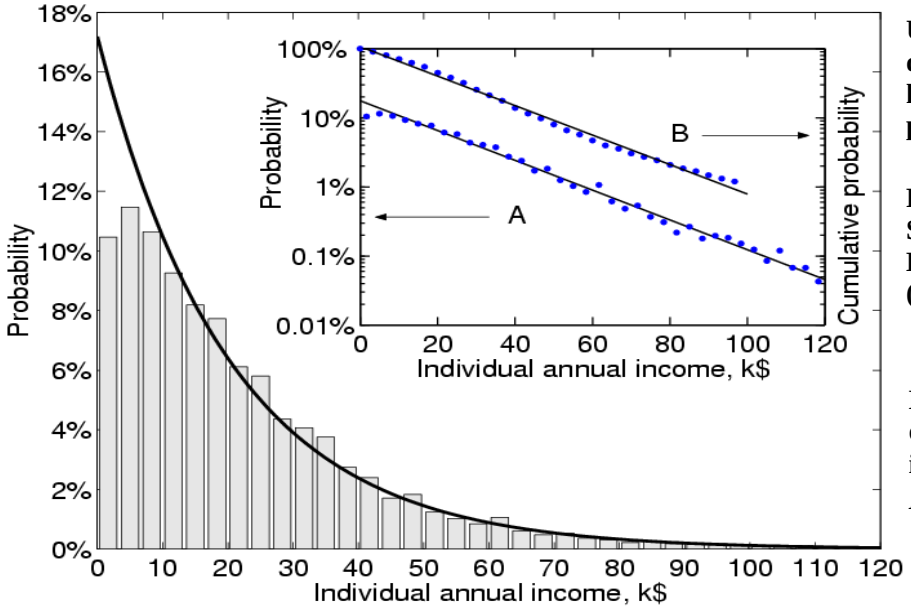

US Census data 1996 – histogram and points A

PSID: Panel Study of Income Dynamics, 1992- (U.Michigan) points B

Distribution of income *r* is exponential: $P(r) \propto e^{-r/T}$

Figure 12 - Probability Distribution of Individual income in the USA
A. A. Agulescu and V. M. Yakovenko, Euro. Phys. J. B 20, 585 (2001) [cond-mat/0008305].

[183] Dr˘agulescu, A A., and V M. Yakovenko. European Physics. J. B, 2000.

2.3.3: Lorenz Curve and Gini Coefficients

A standard way of representing the 'whole income' is by using the Lorenz curve shown in figure 13. The Lorenz curve is a graphical representation of the cumulative distribution function (CDF) of a probability distribution; it is a graph showing the proportion of the distribution assumed by the bottom percent of the y-values. The horizontal axis of the Lorenz curve, x(r), represents the fraction of population with incomes below r, and the vertical axis y(r) represents the fraction of the total income the population accounts for. As r changes from 0 to 1, x(r) and y(r) change from 0 to 1, it defines the Lorenz curve in the (x, y) space. The diagonal line y = x represents the Lorenz curve in the case where all population has 'equal' income. The inequality of the actual income distribution is characterized by the Gini coefficient.

$$0 \leq G \leq 1$$

The Gini coefficient is the area between the diagonal and the Lorenz curve, normalized to the area of the triangle beneath the diagonal.

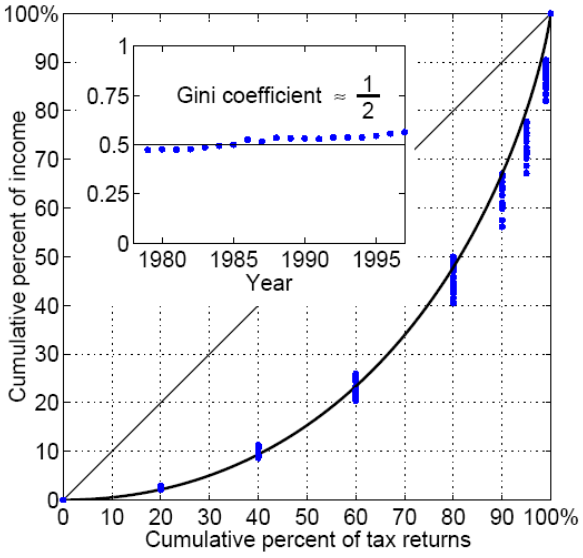

.

Figure 13 - Gini Coefficients and Lorenz Curve
Silva, Christian A., and Victor M. Yakovenko.
"Temporal Evolution of the "Thermal" and "Superthermal"."

For the exponential probability distribution function, the Lorenz curve and the Gini coefficient can be easily calculated:

$$\boldsymbol{y = x + (1 - x)ln(1 - x), G = \frac{1}{2}} \quad \text{Eq. (1)}$$

(1) The solid line in figure 13 shows the theoretical Lorenz curve given by Eq. (1), and the points shown in the income data for 1979–1997. The agreement is quite good. In the first approximation, given the curve of Eq. (1) has no fitting parameters, the inset graph shows that the Gini coefficient is close to the theoretical value 1/2, although the inequality does increase during the last 20 years. Using the Lorentz Curve and the Gini coefficient now for family incomes we can see that inside of the US the overall entropy, which represents the balance of wealth throughout lower-income families in America, has been quite stable over the past 50 years. With no drastic changes the entropy will be proven by G=1/2 for individuals and G=3/8 for families in this income level. [184]

[184] Yakovenko, Victor M. "Research in Econophysics."

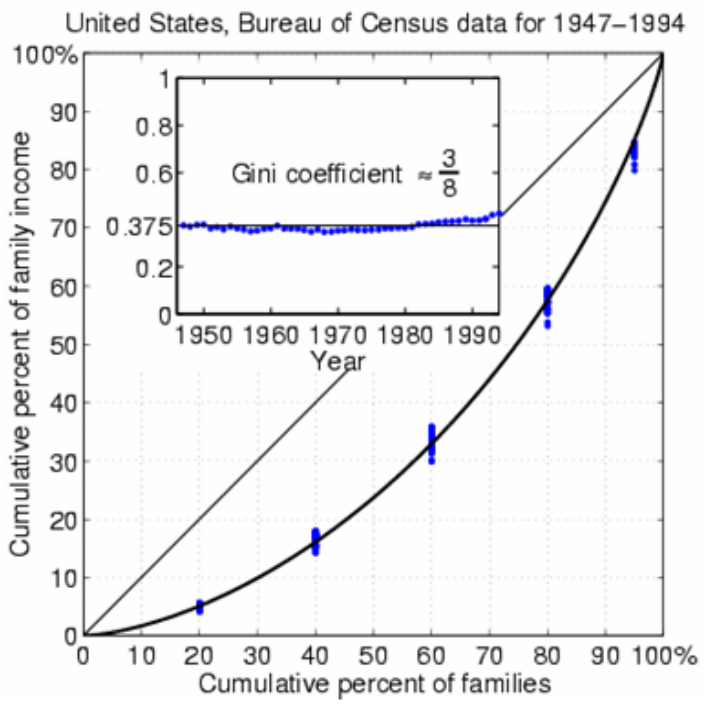

Lorenz curve for families $P_2(r) \propto r \exp(-r/T)$. The calculated Gini coefficient for families is $G=3/8=37.5\%$

No significant changes in Gini and Lorenz for the last 50 years. The exponential ("thermal") Boltzmann-Gibbs distribution is very stable, since it maximizes entropy.

Maximum entropy (the 2nd law of thermodynamics) $\Rightarrow$ **equilibrium inequality**:

$G=1/2$ for individuals,

$G=3/8$ for families.

Figure 14 – Lorenz Curve and Gini Coefficient for Families
Silva, Christian A., and Victor M. Yakovenko.
"Temporal Evolution of the "Thermal" and "Superthermal"."
Univ. of Maryland. 15 Jan. 2005. UMD. 30 Nov. 2007

3: Methodology

3.1: General Approach

The generalized approach to this problem followed six key steps. Each of these steps was used to bring in a proper flow of information that we would be able to utilize to our benefit. The six steps were:

- Collect colleague information
- Perform system level analysis
- Write a brief description of the systems involved
- Look at all the processes
- Develop a strategy
- Create a proposition for improvement.

These steps were used throughout the project to provide fundamental background and to further elaborate on the actual problems. The collect colleague information step was used to set up valuable lines of communication. These lines of communication ultimately allowed us to have an unrestricted access to the people who would benefit from our work. There were instances where it was not possible to collect information from our colleagues. In these cases, we tried other sources of information such as the Bank of America information database.

The systems level analysis was put in as a way of understanding what was going on. Without a solid foundation of what systems were talking to one another, it would seem foolish to try to solve the problem. This step was specifically designed to allow us to have a physical understanding of the data flow and the communications that take place. This was also a precursor to the description step as each system had to be identified, in order to determine it's relevance to the entire system.

The writing of a brief description of the systems involved was designed specifically to understand at a technological level what was going on. This did not promote insight on the

physical processes but it did allow us to understand what type of information was being passed in between the machines. This also allowed us to know what type of information was being inputted into the machine. We used this step extensively as a way of familiarizing ourselves with the P&L reporting systems. It was also a good way of testing out our contacts and pull more information from them that wasn't readily available.

The look at the processes section was specifically an examination at how the information was being used. This step also hit a wall because we were being isolated from middle office. This prevented us from seeing how they truly reconciled the Actual P&L. However, we were still able to go around with the amount of information we knew. We were able to look at what went into middle office and ultimately what came out of middle office. This helped us by preparing our questions in such a way that we would be able to get the maximum amount of information out of a contact.

The strategy development step was designed to not only prepare for meetings but to also figure out what information would be necessary to solve the problems presented to us. A lot of the strategy based work was around figuring out who of our contacts would be the correct person to ask a specific question. It was also figuring out how to pull a solution out of someone in the most efficient manner as we were on borrowed time. This approach produced a better feel for what was actually needed by Bank of America.

The proposition for improvement was truly the last real step in the process. This is where all of the ideas come together and lead to a finalized product. This finalized product is the summation of the solutions to the problems as we interpreted them. This is the true fruit of the project as the rest is just planting the seed, watering it, and watching the tree grow. These

proposals for improvement add up to what was asked for us by Abhishek Vashishth on behalf of the technology department in Bank of America - London.

3.2: Systems Engineering

When undergoing a task of this magnitude, there are many ways to start. We decided to do an overall systems analysis. This allowed us to be able to understand the reporting systems that were in use by Bank of America. This basic knowledge enabled us to be able to converse on a similar level as the people we were working with. The systems that are involved in the Flash and Actual P&L reporting process are detailed in this following section.

The Flash P&L Workflow components includes Odin, Camden, Results Database Application (RDA), eBlotter, Alice, Risk Engine (Risk Engine), Results Database, Flash Controller/Calculator, Flash P&L Database, Flash Viewer, Bloomberg AF8, Cash MO Templates, GCP Flash Controller, and GCP Flash Summaries. Camden supplies market data, bond prices, CDS spreads, volumes, and correlations to the Flash P&L. The eBlotter is a program that supplies lists of trade entries and confirmations to the Flash P&L. Alice is used for bond references that are entered into the Flash P&L Database. Flash Viewer is used to view the Flash P&L Database in an Excel format. The Flash P&L database sends the data to the GCP Flash Controller for Flash Consolidation.

For further depth to what goes into the P&L here are a few details. Credit Advantage Market Data Environment (Camden) is the repository for non-interest rate market data used in the pricing of credit derivatives. Camden is an “open service” that is able to provide credit market data to any application requiring it within a bank. It contains functionality for viewing credit spread history, bonds, and bond price history. It also allows distribution of market data,

bond prices, CDS Spreads, Volumes, and asset correlations. This data is distributed to Risk Engine, Flash P&L, and Front Office Publishing. Currently Camden is a stand-alone application; in a later release it will be available as a component on the Advantage desktop. Camden consists of a central server, regionally deployed web services, and a locally installed user interface which holds an editor for maintaining asset correlation data and abstract data. Once Camden is done collecting data, the data is sent to both the Risk Engine and Flash Controller for further analysis.[185]

The RDA is used by the Middle Office as a repository for historical risk and P&L data that is generated by the Risk Engine. The overall objective of this system is to maintain a controlled environment for risk and P&L reporting. Risk and P&L numbers are generated by the Risk Engine during an overnight batch, which is started after the close of business in the U.SA. The Risk Engine is a computation farm which uses trade and static data from other systems, such as Alice, to calculate both Risk and P&L numbers. These numbers are then fed into a central database and are able to be viewed and manipulated by the users. Because the numbers from the Risk Engine are generated overnight, any intra-day amendments to the system are not affected immediately; therefore, a mechanism is required to perform this function, hence the RDA. [186]

Alice is also another component mentioned that passes on bond references to the Flash P&L Database. Alice is built on a Sybase Database at its core. This database maintains all of the information collected by Alice. A collection of Java Enterprise services is used to facilitate the management of changes to this data. Currently, separate services exist for bonds, loans and

[185] Ogrinz, Michael, and Kathiresan Solaiyappan. "Camden Home." Bank of America.Com. 8 Nov. 2007. Bank of America. 22 Nov. 2007 <http://gmwiki.bankofamerica.com/display/camden/Camden+Home>.

[186] Reporting Group, Risk Engine. "Risk Engine User Guide." Risk Engine. 3 Oct. 2007. Bank of America. 16 Nov. 2007 <http://crpscsgap105/RiskEngine/servlet/FileRetriever?filename=%5C%5CCRPNYCDEV22%5Cp4store%5Cvob_adv_riskengine_docs%5CUser+Guides%5CRisk+Engine-User+Guide.pdf>.

issuers. These services are independent at this point, although bonds and loans do depend on the issuer's information. Future plans to provide a similar import mechanism from Bloomberg are currently being researched. The user application that manages the loans is delivered by using small web-based applications. These small web-based applications are used to manage the definitions of the loan tranches and credit agreements.[187]

The Risk Engine was developed by Bank of America's Risk Engine and Reporting group to provide a single tool to compute risk, P&L Explain, and General Ledger (GL Jobs) based on market scenario capabilities. The results computed by the Risk Engine are saved to the Risk Engine Results Database and used by many other Bank of America financial applications. The Risk Engine can be configured to run both standard and calculations (called "jobs") that produce risk figures, P&L Explains, and GL Jobs. Each job uses a set of trades, and a set of pricing rules. At the point of calculation, each job activates the database's refresh rules to standardize the retrieval of data from the data sources. The calculation results are then saved to the Risk Engine Results Database as part of a results group called a "data set". [188]

Within the Risk Engine process there are jobs. A job is information submitted intraday for processing. Users can view the current status of each job submitted for intraday processing. If there are several jobs that can use the same pricing configuration, data set, and caching policy, you can group these similar jobs into a single entity for processing called a "batch". Standard batches or jobs are scheduled to run overnight, but prototype batches can be submitted intraday. For batches related to calculating risk, you can specify adjustments to the actual trade called

---

[187] Fisher, Susie. "Alice 4.0." Gmwiki. 2 Aug. 2007. Bank of America. 25 Nov. 2007 <http://gmwiki.bankofamerica.com/display/GSPBAT/Alice+4.0>.

[188] Reporting Group, Risk Engine. "Risk Engine User Guide." Risk Engine. 3 Oct. 2007. Bank of America. 16 Nov. 2007 <http://crpscsgap105/RiskEngine/servlet/FileRetriever?filename=%5C%5CCRPNYCDEV22%5Cp4store%5Cvob_adv_riskengine_docs%5CUser+Guides%5CRisk+Engine-User+Guide.pdf>.

"scenarios". Scenarios are used in the calculations process to project the impact of a particular trading market data event. [189]

Scenarios are primarily used for analytical use prior to determining job submissions. Risk Engine tracks all user changes that have been saved to the batches, jobs, data sets, pricing configurations, and caching policies. All users can see the history of each data element that has been changed. These user changes are known as "audits". In addition, Risk Engine provides a tool for monitoring the status of all submitted jobs, batches, and audits. This monitoring process allows for the Risk Engine to summarize any errors that may have occurred throughout the Flash process.

With this knowledge of the systems and our in-depth look at financial terms, we were able to converse with our fellow employees on a level that allowed us to pull the maximum amount of information from them. We used this information to fuel the pursuit of being able to find acceptable solutions based on the problem at hand. This helped us understand the vital roles of components such as Odin and eBlotter. All of this combined allowed us to be able to have a firm grasp of what was to be expected of the overall project.

3.3: Flash P&L

3.3.1: Approach to the Problem

Flash P&L is the Profit and Loss statement that is received in the morning by the traders in the front office. The traders use this statement to compare their calculations that were done the previous night by Odin. The workflow diagram depicts how the P&L process works. The

[189] Reporting Group, Risk Engine. "Risk Engine User Guide." Risk Engine. 3 Oct. 2007. Bank of America. 16 Nov. 2007 <http://crpscsgap105/RiskEngine/servlet/FileRetriever?filename=%5C%5CCRPNYCDEV22%5Cp4store%5Cvob_adv_riskengine_docs%5CUser+Guides%5CRisk+Engine-User+Guide.pdf>.

diagram shows the flow of information from the front office, through the Risk Engine and middle office, and into the Flash Consolidation Process. Briefly the diagram portrays how market data, Bonds, and CDS Spreads are brought in from various repositories into the Risk Engine. Overnight the Risk Engine runs analysis on the risk, P&L, and GL Jobs from the daily trades made by investors and traders. These results are then stored in a results database. Eventually, these results are then transferred to a Flash P&L Controller/Calculator that begins calculating the Flash P&L. This data is then stored in a Flash P&L Database where the information can be viewed in Excel format by users, hence the use of the Flash Viewer.

The Front office is having issues with the accuracy between the Flash P&L and their own personal P&L that was calculated after the inputs of Odin. They were having issues with various things throughout the Flash process, such as partial unwinds, cancelled / removed trades, duplication of new trades, the NewBus P&L tab and overall timeliness and system performance of running the Flash. The current major partial unwinds problem is a current split/confusion between NewBus, Credit Spread, and edit inputs by traders. When the inputs become error prone, the traders' predictions and P&L do not match. This leads to further problems and confusion between Middle office and Front office. Middle office will provide the P&L and request clarification on the differences with the P&L statement. When Front office has trouble recovering the data due to confusion in the inputs, it slows down the front office workflow and errors have to be sorted out. Other current problem varies among different offices. Some believe the calculations for the Flash are too elaborate and need to be simplified, and some believe the calculations are already extremely basic but the process just needs to be sped up and made more efficient.

The Flash is on a desktop and uses up CPU resources making the computer unusable. Our suggested solution is to bottle-neck the program which allows the CPU to be used by other programs. Also, in this process we are trying to simplify the layout more. We are suggesting the combination of the brokerage, carry, and curve shift tab into one single tab to clean-up the layout. Another question at hand is if the Flash P&L is even needed? As briefly spoken about above, we ask ourselves if a simple version of the Actual P&L can be pumped out within 3-4 hours of the Middle office receiving the information from Risk Engine. If so, can it be including in the Odin and/or the non-reconciled Actual versus having an entire separate process for calculating the Flash, which is actually slowing down the CPU for other programs.

An area for improvement that's present, having the Flash only report the three major categories of P&L while the others should be grouped together. The three major categories are credit spread, edit, and new business. Ben Cooper has suggested that all of the others besides these three be grouped together. Priscilla Wong has also noticed that these three do seem to be the major players. Current research, based on the concerns of developing ways to make the Flash viewer more user friendly, allow adjustments to be made manually, and by granting Flash the ability to ignore cancelled/removed trades.

3.3.2: Solutions to the Problem

With potential upgrades directed towards Flash upgrades .Net program written in C# is being developed to make viewing the Flash easier and more user friendly. Priscilla Wong is currently looking for opinions on what should and shouldn't be enclosed inside of the program. This program would make dealing with the Flash more streamlined by offering the user capabilities of customizing what they would like to see.

Net program is also under development as an adjustment tool for fixing problems with the Flash. This comes from a request made by Alex May because there are times when he needs to change things manually and would like to be able to do it without haven't to bother the Flash development team. The .Net program appears to be a good step towards making the overall Flash process better.

Cancelled / removed trades are being addressed in the release of the Flash P&L 3.10. This makes Flash more robust by being able to ignore cancelled and removed trades that would not run in the Risk Engine. The trades coming through with the "Blot:####" ID that appear in the NewBus P&L are also being removed as these will cause errors in the Flash.. Trades that are not uploaded to IRP/ADV should not appear in Flash. Trades that are assigned the wrong CDS curve in Flash will be getting an enhancement where they pull additional fields from eBlotter in order to identify the proper curve.

In the Flash 3.10 P&L has new features that have been introduced to eliminate errors, solve current problems and also increases business workflow. The first major concern that the Flash 3.10 introduces is that of multiple partial unwinds. Flash 3.10 has now a better handle on multiple partial unwinds per day. It has found a way to sum all unwind fees for the day if the last eBlotter input event is a Partial Unwind. These new features decrease the amount of manual adjustments that are required to compensate for the unwind problem. This adjustment was requested by Alex May (middle office) and Jennifer Ornellas.

The 2nd new feature in the Flash 3.10 is the ability to correct CDS curves for new trades. Additional fields have been added to eBlotter. Now the Flash can look up the correct curve so trades can be mapped correctly. The CDS curves will now have more information that will be

used to map the curves in Flash properly. The proper mapping will correct credit P&L in Flash for new trades and will benefit both the middle office and front office of Bank of America. This new feature was requested by Alex May.

The 3rd new feature blocks intraday trade cancels/removes that don't make it to Risk Engine from showing up in Flash calculations. This upgrade will remove cancelled/removed trades from the NewBus-P&L tab and will allow the Flash population to be more consistent with the Actual P&L calculations. To exclude the cancelled/removed trades was a request made by Alex May and the middle office.

The 4th and final new feature that was noted by Priscilla Wong is that of eliminating any eBlotter trades that were not STP-ed from NewBus_P&L tab, from appearing in Flash. This in more detail shows that any eBlotter trades that were not uploaded to IRP/ADV will not be calculated in the Flash P&L. This upgrade feature once again will eliminate more inconsistency between the Flash and Actual P&Ls. This final feature was a dual request by Alex May and Bryan Stonebraker.

The Flash server option immediately takes the strain off of the desktop as it would be run on a more robust server. This allows for information to be quickly pulled from the server and for the jobs to be run in a friendlier environment. More jobs could be run by the Flash server as it would not be bogging down a desktop. This is because with it being server based, the user will feel more confident in requesting more information or reprocessing information. This allows for great flexibility and promotes double checking when necessary.

On a side note, there are two options for the computing apparatus that are explored in the Risk Engine section of this project that can be applied to the Flash's goals of becoming a server

based application. These solutions are the 2 processor 3.16 GHz X5460 Hewlett-Packard DL360 G5 and the Cell-based solution. The time frame of implementation is critical for this improvement is critical. This is because if it is before mid-2008 then the X5460 HP DL360 G5 is a better option. This is based on knowledge of financial calculations being based around double precision and the current Cell processor does not offer what the HP DL360 G5 would. The code can be written in such a way for the Cell processor to take advantage of the single precision performance but this would put a larger workload on the current programmers.

If the server option is to be implemented after mid-2008, this allows for the double precision enhanced hardware to be released. This would push the Cell-based solutions past the quad-core X5460s and offer a significant amount of computational power for a reasonable price. However, as discussed in detail in the Risk Engine section, the uncertainty in developing for the Cell still pushes the quad-core X5460s. Overall, when erring on the side of caution; it is safer to go with the quad-core X5460s until more software is developed by third parties to take full advantage of the Cell.

3.4: Actual P&L

3.4.1: Approach to the Problem

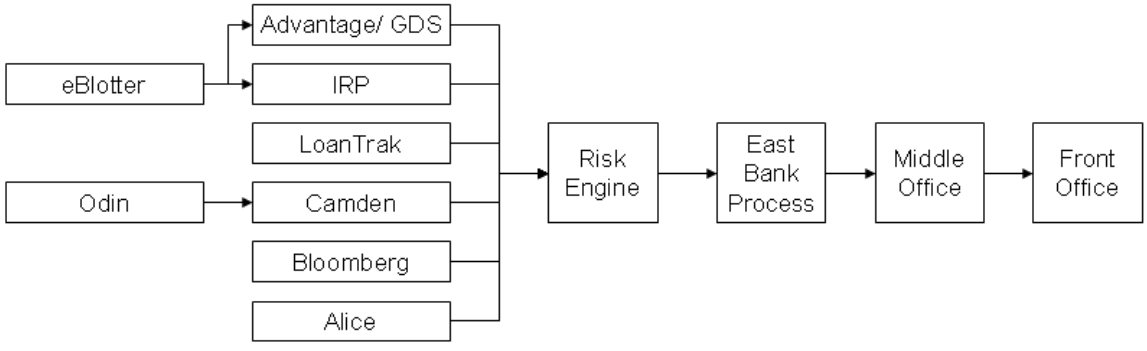

Figure 11 - Interconnected Systems for the Actual P&L

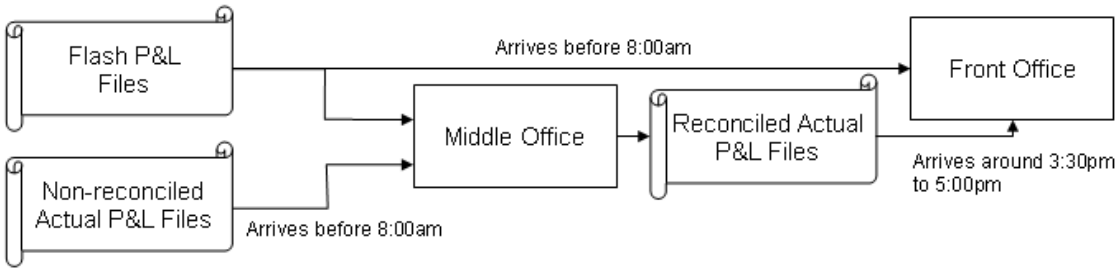

Figure 12 - Flow of P&L Files

The Actual P&L is a report that is generated by the Risk Engine. The Risk Engine is a cluster of 160 computational cores working together to compute three categories of information. The three categories are risk, the Actual P&L, and the general ledger. This system shares its computational time between these three pieces. The Actual P&L is released to the Risk Engine results database. Here it is stored as a text file. This text file is then pulled by the East Bank Process and compiled into an excel spread sheet. This excel spread sheet is published to middle office. This is the non-reconciled Actual P&L. Middle office will then spend most of the day reconciling the Actual P&L. The reconciled Actual P&L will arrive between 3:30pm and 5:00pm British Time. At this time, it becomes little help to the traders.

This immediately became a concern after looking at a timing diagram for the entire P&L reporting process. Middle office has access to the non-reconciled Actual P&L files by 8:00am in the morning usually. This means that it takes a total of 7.5 to 9 hours to reconcile the Actual P&L. We checked with Pete Spencer to see if the Risk Engine could be sped up in order to allow middle office access to the files sooner. This lead us to the fact that the London P&L files are usually delivered to the East Bank Process computers by 5:00am British Time. The East Bank Process had then been identified as a 50 minute process that just put text files together into excel spreadsheets.

This means that if everything works as it should, then the P&L figures are available by 6:00am British Time. However, after discussing this with our colleagues it seemed that 8:00am was a better starting point for access to the P&L. This allowed the people in middle office to show up to work before expecting labor to begin. This led us to looking at the entire picture to see if there were any other minor aspects that we could improve upon.

This allowed us to be ready to start pooling in information from middle office. Unfortunately, we were not able to collect colleague information from the reconciliation team for profit and loss. In lieu of using the information from the reconciliation team, we went to front office to see their main concerns surrounding P&L.

While middle office performs the reconciliation process, front office is the team that will use the results from middle office. This provided us with a black box. Much like in an electrical engineering signals class, we were given the data that went in and the data came out. However, instead of finding the transfer function, we would have to figure out ways of helping middle office without their direct input. This is when we specifically turned to Ben Cooper. We would meet with Ben Cooper weekly to discuss problems that occurred with the Actual P&L reporting process.

One of the first things that was established was when a discrepancy arose during the reconciliation process of the Flash and the Actual P&L, middle office would approach front office to find the solution. This obviously takes the time of both the front and middle office. As we looked farther into the problem, our efforts were ultimately hindered by not being able to see into the middle office's reconciliation process. We had to find other means of getting the information. This came from holding more meetings with Ben Cooper, Abhishek Vashishth, and

Priscilla Wong. Eventually, we were able to construct better questions to pull pieces of information that these three people possessed.

The problems did not stop, as Ben outlined that there is a consistent problem with the Actual P&L. When the Actual P&L is calculated partial unwinds, full unwinds, and assignments will cause reporting errors. A partial unwind is when a portion of a deal is settled but the rest is unsettled. This would be like saying if you have a credit default swap for $1,000,000. You can partially unwind it and settle any amount. For this example, you decide to partially unwind the deal for a total of $500,000. This still leaves $500,000 that hasn't been settled. This unwind process should be completely reported as credit spread P&L. Unfortunately, this is not the case.

What happens with partial unwinds and full unwinds is that some of the P&L is reported in credit spread and some of it is reported in edit P&L. These categories of P&L are very specific and shouldn't share their P&L across separate portions. Edit P&L is specifically reserved for changes that are made because of a human error when inputting the information into the system. This gray area in between the P&L categories had to be clearly addressed. However, assignments are subject to a very similar problem. Assignments should be registered completely as new business. Unfortunately, assignments will cause the Risk Engine to report some of the P&L inside of the new business category while the other portion will go into edit P&L.

In addition to incorrectly reporting P&L by the Risk Engine, one of the tools that was designed to help the traders is being ignored. Many members of the front office believe that the Flash P&L is inaccurate and they've lost faith in the reports generated by the Flash P&L. The Flash P&L's purpose is to give the traders an idea of where they stand before the Actual P&L is

released much later in the day. However, the Flash will often have large discrepancies when compared to Odin.

The traders put their faith into Odin because it allows them to mark their curves and set up their predictions. This puts responsibility on the traders to correctly mark their curves. However, the differences between Odin and Flash became evident even though Odin publishes to Camden. Flash receives its information about credit spread from Camden and ideally Flash should report exactly what Odin reports. This isn't the case as we found out in a discussion with Ben Cooper, Priscilla Wong, and Michael San Jose.

Odin operates independently of eBlotter. This prevents Odin from knowing if there were partial unwinds, unwinds, or assignments. This will ultimately force the traders to mark their curves inherently incorrectly. They are not able to appropriately account for the partial unwinds or assignments and this forces their predictions to be incorrect. This became an area of concern because the traders are relying on their predictions and then the reconciled Actual P&L. This prevents them from having the insight of being able to accurately predict the Actual P&L. Ben Cooper clearly established that in an ideal world Odin, Flash, and the Actual P&L would all be exact matches. This highlighted an area of concern for Odin as the traders basically had a tool that in principle works but doesn't allow them to provide accurate and consistent data.

3.4.2: Findings and Proposed Solutions

3.4.2.1: Odin Upgrades

Upon finding out Odin's shortcomings, it became obvious that Odin needed to pull in more information than it currently does. Currently, Odin only calculates what is called deltas. These, as stated in the financial section, are a measurement the sensitivity to changes in the price of the reference asset. However, the Flash takes this a step further and also calculates the gamma

also. The gamma is the derivative of the delta. The problem with ignoring the gamma is that if there is a large credit spread shift, Odin will not be able to allow the traders to properly mark their predictions.

This inability to properly calculate the credit spread P&L in Odin prevents the traders from making accurate predictions. This ultimately slows down the entire reconciliation process because if Odin correctly reported credit spread P&L, it would be much easier to locate the error in the P&L. It also makes the predictions almost pointless in the fact that they can't properly show the gains or losses.

Another aspect of Odin is that it doesn't take in information from eBlotter. This problem causes Odin to not be able to take into account unwinds, partial unwinds, and assignments. This makes reconciliation harder because the traders are unable to accurately predict their credit spread P&L. The inability to not correctly predict something may sound like a logical and realistic expectation. In the financial world though, traders should be able to make predictions based on knowledge of the market. The marks won't be on target but they will be within an acceptable range. Ben had suggested that around $5,000 was an acceptable prediction range. With the very large sums of money that are being exchanged, this seems reasonable on a fundamental basis also.

3.4.2.2: East Bank Process

The East Bank Process was originally intended to be a short-term program that has lasted over five years. It is now two desktops that sit inside of middle office and each handles an almost equal workload. This can be condensed to a single server based solution. Phil Mitton had requested a system that was robust, timely, secure, and flexible. This focused around having a system that is an SQL server based solution.

This server would have to allow users to randomly throughout the day put together queries and fill out reports. This would have to be done in a user-friendly method that didn't boggle down the users system or the server. This could be done through a program that allows a person to modify what type of report they would like to see and what sort of things on the report they want. This would allow for the flexibility of not having to ask someone else to get a specific report. The most used types of reports would be saved as standard templates. The security aspect would come from the Windows authentication service.

3.4.2.3: Flash and Actual P&L Comparison Program

It is important to realize the traders in front office have lost their faith in the Flash P&L reporting system. This stems from Odin not matching up with Flash and also a misunderstanding of the Flash. The original idea for optimization was to cut middle office out of the loop by automating the reconciliation process through the use of a sophisticated program. However, we soon learned that middle office's purposed was to promote transparency inside the company and serve as a checkpoint. This obstacle presented a unique opportunity and a different way of thinking.

Even though middle office serves as a checkpoint, there had to be a way to get more information to the traders. This spurred the Flash and Actual P&L comparison program. This program would take the Flash and non-reconciled Actual P&L and perform a comparison between the Flash and the non-reconciled Actual P&L. Differences would have highlighted backgrounds in different colors according to their severity. Red would be a difference of over $100,000. Yellow would be a difference between $25,000 and $99,999. Green would represent $2,000 to $24,999. Anything less than $2,000 would not have a color associated with it. This

would let traders immediately see where the discrepancies are and to have more sources for their P&L while they wait for the Actual P&L to be released.

This program would serve other purposes. The first of which, would be a feature that middle office could use extensively. On top of comparing the Flash and the Actual P&L, it would have to be designed to be able to pull the appropriate information to tell where the problem originates from. This is to say, that if there was a problem with a credit spread P&L line that information from Odin would be readily accessible. This would allow the middle office workers to have most of the information that they need in one program. Since this program would automatically search for specific information, it will ultimately cost less time for middle office as they won't be searching as much for where the numbers originated from.

This won't necessarily cut down on middle office going to front office and asking what happened. On the other hand, it will allow middle office to ask more precise questions by having all of the data readily available. Also, with the Odin upgrade, this should cut down the need for middle office to question front office as Odin will allow traders to predict more accurately.

The program could also be designed to be able to compare the reconciled at the end of the day so that traders can see how the Flash and non-reconciled Actual P&L were altered. This would give the traders a better feel of the overall process by quickly showing where the differences and discrepancies occurred. This would also allow traders to have a better understanding of the entire Flash process and promote its use.

The last feature that would be extremely helpful in minimizing the reconciliation process would be a data logging feature. This would log the information of how things were changed. These logs would be able to stem into charts and graphs of where the discrepancies occurred,

how much of a discrepancy occurred, and what was the corrected solution. This data could be collected over a few months time in order to be analyzed for future improvements to the Risk Engine and the Flash algorithms.

3.4.2.4: A New Viewer for P&L

The Flash has a .Net viewer being written in C# that'll allow traders to use the Flash easier. This should be taken a step further and be able to use the files from the comparison program and also the reconciled P&L. This will give traders an immense wealth of information at their finger tips that they can customize to their liking. This would allow traders to put categories together that didn't have large shifts and focus on the main ones (credit spread, edit, and new business).

Since Priscilla Wong's team is already working on this viewer when it is completed it can be ported to the London team to add the additional features. This would save the team development costs and still be able to provide a better product. This could effectively be pushed back to the United States if desired. This would enable traders to be able to quickly pull up information that they need to be able to compare quickly.

3.4.2.5: Process Changes for the P&L Workflow

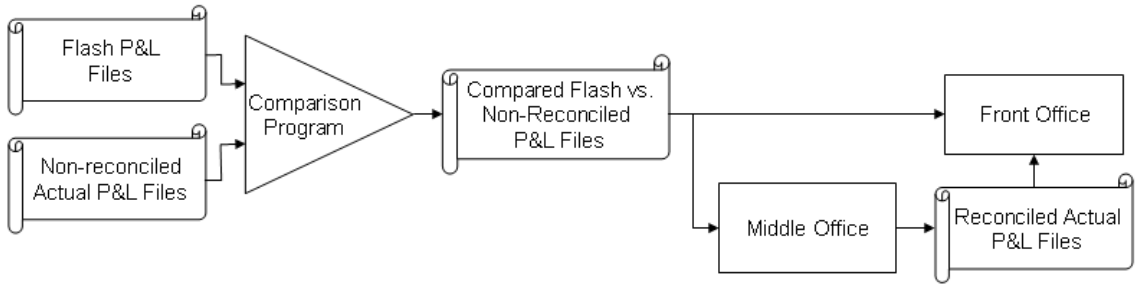

Figure 13 - New Process Flow for the P&L

This process change is directly related to using the comparison program. Instead of front office receiving just the Flash P&L in the morning, they would have access to the Flash and non-reconciled Actual P&L files. This would be provided by the comparison program and would show them the discrepancies. This allows front office to be able to make adjustments accordingly based on how different the Flash and Actual P&L are. This will also make front office aware of the discrepancies. By doing this, front office will be prepared to help out middle office with the reconciliation process in a more timely manner.

Middle office would still go about the reconciliation process as usual but they would also have access to the information from such key systems as Odin and eBlotter readily available. By using the comparison tool in conjunction with their standard practices for reconciling the Actual P&L, they will be able to ask precise questions. This would lead to a speed up by being able to ask well-informed questions instead of asking why the numbers didn't match up. This also promotes transparency in the company by allowing middle office to easily access all of the applicable information.

The concern of middle office being bypassed is also avoided since front office would only be getting the Flash and Actual P&L reconciliation at the beginning of the day. When middle office was done with the reconciliation, they would then release the reconciled Actual P&L to the front office. This keeps middle office as a checkpoint entity for front office and prevents the traders from consistently marking their own trade information incorrectly. Ultimately, this saves the company money by allowing both offices to have an increased flow of information that is vital to their operations.

3.5: Econophysics

3.5.1: Freezing – Bose-Einstein Condensate on Personal Income

A discrepancy can be noticed when looking at the Lorentz Equation and Gini coefficients for short periods of time of individuals on similar income levels. This discrepancy becomes clear when we look at the cumulative distribution of income up to $1,000,000 /year shown in figure 18. It is clear from figure 18 that income distribution for the majority of our population, more than 97%, is described by the exponential Boltzmann-Gibbs law. The extra income in the upper tail of the distribution can be considered as the "Bose-Einstein Condensate".

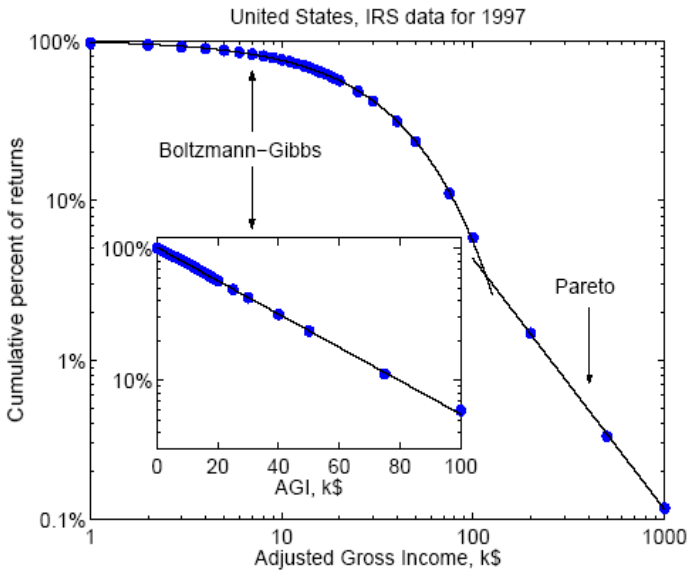

Figure 18 – Cummulative probability distribution of individual income in log-log (main panel) and log-linear (inset)
A.A. Agulescu and V. M. Yakovenko, condmat/0211175.

"Bose–Einstein Condensate" (BEC) in physics terminology is a state of matter formed by a system of bosons confined in an 'external potential' and cooled to temperatures very near to absolute zero (−273.15 °C). Under these super-cooled conditions, a large fraction of the atoms collapse into the lowest quantum state of the external potential. When relating BEC to the

probability distribution function in econophysics, it expresses that when the probability distribution function curve reaches the upper tail, wealthier individuals of America, the distribution observes scenarios that suggest that most wealthy market participants tend to invest similarly to other wealthy participants. This is due to a choice of investing strategies according to the opinions of the wealthy majority. Relating this to the turbulence of water molecule effect, this is also due to a "condensation effect". Investors in investment groups, with potential investment strategies come together and discuss future stock market predictions. Do rich minds think alike? Overall the so called "condensation effect" in economy and the Lorenz curve should be modified as:

$$\boldsymbol{y = (1 - f)[x + (1 - x)ln(1 - x) + f \pm (1 - x)} \qquad \text{Eq.(2)}$$

where the last term is the delta-function, and 'f' is the fraction of income in the "Bose condensate". As shown in Fig. 5 at the end of this section, Eq. (2) depicts the distribution of individual income. The Eq. 2, represents the accurate "condensation effect" with f = 16% in 1997. When it comes to wealthy investors with income in the top 16% of America, their great minds do think similarly. A small 3% of Americans represent almost 16% of the American wealth; therefore, if bad decisions are made with all the wealthier Americans thinking similarly, then not only are they affected, the whole world's economy is affected. These graphical representations in our project only represent America. Imagine in various countries around the world, when poor decisions are made by the small percent of people who control more of a country's wealth. They all usually think similarly when it comes to investing large amounts of money, then not only does their own country feel the wrath of their poor decisions, in various ways countries around the world are affected by their decisions!

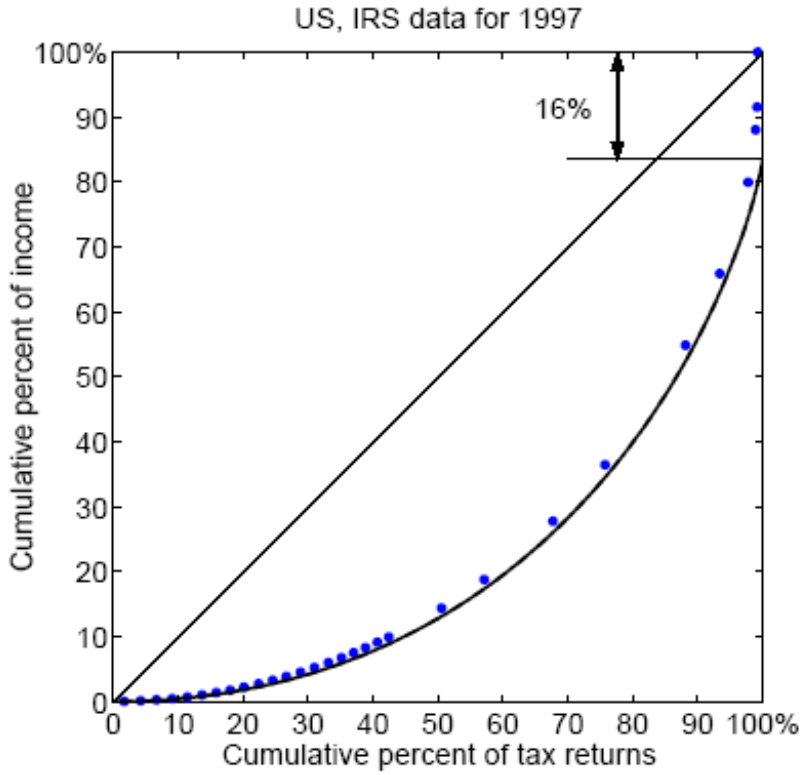

Figure 19 – Theoretical Lorenz Curve compared with the data point from IRS
Silva, Christian A., and Victor M. Yakovenko.
"Temporal Evolution of the "Thermal" and "Superthermal"."
Univ. of Maryland. 15 Jan. 2005. UMD. 30 Nov. 2007

3.5.2: HOT! HOT! HOT! - Thermal and Super Thermal

Solely speaking about America, personal income distribution in the USA has a well-defined two-class structure. The majority of population (97–99%) belongs to the lower class characterized by the exponential Boltzmann-Gibbs ("thermal") distribution, whereas the upper class (1–3% of population) has a Pareto power-law ("superthermal") distribution. By analyzing income data for 1983–2001, studies have shown that the "thermal" part is stationary in time, except for a gradual increase of the effective temperature which relates economically to inflation. The "superthermal" tail swells and shrinks following the stock market activity. Discussions around the world on these concepts of "thermal" and "super-thermal" personal income

distributions are flourishing. The study of social equilibrium/inequalities, principles of maximal entropy, and quantitative studies showing that these concepts apply to the majority of the world are being looked over by students at universities such as Cambridge, University of Maryland, and the University of Houston.[190]

Boltzmann-Gibbs probability distribution $P(\varepsilon)\propto\exp(-\varepsilon/T)$, where ε is energy, and $T=\langle\varepsilon\rangle$ is temperature.

Pareto probability distribution $P(r)\propto 1/r^{(\alpha+1)}$ of income r.

Analogy: energy $\varepsilon \leftrightarrow$ money $m \Rightarrow P(m)\propto\exp(-$

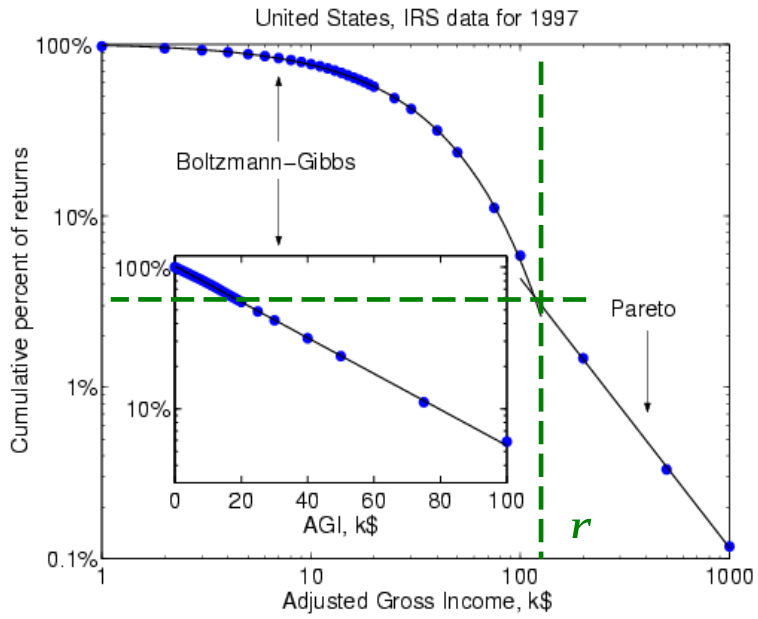

Two-class society

Upper Class

- Pareto power law
- 3% of population
- 16% of income
- Income > 120 k$: (investment capital)

Lower Class

- **Boltzmann-Gibbs exponential law**
- 97% of population
- 84% of income
- Income < 120 k$: wages, salaries

"Thermal" bulk and "super-thermal" tail

Figure 20 – Thermal Social Classes
Silva, Christian A., and Victor M. Yakovenko. "Temporal Evolution of the "Thermal" and "Superthermal"."

[190] Silva, Christian A., and Victor M. Yakovenko. "Temporal Evolution of the "Thermal" and "Superthermal"." Univ. of Maryland. 15 Jan. 2005. UMD. 30 Nov. 2007 <http://www2.physics.umd.edu/~yakovenk/papers/EPL-69-304-2005.pdf>.

The origination of these two classes can be derived from the difference in the 'source of salary'. For the lower class, thermal income is based solely on salaries and wages. For the upper class, superthermal income is based on capital gains, commission, and investments. The income between these two classes is very dynamic and can be represented using econophysics. The income dynamics are described using both additive and multiplicative diffusion. This method requires more advanced mathematics. For an easier social generalization, people tend to think of the classes as a separation between the 'employees' and the 'employers'. This generalization is described throughout Karl Marx's work. Many econophysicists and mathematicians are trying to use computer simulations to show the change from initial 'equal agents' and see how over time an emergence of social classes evolves into Social Capitalism. These basic functions and models ultimately depict the basic concepts that rest behind the base architecture of Social Capitalism using thermal econophysics.

The use of econophysics has evolved within the last decade and will continue to be used to analyze economical issues that affect our world. Combining math, physics, computer science, and economics is essential to studying econophysics. There are econophysics conferences being held worldwide. Within America the University of Houston offers a PhD in Econophysics, and around the world various countries are using it to study their own national economics. Mistakes or misconceptions of our investment markets can prove to be detrimental to everyone on this Earth. We must be prepared, as a world, for social problems that may lie ahead. With the help of econophysicists, I believe we can be prepared to make correct future decisions and also help various countries get back on their feet economically.

3.6: Risk Engine Review

3.6.1: Approach to the Problem

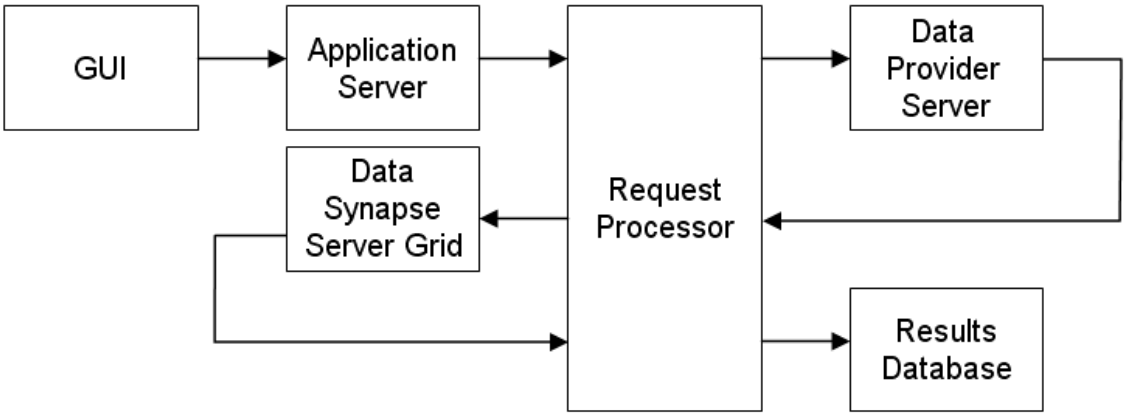

Figure 14 - The Risk Engine Architecture Block Diagram

The Risk Engine is an over encompassing term for a system of servers that work together to produce risk calculations, the Actual P&L, and the general ledger. This can be seen in the figure above. The section that this upgrade focuses on rests inside of the Data Synapse server grid. The entire server grid has 960 cores that are working in tandem. The risk engine utilizes 160 of these cores. That is where the focus for optimization will be spent.

In order to understand what is going on, however, it is important to elaborate on the above block diagram. The GUI (graphical user interface) is a client-side program. This program can be installed on a work station inside of Bank of America and it allows people to put in request for jobs. Jobs are certain tasks that encompass a specific set of information. An example of a job would be calculating the London credit spread P&L. This job would be specifically set on calculating the credit spread P&L for London and would not encompass the other information. When the London P&L is calculated, it is not done all at once. This is because if a part of a job fails to run, they still want the other information to be accessible. This provides a failsafe for the other portions of the P&L. If a job fails to run, it will be manually ran again when the error has

been detected. The application server servers a small purpose but it directly talks to the request server. The application server is designed to be the translator for the GUI and the request processor. That is to say, it will take the GUI information and send it to the request processor.

The request processor is responsible for issuing the request, monitoring the request, and finalizing the request. In this sort of environment, it is a lot like a manager. It will receive the information from the application server about a job that needs to be done. The request processor will then put it into a queue. Once the system is ready for the request processor, it will request the proper information from the data provider server. After the information is retrieved, the request processor then acknowledges that it has the information and forwards the information and the job to the grid. The proper processors inside of the grid will have their resources pooled together to calculate the job. Once the job is complete, the request processor will be notified and this will publish the information to the results database.

The data provider server is what directly talks to the database systems: IRP, Advantage, GDS, LoanTrak, Camden, Bloomberg, and Alice. The data provider server will take the request from the request processor and will go out and find the appropriate information to compute the requested job. For example, if London credit spread P&L is in the request then it will request the information from Odin and IRP. It calls Odin for the marks made by the traders but IRP will have the information from eBlotter. This allows the Risk Engine to be able to accommodate for unwinds and assignments that should effect Odin originally. As stated before, this will send the information to the request processor saying that the data is ready and this will get passed onto the Data Synapse server grid.

The server grid will then appropriate the portion of the grid that has been allocated to the specific job type. For Risk Engine, 160 of the cores that are available will be dedicated to computing the job. This enables Risk Engine to produce London P&L files very quickly as each job will take a matter of 2-10 seconds. The most time consuming section of the Risk Engine is computation of risk. The 160 cores that are used are stored on 40 1U rack servers. 1U is 1.75 inches tall and 19 inches wide. A rack server is a server that's chassis is specifically designed to be mounted in a rack. Common rack sizes are 24U and 42U. Each of these 1U rack servers is a Hewlett-Packard ProLiant DL360 Generation 4(G4p). The HP ProLiant servers that are being used currently hold two dual-core 2.8GHz Xeon processors.

Once the engine has completed the job it will send the information to the request processor and inform the processor that it is done. This will allow the information from the engine to flow to the results database. The results database will then store all of the information produced by the risk engine. This information is saved in a simple text file and will later be pulled by the East Bank Process to be put into an Excel based file.

We have the foundation of the Risk Engine and the area for optimization has presented itself as the 160 cores that the current Risk Engine is using. The space used by the Risk Engine's rack mounted servers can easily fit inside of a 42U rack. This leaves room for improvement as there is now technology available that can give equal processing capability and use less space. With a compression of rack space, a logical step would be a reduction in power consumption. This becomes the main area of interest for the optimization of the Risk Engine. A look at reducing the power consumption while offered a comparable computing apparatus is where this Risk Engine upgrade review comes from.

One of the first aspects that was investigated is finding a newer version of the HP DL360 servers. This revealed that there is a 5th generation version of the HP DL360 servers that utilize quad-core Intel Xeon Processors with an upwards clocking speed of 3.16GHz. This would be one avenue to explore. However, to get a feel for this option it would have to be weighed against other options. This lead to a search for an alternative type of computing apparatus that offered amazing benchmark numbers. This led to an investigation of the Cell processor as a means of comparing an alternative option to a simple upgrade in the system.

The problem that this presented was how the systems would be compared if they didn't have a similar architecture. The answer comes from an understanding of financial calculations. Most financial calculations are done using double precision floating point arithmetic. This leads to using floating point operations per second (FLOPs) as a comparison point. With further investigation, today's computers are able to process in the giga-FLOPs or GFLOPs. This comparison would have to be combined with a power analysis and gives us GFLOPs per watt. However, finding the double precision FLOPs rating isn't exactly easy as each processor has its own way of seeking the optimized amount of calculations. This would require some sort of benchmark to be able to compare across the board.

By looking at the Cell's ratings, the benchmark of choice was LINPACK benchmark. The LINPACK benchmark is used as the benchmark for rating the top supercomputers of today. It is also a rating of a system's capabilities to perform floating point operations. This seemed to lend itself to a fair comparison point as the LINPACK benchmark is based around the ability to

solve a very dense system of linear equations. [191] This would be the basis for comparing two different options with a common benchmark.

A Xeon core is capable of approximately capable of providing 2 single or double precision floating operations per cycle.[192] This sets the 2.8 GHz Xeon processors that Bank of America is currently using at 11.2 GFLOPs. This is the theoretical maximum and LINPACK data will often have processors that fall below their predicted capabilities. However, for the sake of making the argument for an upgrade, it will be assumed that these Xeon processors are able of consistently putting out 11.2 GFLOPs per processor (5.6 GFLOPs per core). This means that each rack server is capable of outputting 22.4 GFLOPs of single and double precision floating point operations. This is so because there are two of these dual core 2.8 GHz Xeon processors in each of the rack servers.

Next, it is important to take account of the power that is being consumed by these 2.8 GHz Xeon processor rack servers. These dual-core processors are rated for 135Ws. Two of them would mean that there is a total of 270Ws being used between the two processors. The power supply is capable of outputting a peak of 585W while its steady-state power is 535W. By using the power calculator provided by Hewlett-Packard the power used when the system is idle is 237W and when it is at a 30% workload is 295W. At a 70% workload it is 376W and at full workload it is 436W. These results were achieved by selecting the appropriate processor, setting the processor setting to two, using 1 PCI card, and assuming there are 1GB of RAM and 1 72GB hard drive. This gives us four points of comparison; however the idle power is only necessary

[191] "The Linpack Benchmark." Top 500 Supercomputing Sites. 27 Nov. 2007 <http://www.top500.org/project/linpack>.
[192] Dongarra, Jack J. "Performance of Various Computers Using Standard Linear Equations Software." Netlib. 20 Nov. 2007. University of Tennessee. 27 Nov. 2007 <http://www.netlib.org/benchmark/performance.pdf>.

when the system is not performing any calculations. These are also approximations but this gives us comparison points. We also know that the system cannot exceed 585W.[193]

By putting together this information, we can then make a comparison of GFLOPs versus watts. If the system is running at the full workload power of 436W and we compare it to 22.4 GFLOPs, the ratio is 0.051 GFLOPs per watt. This gives us not only the single precision but also double precision GFLOPs per watt of the system. The Cell will vary from this and offer two different GFLOPs per watt because it offers a much faster processing of single precision FLOPs.

The next important aspect of this was to find a price. However, this wasn't an easy task as my connections didn't have the price nor did they know who would know. However, to solve this problem I looked at the base price that HP claims their ProLiant DL360 G4p servers are priced at, which is $2,644.[194] This sets up a basis for the comparison. Since the Risk Engine DL360 G4ps have two of these dual-core processors and 1GB of memory, the price that will be used as a comparison will be $5,100. This is a fair price as server distributor websites have a single CPU version priced at approximately $3,500. This is a difference of $806 from the base price. This difference was multiplied by two and added to the base price.[195] This price also accommodates the fact that technology becomes cheaper as time advances. The $5,100 price tag, allows a comparison of price with the GFLOPs available. This gives us 0.0044 GFLOPs per dollar. This gives us three measures that we can use to compare the G4p to the G5 and to the Cell processor boards provided by Mercury Systems.

---

[193] "HP ProLiant DL360 Generation 4p (G4p)." Hewlett-Packard. 27 Nov. 2007 <http://h18004.www1.hp.com/products/quickspecs/12166_div/12166_div.pdf>.
[194] "ProLiant DL300 Series." Hewlett Packard. 27 Nov. 2007 <http://h10010.www1.hp.com/wwpc/ca/en/sm/WF04a/12146350-12146352-12146476-12146476-12146498.html>.
[195] "HP ProLiant DL360 G4p 2.8GHz Intel Xeon." Bizrate. 27 Nov. 2007 <http://www.bizrate.com/desktopcomputers/hp-proliant-dl360-g4p-2.8-ghz-intel-xeon--pid469987842/compareprices.html>.

3.6.2: Solutions to the Problem

3.6.2.1: The 3.16GHz Xeon X5460 HP DL360 Server Option

With the DL360 series by Hewlett-Packard, there are always two slots for processors to be placed into. With the fifth generation of the HP DL360 comes support for quad-core Intel processors. This allows a total of 8 cores to run inside of one HP DL360. As stated before, a Xeon is theoretically capable of performing 2 floating point operations per clock cycle. When applied to one Xeon core clocked at 3.16GHz, this results in 6.32 GFLOPs. When applied to the four Xeon cores that are present in the X5460, this results in 25.28 GFLOPs. This is for double and single precision.

This answer is already larger than the two dual-core processors in the DL360 G4p that was discussed above. This server is already 2.8 GFLOPs faster. This comes directly from having Xeons that are being clocked at higher frequencies and having multiple cores on one chip. When the entire board is evaluated with 2 quad-core processors, this yields a result of 50.56 GFLOPs per DL360 G5s. This one DL360 G5 based on GFLOPs equivalence is approximately 2.3 DL360 G4ps.

For the power comparison, the HP DL360 G5 power supply can output a maximum of 854W. The steady-state power capabilities are rated at 700W. The quad-core Intel Xeon Processors operate at 120W. This is already an improvement of the current dual-core Intel Xeon Processors that are being used in the Risk Engine's DL360 G4p, which are 135W. By using the power calculator and using similar settings to the HP DL360 G4p power supply, at a maximum workload the power consumed is 413W. At a 70% workload, the power requirement is 361W. While at a 30% workload, the power required is 290W and when the system is idle the power requirement is 233W. The idle power requirement is 4W lower with the G5 than with the G4.

This result makes sense because the processors only use less energy and the design of the G5 has probably been refined.

Assuming that the maximum GFLOPs can be pulled out of this system and using the 100% workload power rating, we get a ratio of 0.122 GFLOPs per watt. This is an improvement over the 0.051 GFLOPs per watt of the current G4p series that is used. This puts the G5 at being able to provide 2.37 times the GFLOPs per watt rating of the G4p series. This makes sense as there are four additional cores in the G5 version that was selected and they are clocked at a higher frequency. When the fact that there are four cores present is factored in, it is easy to see that each core is only using 30W of power. This is an improvement over the G4p's dual core processors where each core requires 67.5W.

The price of a G5 with two of the X5460s is approximately $5,530.[196] This was given by increasing the ram to 1GB, adding a 72GB hard drive, and upgrading the processors to 2 X5460s. The result that is yielded from this is a rating of 0.0091 GFLOPs per dollar. This is an increase over the 0.0044 GFLOPs per dollar of the original option. It is actually over 2 times that of the original. Thus with this option, not only is there more value in the product given by the GFLOPs per dollar rating but space will be condensed into smaller racks and less power will be used.

[196] "HP Proliant DL360 G5 Entry." Cost Central. 27 Nov. 2007 <http://www.costcentral.com/proddetail/HP_ProLiant_DL360_G5_Entry/416559001/K23560/>.

G4p Workload	Per Hour	Per Day	Per Week	Per Month	Per Year	Per 5 Year	Per 10 year
Idle	1.40	33.60	235.23	1,008.14	12,265.72	61,328.58	122,657.17
30%	1.74	41.83	292.80	1,254.86	15,267.45	76,337.27	152,674.54
70%	2.22	53.31	373.20	1,599.41	19,459.53	97,297.67	194,595.34
100%	2.58	61.82	432.75	1,854.64	22,564.78	112,823.89	225,647.79
G5 Workload	**Per Hour**	**Per Day**	**Per Week**	**Per Month**	**Per Year**	**Per 5 Year**	**Per 10 year**
Idle	0.62	14.87	104.07	446.01	5,426.42	27,132.08	54,264.15
30%	0.77	18.50	129.53	555.12	6,753.91	33,769.54	67,539.07
70%	0.96	23.03	161.24	691.02	8,407.45	42,037.25	84,074.50
100%	1.10	26.35	184.46	790.56	9,618.50	48,092.48	96,184.96

Table 11 - Comparison of power consumption over various time periods

In the above table, the G4p is the price for running all 40 servers while the G5 is running a processing equivalent with 18 servers. The real comparison to be made is when this option will pay itself off in the power that it saves. This option fails to pay itself off before a 10 year mark with the assumption that power is 14.77 cents per kilowatt hour.[197] It will easily pay for itself in 5 years if power rises to 23 cents per kilowatt hour but otherwise it won't. The likeliness of this event is very uncertain due to the depreciation of the United States dollar and the rise in oil prices. Although, despite power will be saved, this option does not pay for itself in the near future. It is best to wait for the sixth generation of the DL360 that would more than likely introduce more cores and faster overall processors.

3.6.2.2: The Cell Broadband Engine Processor (Current)

The Cell processor has many different aspects and it is easy to get lost in the tidal wave of information that is presented by IBM, Mercury Systems, Sony, and Toshiba. That is why, there will be extra care taken in this section to explain where the numbers and figures are coming from. This section will evaluate multiple options, the Mercury System 1U Cell rack server, the Mercury Systems Cell blade server, and the IBM BladeCenter QS21.The IBM BladeCenter

[197] "Average Retail Price of Electricity, March 2007." Public Policy Institute of New York State, Inc. 27 Nov. 2007 <http://www.ppinys.org/reports/jtf/electricprices.html>.

QS21 will only be evaluated on its price as they do not provide adequate information for a power comparison.

A single cell processor is approximately capable of computing 230 GFLOPs for single precision and 21 GFLOPs for double precision. These numbers includes the PPE and SPEs. Most quoted numbers are just the SPEs combined capabilities of computation. The SPEs alone can provide 204.6 GFLOPs single precision and 14.6 GFLOPs for double precision. It is immediately clear that the double precision capabilities of the Cell do not offer a significant increase when compared to the Xeons that have been presented before. There is hope for the double precision processing of the Cell. This comes from the fact that IBM, Toshiba, and Sony have come up with a way through code to exploit the single precision calculations, in order to perform double precision floating point operations. Through this process, the Cell is capable of performing 100 GFLOPs of double precision calculation.

	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
Single Precision GFLOPs	230.4	204.8	204.8
Double Precision GFLOPs	20.8	14.6	100

Table 12 - Summary of 3.2 GHz Cell GFLOPs capabilities

The Mercury Systems 1U Dual Cell-Based System 2 has a typical power consumption of 575W. This immediately stands out at significantly higher than the G4p and G5 systems. The G4p at a maximum workload under the conditions provided operated at 436W and the G5 operated at 413W. Each of these systems store two Cell processors on board and offer their combined computing power. When comparing the electrical power to the GFLOPs of the Cell with the PPE and 8 SPEs working together this results in 0.801 GFLOPs per watt for the single precision and 0.072 GFLOPs per watt for double precision calculations. The Cell destroys the

quad-core Xeon processor with an impressive 6.7 times increase. On the other hand, when the double precision calculations are compared, the Cell loses to the quad-core Xeon with a 0.6 times decrease.

With only the 8 SPEs operating, this yields a GFLOPs per watt ratio of 0.71 for single precision and 0.05 for double precision. This shows one of the flaws with the Cell. Against conventional hardware, the double precision is not capable of performing at the same level. This can be assisted through code to perform double precision calculations by exploiting the single precision calculation speed. The results that follow are a GFLOPs per watt ratio of 0.71 for single precision and 0.35 for double precision. This allows the Cell processor to lead the G5 set up by 2.9 increase in computational ability. The problem with this approach is that the coders will have to write more elaborate code to be able to exploit this single precision upgrade.

The option above only requires a rack to be placed into while the Mercury Systems Dual Cell-Based Blade 2 is designed to be seated in a BladeCenter H chassis (IBM). The BladeCenter H chassis is 9U and will hold 14 of the Dual Cell-Based Blade 2s or 14 of the IBM BladeCenter QS21. The Dual Cell-Based Blade 2 operates at a typical drain of 380W. This provides a much greater gain for the Cell but this means an approximately $4,000 purchase of an IBM BladeCenter M chassis. When the PPE and 8 SPEs are in full use, this puts the GFLOPs per watt rating for single precision at 1.21 while the double precision sits at 0.11. The Cell has blown away the quad-core Xeon in single precision capabilities and has closed its gap to the Xeon's double precision capabilities.

With only the 16 SPEs running on the Dual Cell-Based Blade 2, this will yield a similar result where the single precision is much higher at 1.07 GFLOPs per watt while the double

precision is resting at 0.08 GFLOPs per watt. With the specific double precision enhancement through single precision exploitation, the Cell yields a respectable 0.53 GFLOPs per watt while the single remains at 1.07 GFLOPs per watt. The problem with this is that the coders will have to follow the process defined by IBM, Sony, and Toshiba to be able to exploit the processor in this way.

1U Dual-Cell Based System 2	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
Single Precision GFLOPs/watt	0.80	0.71	0.71
Double Precision GFLOPs/watt	0.07	0.05	0.35
Dual Cell-Based Blade 2	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
Single Precision GFLOPs/watt	1.21	1.08	1.08
Double Precision GFLOPs/watt	0.11	0.08	0.53

Table 13 - Summary of 3.2 GHz Cell GFLOPs per watt

	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
SP Comparison to QC X5460	6.68	5.94	5.94
DP Comparison to QC X5460	0.60	0.42	2.90

Table 14 - Summary of 3.2 GHz Cell GFLOPs compared to 3.16GHz X5460 Xeon (GFLOPs/W)

The cost of a 1U Dual Cell-Based System 2 is approximately $15,000. This price came from a representative inside of Mercury Systems specifically for this analysis. This price would vary with the volume of the actual purchase but this is the best estimate for how much each would cost. With this we will compare the single precision range of 230.4 GFLOPs and 204.8 GFLOPs, to receive a rating of 0.0307 and 0.0273 GFLOPs per dollar respectively. While the double precision for the PPE and 8 SPEs yields a result of 0.0028 GFLOPs per dollar and with the 8 SPEs alone this results in a 0.0019 GFLOPs per dollar ratio. If the code is optimized to exploit the single precision speed for use of double precision calculations, the ratio is 0.0133 GFLOPs per dollar. This is greater than the 0.0091 GFLOPs per dollar of the X5460 processor.

The cost of a Dual-Cell Based Blade 2 is approximately $12,000. This price was obtained through similar methods as the above. This results in a rating of 0.0384 GFLOPs per dollar for single precision Cell calculations with the PPE and 8 SPEs and 0.0341 GFLOPs per dollar for single precision Cell calculations with only the 8 SPEs. For double precision, the PPE and 8 SPEs option yields 0.0035 GFLOPs per dollar, the 8 SPEs option yields 0.0024 GFLOPs per dollar, and the 8 SPEs with the code enhancement results in 0.0167 GFLOPs per dollar. The price of the BladeCenter H chassis was not figured in because it throws off a one on one calculation. This is especially true if one is going to buy the 14 blades that will fill the entire BladeCenter H chassis. This'll spread out the amount further than if 1 blade is purchased.

The cost of a BladeCenter QS21 is priced at $9,995 by IBM's website. This does not include shipping or tax. This option is more affordable than Mercury System's blade option as can be seen from the previous section. While utilizing the PPE and 8 SPEs, this results in a 0.0461 GFLOPs per dollar ratio for single precision and while only using 8 SPEs results in a 0.0410 GFLOPs per dollar ratio for single precision. Double precision with the PPE and the 8 SPEs yields a result of 0.0042 GFLOPs per dollar ratio. With only 8 SPEs, this yields a result of 0.0024 GFLOPs per dollar. The single precision code exploit for double precision calculations allows this version to reach 0.0200 GFLOPs per dollar. This is approximately 2.2 times larger than the X5460's 0.0091 GFLOPs per dollar.

1U Dual-Cell Based System 2	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
Single Precision GFLOPs/Dollar	0.0307	0.0273	0.0273
Double Precision GFLOPs/Dollars	0.0028	0.0019	0.0133
Dual Cell-Based Blade 2	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
Single Precision GFLOPs/Dollar	0.0384	0.0341	0.0341
Double Precision GFLOPs/Dollar	0.0035	0.0024	0.0167
IBM BladeCenter QS21	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
Single Precision GFLOPs/Dollar	0.0461	0.0410	0.0410
Double Precision GFLOPs/Dollar	0.0042	0.0029	0.0200

Table 15 - Summary of 3.2 GHz GFLOPs per dollar

1U Dual-Cell Based System 2	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
SP Comparison to G4p	6.98	6.21	6.21
DP Comparison to G4p	0.63	0.44	3.03
SP Comparison to G5	3.38	3.00	3.00
DP Comparison to G5	0.30	0.21	1.47
Dual Cell-Based Blade 2	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
SP Comparison to G4p	8.73	7.76	7.76
DP Comparison to G4p	0.79	0.55	3.79
SP Comparison to G5	4.22	3.75	3.75
DP Comparison to G5	0.38	0.27	1.83
IBM BladeCenter QS21	Cell (PPE + 8 SPEs)	Cell (8 SPEs Only)	Cell (8 SPEs with DP Code Enhancement)
SP Comparison to G4p	10.48	9.31	9.31
DP Comparison to G4p	0.95	0.66	4.55
SP Comparison to G5	5.07	4.50	4.50
DP Comparison to G5	0.46	0.32	2.20

Table 16 - Summary of 3.2 GHz Cell compared to the G4p and G5 (GFLOPs/dollar)

Since finance applications are usually performed with double precision calculations the Cell doesn't truly shine in this area. Although, if the code is properly written to exploit the single precision capabilities for calculating double precision floating point operations then the Cell truly shines. This sort of analysis is filled with gray areas because comparing different processors is very difficult. A processor that excels for one thing may not excel in another area. This is true with GPUs as they're very good at doing vector calculations but they aren't design to handle the generalized processing that a CPU will deal with. The Cell offers a unique combination of providing immense calculating power while still serving as a general processor.

1U Dual-Cell Based Server	Per Hour	Per Day	Per Week	Per Month	Per Year	Per 5 Year	Per 10 Year
Cell SP (PPE + 8 SPEs)	0.17	4.08	28.54	122.30	1,487.93	7,439.65	14,879.30
Cell SP (8 SPEs)	0.25	6.11	42.80	183.44	2,231.89	11,159.47	22,318.95
Cell SP (8 SPEs Enhanced Code)	0.25	6.11	42.80	183.44	2,231.89	11,159.47	22,318.95
Cell DP (PPE + 8 SPEs)	1.87	44.84	313.89	1,345.25	16,367.23	81,836.14	163,672.28
Cell DP (8 SPEs)	2.72	65.22	456.57	1,956.73	23,806.88	119,034.38	238,068.77
Cell DP (8 SPEs Enhanced Code)	0.43	10.19	71.34	305.74	3,719.82	18,599.12	37,198.25
Dual-Cell Based Blade	**Per Hour**	**Per Day**	**Per Week**	**Per Month**	**Per Year**	**Per 5 Year**	**Per 10 Year**
Cell SP (PPE + 8 SPEs)	0.11	2.69	18.86	80.82	983.33	4,916.64	9,833.28
Cell SP (8 SPEs)	0.17	4.04	28.29	121.23	1,474.99	7,374.96	14,749.91
Cell SP (8 SPEs Enhanced Code)	0.17	4.04	28.29	121.23	1,474.99	7,374.96	14,749.91
Cell DP (PPE + 8 SPEs)	1.23	29.63	207.44	889.04	10,816.60	54,083.01	108,166.03
Cell DP (8 SPEs)	1.80	43.10	301.73	1,293.14	15,733.24	78,666.20	157,332.40
Cell DP (8 SPEs Enhanced Code)	0.28	6.74	47.15	202.05	2,458.32	12,291.59	24,583.19

Table 17 - Summary of power usage by the Cell over various time periods

In the above table, each is set to an equivalent amount of Cell devices to equal the computing power of the G4p. If the code is written to exploit the single precision calculation speed for double precision calculations, the Cell pays for itself in approximately 3 years. This would be because 5 Cell blades would only total approximately $64,000 and with the BladeCenter Chassis H that puts the total at $68,000. That is because a Cell that is operated in this fashion for a year will only consume $2,458 worth of power while the current system would approximately consume $22,564 worth of a power. This doesn't include the cost of development that would directly follow converting the architecture. Although, there is an entire team dedicated to Risk Engine and its constant improvement. This could be applied to a side project for them to bring the code over to the Cell processors. The difference is clear, consume less power by using Cell based options and save in power bills.

The part that is unclear is how much effort is required to truly maximize the potential of the Cell and how difficult is it to port the code from the RISC structure that is currently in use to a CISC structure. There is hope as games made for personal computing are transferred to the Playstation 3 on a regular basis. This isn't without snags as certain games like Half-Life 2: The Orange Box and Assassin's Creed have had problems running on the Playstation 3 while their PC and XBOX 360 counterparts run smoothly.[198] Sony has always introduced hardware that needs time to fully realize its potential. This has held true with Playstation 1 that had Final Fantasy VIII, which pushed full motion videos to an entirely new level. The Playstation 2 had God of War and Metal Gear Solid: Snake Eater, which both pushed the graphical limits of the Playstation 2.

It is clear that the Cell needs to be able to perform consistently in the area of double precision. The optimization of double precision by taking advantage of the single precision capabilities through code will be a time consuming rewriting stage. It will already be a challenge to transfer the current code over to the Cell processor. Using the single precision capabilities for double precision calculations is a work around and this is not acceptable for a system that is being built to last. That is why, at this time, the Risk Engine should not be completely transferred over to the Cell structure. An effort, on the other hand, should be made to further investigate the Cell structure. Mercury Computer Systems now offers a software development kit that has numerous mathematical libraries that have been optimized for the Cell technology.

[198] Ellis, David. "Half-Life 2: the Orange Box - This Box Has Problems." 1up. 21 Nov. 2007. 2 Dec. 2007 <http://www.1up.com/do/previewPage?cId=3164539>.
"Ubisoft Acknowledges Assassin's Creed Freezes on PS3." Gamekiq. 2 Dec. 2007 <http://www.gamekiq.com/wordpress/2007/11/18/ubisoft-acknowledges-assassins-creed-freezes-on-ps3/>.

3.6.2.3: The Cell Broadband Engine (Near Future)

The Cell will be receiving an update mid-2008. This update will include enhanced double precision support to push the Cell's computation of these types of floating point operations to the 100 GFLOPs range.[199] From the previous section, this allows the Cell to shine on a strictly hardware basis without resorting to specific coding techniques to exploit the single precision capabilities of the Cell. This would provide 200 GFLOPs for double precision calculations on the Dual Cell-Based server options. Once this new Cell and hardware that supports it is released, the Cell's benefits will become overwhelming.

Once the new iteration of the Cell is released that is when it will become a viable option for starting the shift from conventional processors to the Cell. This shift over to the Cell does not have to be sudden and abrupt. Instead, the Cell can be programmed to work well with Windows based servers. Mercury Computer Systems has mentioned this combination of Cell processors with conventional processors is usually very effective.[200] This would allow the Risk Engine team to start developing specific routines for the Cell processor that are equivalent to the current Risk Engine software.

As time progressed, more Cell based solutions could be added to the racks and more code could be converted over until the system has been optimized for power consumption and for computing capabilities. On a side note, in the future IBM, Sony, and Toshiba have plans of creating a 2 PPE and 32 SPE version of the Cell processor. This will be their attempt at offering 1 TFLOPs for a single processor. With an effort being placed on the next iteration of the Cell, Bank of America could easily implement the 2 PPE and 32 SPE version of the Cell processor with a little extra effort. If the chip lives up to its promise, the power consumption by Bank of

[199] Mercury Computer Systems. Telephone interview. 15 Nov. 2007.
[200] Mercury Computer Systems. Telephone interview. 15 Nov. 2007.

America in the area of Risk Engine computation will be dramatically decreased as will the space required.

With these hopes for the future, the Cell-based option becomes a reality for improvement in mid-2008 and picks up momentum as the years go on. That is why this option should be revisited in mid-2008 to perform a cost analysis and the feasibility of starting to implement a Cell-based solution. Programmers could be taken from the video game industry and applied to working on Cell-based solutions as more and more will be trained to program for the Cell. This will add to the overall value. All of these reasons form why the Cell isn't a viable option right now but in the near-future it will have great potential and should be explored further.

3.6.2.4: Data Synapse Grid Servers

This portion of the project was outside of the bounds of the original goals. The reason this is being included is because there is a great opportunity for optimization here. Currently, the Data Synapse server grid is a bunch of computers that have various form factors. These form factors are blades, server racks, and standalone work stations. Each of these is combined to work in parallel with one another pooling together their computation capabilities. The real reason for the interest in this particular region is that these computers are all running 1.2 GHz Pentium IIIs with 1 GB of RAM. A Pentium III is capable of performing 1 FLOP per cycle. This puts the computing capabilities at 1.2 GFLOPs per processor in the Data Synapse server grid.[201]

Knowing that there are approximately 800 processors being run, for each type of form factor it will be assumed based on conversations with Crystal James from Bank of America Distributed Computing Development team that there are four Pentium III processors in each type of form factor. That leaves 200 machines that are apart of the server grid. The assumption will be

[201] Dongarra, Jack J. "Performance of Various Computers Using Standard Linear Equations Software." Netlib.

that each machine consumes 120W. This is based on assuming that power is not consumed by any hardware but the actual CPUs.[202] This is assumption is also based on the fact that the server grid is a very mixed collection of computing assets.

Understanding that there are approximately 800 Pentium III processors being closed at 1.2 GHz this yields a shared computing capability of 960 GFLOPs while consuming 96kW. Using this information again gives a result of 0.01 GFLOPs per watt. As stated above, the G4p provides 0.051 GFLOPs per watt, the G5 provides 0.122 GFLOPs per watt, and the Cell-based solution can provide up to 0.53 GFLOPs per watt (double precision). This server grid was definitely a great idea at one point in time because it allowed them to recycle computers during an upgrade. Now it is a great area for optimization.

	Per Hour	Per Day	Per Month	Per Year	Per 5 Year	Per 10 Year
Current Server Grid	14.18	340.30	10,209.02	124,209.79	621,048.96	1,242,097.92
G4p 2.8 GHz Dual-Core Option	2.77	66.46	1,993.74	24,257.14	121,285.69	242,571.37
G5 3.16 GHz Quad-Core Option	1.16	27.82	834.48	10,152.86	50,764.28	101,528.57
Current Cell 1U MCS Rack Option (SP)	0.25	6.11	183.44	2,231.89	11,159.47	22,318.95
Current Cell MCS Blade Option (SP)	0.17	4.04	121.23	1,474.99	7,374.96	14,749.91
Current Cell 1U MCS Rack Option (DP)	2.80	67.26	2,017.88	24,550.84	122,754.21	245,508.42
Current Cell MCS Blade Option (DP)	1.85	44.45	1,333.55	16,224.90	81,124.52	162,249.04
Future Cell 1U MCS Rack Option (SP)	0.25	6.11	183.44	2,231.89	11,159.47	22,318.95
Future Cell MCS Blade Option (SP)	0.17	4.04	121.23	1,474.99	7,374.96	14,749.91
Future Cell 1U MCS Rack Option (DP)	0.42	10.19	305.74	3,719.82	18,599.12	37,198.25
Future Cell MCS Blade Option (DP)	0.28	6.74	202.05	2,458.32	12,291.59	24,583.19

Table 18 - Power comparison for the Data Synapse Server Grid

It's easy to see that the current server grid is chewing up power. While systems that provide similar computing capabilities on the basis of FLOPs, will save a lot of money. It would take 43 G4ps, 19 G5s, 3 current Cell-based solutions (single precision), 33 current Cell-based solutions (double precision without single precision optimization), and only 5 for the future Cell-

[202] Shimpi, Amand L. "Cool Runnings: the Intel Bobsled Team." AnandTech. 30 July 2001. 29 Nov. 2007 <http://www.anandtech.com/showdoc.aspx?i=1509&p=3>.

based solutions with double precision enhanced hardware. The assumptions made in the above power comparison are power costs 0.1477 cents per kilowatt hour, GFLOPs is an acceptable way of matching performances of systems, and that the minimum number of new servers is acceptable to run the current tasks.

This analysis shows that what once was a good idea and worked well in the past by reusing hardware is starting to show its age. There are now more efficient ways of providing computing capabilities with less power being required. This would also save a tremendous amount of physical space by using these more efficient servers. The Risk Engine team had the right idea by trying to save energy and offer more computational capabilities by using the dual core processors. It is time for the Distributed Computing Development team to follow suit and perform a power analysis of their system, in order to save Bank of America thousands to hundreds of thousands of dollars.

4: Conclusion

4.1: Profit and Loss Workflow Gap Analysis

For the Flash P&L reporting system, the server based option would allow the Flash to perform complex calculations faster. This is because it would not have to share runtime with other processes like the current Flash does. This would enable Bank of America's London Flash P&L team to run jobs faster than before and be able to share information. While the .Net Flash viewer should also include the feature of grouping the tabs together. This enables the traders to group P&L categories that shift slowly and see the major changes quickly.

The Actual P&L reporting system consists of a multitude of different problems. Odin doesn't allow the traders to accurately predict their credit spread P&L if there are unwinds, assignments, or large changes in the credit spread. This should be changed in a future edition of Odin by allowing Odin to take information from eBlotter or IRP and to make corrections accordingly. It would also be advantageous to allow Odin to calculate the second derivative and to be able to account for convexity of the credit spread.

The East Bank Process should move to a server based option since it was originally designed to be a temporary solution and because it can be optimized further. The proposed solution for the East Bank Process includes allowing middle office to easily request reports while also being able to customize the fields that they want to look at. This would allow for a great variety of reports that could be generated and middle office wouldn't have to depend on anyone else to obtain a specific report.

A Flash & Actual comparison program allows the traders to have discrepancies quickly highlighted between the Flash and non-reconciled Actual P&L. This would also give the traders

access to the non-reconciled Actual P&L and enable them to be prepared to reconcile the differences with middle office. The report generated from this program would also help middle office by clearly highlighting the areas of difference. With the data pull feature, middle office would quickly be able to look at a trade and then see the original information that the value was based off of. This enables middle office to not have to manually pull up information and also allows them to consistently have information readily available.

The comparison program would also benefit from being able to compare the Flash, non-reconciled Actual, and the reconciled Actual. This would allow traders to quickly see the differences of the Flash and non-reconciled Actual when compared with the reconciled Actual. The data logging feature would be a great enhancement as it would allow for an analysis to find problems in the current Flash and Actual reporting systems. The data logging feature could be run for days, months, or even years and collect all of the appropriate information. The graphing feature would allow for quick visual analysis of the situation.

The new viewer for P&L could be based off of the currently in production .Net Flash viewer. This would put a tool that is currently being produced and add additional features to it, in order to save money in development costs. This would add value to the program as it would take on additional features and become more flexible.

All of these solutions would take time to develop but we believe that they would ultimately be beneficial to the company. This gap analysis has highlighted areas that need improvement and suggestions for how to develop upon the current system. It has also summarized the concerns of the current users of the Flash and Actual P&L.

4.2: Risk Engine Upgrade

This project has shown that while the Cell processor is extremely capable as shown previously, the uncertainty behind the difficulty of porting the code weighs negatively on it. This is especially true with the current iteration of the Cell hardware as its double precision capabilities are currently inferior to a quad-core processor. That is why if an additional server is needed, it should be a quad-core processor like the X5460 HP DL360 G5 that is reviewed above. This enables the Risk Engine team to make minor adjustments to their code and enjoy the benefit of top-of-the-line hardware.

This project has also shown that Bank of America can save money in the future by switching either the Cell or the X5460 HP DL360 G5. However, the HP DL360 G5 did not pay for the hardware investment within five years. The Cell did pay for itself within three years but the development costs behind switching to the Cell left an uncertainty on the feasibility of completely switching.

The additional examination of the Data Synapse server grid showed an area where considerable improvement could be made. The current systems consume a lot of power and while they were a great use of older machines in the past, their age is starting to show. Technology has made leaps and bounds since the Pentium III era. For this problem, both the X5460 HP DL360 G5 and the Cell-based solution pay for themselves. The Cell pays for itself in a much shorter span of time but the continual recommendation of this report is to wait until the double precision optimization through the actual hardware is released. This will allow for a simpler coding solution instead of constantly trying to exploit the single precision capabilities of the Cell to perform double precision calculations.

4.3: Econophysics

Interdisciplinary studies of physics and economics is a beautiful field that is growing around the world. With plenty of exposed research our world is taking a great step towards eliminating monetary turbulence and establishing a stabilized study of economics. With only basic ideas of econophysics described here in our project and solely directed in the world of thermal physics, it should be known that there is a dynamic world of research to be seen in this field. Our project has illustrated the basic concepts of only one single element. As a physicist interested in economics, the world should know that economic problems can be predicted. Everyone, even those on the social boundaries of our world, is affected by financial decisions that are made day to day. With this being said, as a world, we should be constantly studying our economy and trying to derive proper ways to predict future problems in our society to prevent naive financial decisions.

5: Recommendations and Future Work

We recommend that Bank of America takes our solutions and implements the most crucial ones first. This would include the reconciliation process renovation since the traders do need more information while they're waiting for the Actual to be reconciled. These solutions should be compared with the American P&L Program to see if it can be exported to London and modified for these purposes.

If the comparison program is implemented with the data logging tool, this could be used in a future project to highlight areas that need improvement. This would allow for plenty of data to be easily compared. This would be an excellent project for students if there is a systematic problem in the way things are computed. The project would include an attempt to find where errors consistently originate from and making suggestions based on the analysis of the collected data.

After the next iteration of the Cell processor is released with the double precision floating point hardware upgrade, an effort should be started in attempting to start converting some of the Risk Engine code over to the Cell processors. This would make for an excellent project for 4 Worcester Polytechnic Institute students as they are much cheaper than regular employees. Each of these students should have a strong background in computer science. Acceptable majors also include computer engineering and technical interactive media and game development.

Appendix A - Terminology

Asset Swap – This is when the interest-rate risk is eliminated from a trade and all that remains is the credit and funding aspect.

Assignment – Bank of America uses assignments when it takes a trade or a deal and moves it from one person to another.

Arbitrage – Under normal conditions, all markets should have the same price listed for their assets. Arbitrage is when there is a discrepancy in the market and one can buy at a lower price and then sell to the other market at a higher price. This usually only works with large trades as the fees associated with trading will need to be accounted for.

Basis Point – Abbreviated as bp and is equal to one-hundredth of a percent.

Bond – A form of credit where a buyer of a bond buys into the debt of an entity.

Cell – Shorthand for Cell Broadband Engine Architecture developed by IBM / the Sony Group / and Toshiba.

Credit Default Swap (CDS) – A form of credit derivatives where the risk of default is separated and sold to another company that acts like an insurance agency.

Credit Derivatives – A tool for mitigating risk by being able to separate it from the trade and selling it to a protection selling entity.

Credit Spread – The difference between a corporate bond and a government bond that are identical in everything except their credit rating.

CS01 – Credit spread per basis point.

DV01 – Dollar value per basis point. This means how much does the dollar value of an asset change with a one basis point change in the yield.

Exercise Price – The price at which an option can be used at.

Greeks, The – Used for evaluating financial figures and are come from the Black-Scholes model.

IR01 – Interest rate per basis point. This is how much does the interest rate of an asset change with one basis point change in the yield. This is also the common way to talk about how the interest rate has changed by measuring in basis points.

Loan – A form of credit where the lender gives the borrower something in return for a promise of compensation.

Option – A type of derivative that can be purchased that gives the user the right and is not an obligation to buy a stock at the exercise price.

P&L – Profit and loss tracks how much a company lost or made on a certain say. This is the general accepted form by the public.

P/L – Variant of P&L and is used mainly in the United States by Bank of America.

P&L – Variant of P&L and is used mainly in London by Bank of America.

PV01 – Present value per basis point. See DV01.

Strike Price – See exercise price

Reference Asset – A particular asset that is set as the foundation for pricing of a credit derivative and other financial devices.

Total Return Swaps – An off-the-books accounting technique where an asset is completely transferred to another entity who has space on their books to accept the asset. The entity that has the asset will then receive the returns on the asset and will return it when the original organization has room on their books.

Underlying Asset – See Reference Asset.

Unwind – A financial tool for closing or cancelling out an earlier transaction. This results in a settling of a deal and may include a cancellation fee. A partial unwind only cancels or closes out a part of a deal.

Appendix B – Midway Report to Bank of America

Last Update: 09 November 2007

Main Project Summary

Title: Profit and Loss Workflow Analysis

Name: Angelo Chandler
Name: Thomas MacDonald

Task:

To perform a gap analysis of the profit and loss workflow that currently happens at Bank of America in London. This is to prepare a plan for optimization and how to condense the workflow so that the current Actual P&L will be available at an earlier time than it currently is. The goal time is half past eight AM.

Contact List:

Abhi Vashishth – Mentor, Manager, General Overview
Alex May – In charge of the Flash P&L team in Middle Office
Dilruk Jayamanna – In charge of the Actual P&L team in Middle Office
Ben Cooper – Works on the trading floor, has solved certain P&L problems, extremely knowledgeable and helpful in almost all areas
Priscilla Wong – The lead for the Flash P&L development team for technology
Kate Schofield – Has worked with risk engine and is knowledgeable in the field.
Pete Spencer – Is working on revolutionizing P&L
Professor Gerstenfeld – Management and Financial Engineering Advisor (WPI)

Approach:
Defined: Friday November 2nd, 2007

1. Collect colleague information (90%)
    - Set up a list of people to contact for certain information
    -Collect the problems (Still waiting for Dilruk)
    -Collect possible solutions to the problems (Still waiting for Dilruk)
2. System Level Analysis (85-90%)
    -Systems have been identified
    -Connections between systems have been identified
    -Interfaces have been identified
    -Clearer system level diagrams need to be created
3. Brief Description of Systems Involved (65-75%)
    -Preliminary system descriptions have been written

-Review by Angelo and Thomas is imminent
-Flash and Actual systems have been identified
-Risk Engine software manual has been identified
-A better understanding of how things are actually fed information

4. Look at all processes
-Manual (55%)
-Technical (80%)
-Who? (75%)
-What? (85%)
-Where? (80%)
-When? (95%)
-How? (55%)

5. Develop a strategy (10%)
-Fundamental understanding of events has been established
-A grip of what is going on has been established
-In depth look at the reconciliation process is required

6. Proposition for Improvement (5%)
-Outline Areas for Improvement
-Research improvement areas
-Develop improvement plans and methods for each problem

7. Celebrate (0%)
-Celebration begins once the project has reached acceptable range

Flash P&L Progress

Main Drivers: Priscilla Wong and Alex May

Major Concerns:

- Partial Unwinds
- Cancelled / Removed Trades
- Duplication of new trades / NewBus tab
- Full Unwinds with an unsettled Fee
- Trades feeding the NewBus P&L Tab
- Timeliness of Running the Flash / System Performance
- Adjustment Tool requested
- Credit Spread Moves

Summary:

Contacts have been established for most of the most important areas. In Abhi's absence, it will be important to be able to go to the correct people for questions. Alex, Dilruk, and Ben have all been identified as very busy people. However, Ben is more than willing to sit down with us to review information. Ben's time should be used extremely wisely as he is extremely busy. This process is beneficial for both Ben and the Technology department.

A work around for unwinds for the Flash has been to run it through the Risk Engine. However, this takes additional time and defeats the purpose of Flash. Also, it has been made clear that a Flash should be computed faster and stay away from using too much advanced mathematics. There are problems with the reconciliation process as the Actual P&L team will compute the P&L and if the numbers don't match up they'll go to the trader's and ask why there is a difference. This shouldn't be a necessary step as this should be automatically detected and highlighted as a difference.

Flash has been identified as a program that is run on a desktop. When this program is running it will eat up the resources that are available on that computer rendering it unusable. Questions have been raised about bottlenecking the program and allowing other programs to use the CPU. A current work around is to change the definition of the job. If you map fewer strategies and increase the number of jobs, this will make the run a lot less work intensive. A bottleneck is also offered when the computer has to pull the appropriate curves from Camden. Another point of interest is that there is a desire to make the flash a server side program. They would also want to upgrade this to a near real-time or real-time calculator of Flash.

Differences of opinion of Flash P&L have raised interesting aspects of this analysis. A question being thought about now is the proper direction to take the Flash P&L. This is between making it simpler and quicker or making it more detailed. The more detailed option would run on the server side and be a real-time estimate of what is going on. However, by the definition of estimate this implies that it doesn't have to be exactly equal to the actual although it should be close. Currently the Flash P&L has trouble with unwinds and needs to pull more information from eBlotter.

Problems with the Flash P&L are well documented and will be addressed accordingly. Problems with the Actual P&L are still coming in and will need further investigation. Background research on financial topics has concluded. Each student should exhibit a solid understanding of financial engineering.

Ben has suggested that the Flash system should only focus on three types of P&L because most of them experience small changes. He has pointed out that new business, credit spread, and edit P&L are some of the most important as their changes can be very extreme. While he suggested putting carry, convexity, and interest-rate P&L all go into a large encompassing basket.

Actual / Explain P&L Progress

Main Drivers: Ben Cooper and Dilruk Jayamanna

Major Concerns:

- Unwinds
- Assignments

- Middle Office and Front Office Reconciliation

Summary:

Areas of concern have been established. A consistent area of concern is dealing with unwinds. Ben Cooper has also pointed out one of his biggest concerns is the assignment process. This is actually a cause to some P&L appearing in Edit P&L and New Business P&L. This splits up some of the P&L and makes it difficult to trace where the error has occurred. A suggestion has been made of redefining what constitutes New Business P&L. In the case of assignments, all of the P&L that comes from a new assignment should be listed under New Business P&L. The reconciliation time of the Actual P&L versus the Flash P&L has also been identified as a problem. This takes much too long and is the main area for optimization.

When the prediction and the actual P&L don't match, middle office will go to front office and ask why the numbers don't match. This leads to a lot of problems as front office's time will be spent on trying to find where the mismatch occurs. Middle office will basically just provide the P&L. There should be some sort of optimization here that will enable front office to continue along on their way without having to worry about searching for their errors. This could be done with a computer program that discusses the differences between the two. This would enable a quick view of the differences and a person could easily make the appropriate changes.

Investigation of the actual reconciliation process is underway. This is requiring a response by Dilruk for additional information on what actually goes on during the reconciliation process. A wealth of information has been unleashed by being able to access the database that contains a lot of the documents for Bank of America projects. This has been identified as business sensitive information and is not to be leaked out to other parties. Information needs to come from Dilruk. This information is mission critical for establishing the plan for optimization, automation, and solutions to current troubles.

Currently, eBlotter seems to put out an acceptable level of work. The errors coming from eBlotter are kept to a minimum because they effectively have three checks for the information to be correct. A broker will put in a ticket and eBlotter will translate this ticket to a usable format. Then a trade assistant has to put in their ticket and it has to match the original. This is a part of the check and balance system seen in the trades. Any issue with eBlotter is seen as a lesser issue because of the check system.

System write ups are currently being reviewed and edited. They will be soon shown to Abhi for further verification that the understanding of each system's importance is highlighted. Understanding of what is going on between the systems has been pretty well defined. Areas of concern overlap between middle office, the trading floor, and technology. The reconciliation process needs optimization, automation, and extra solutions for what happens if things go wrong. A suggestion is that there is a program that compares the prediction of the trader's office and the

Actual P&L. This will help the people who compute the Actual P&L as they won't have to go to the trading floor to ask why the prediction varies from the Actual P&L.

Direction:

The direction that this project is going is to focus in on the reconciliation process. This will eventually lead to a step-by-step plan for the technology department here at Bank of America to implement. This step-by-step plan will demonstrate understanding of the steps that go into the Flash P&L and the Actual P&L.

Current efforts are focused in understanding the reconciliation process while still pulling out information from sources. After Dilruk's information is received, all of it will be compiled into one format. Partial unwinds will be investigated further in order to aid in the search for a solution. This may be outside of the realm of do-ability as this seems to be a code level problem and understanding how to get more information from eBlotter. However, perhaps because the unwind rate is so low throughout the day, another solution could come of this. This solution could be an unwind client that one could run. This client would be based around marking an unwind. This would tell the Flash P&L to treat the partial unwind differently. Also, there could be a flag inside of eBlotter that tells the Flash P&L that it is a partial unwind. This aspect will be further elaborated on.

The goal is to establish a firm understanding of the processes that are going on between the trader's floor and middle office. This will be a main focus since this will allow us to investigate the reconciliation process and figure out what it takes to get a computer to do the same tasks.

Side Project Summary

Title: Risk Engine Architecture Review

Name: Thomas MacDonald

Task: A self-assigned project to investigate the Risk Engine server that computes the P&L, Risk, and General Ledger information for Bank of America. This project's goal is to be able to provide a plan for optimization to reduce the time jobs take to compute. Also included in this goal are to do basic computer engineering optimizations. This is also comprised of proposing new more powerful hardware, cost estimates, pros of the switch, cons of the switch, power efficiency, and efficient use of space.

Contact List:

Refined

Abhi Vashishth – Mentor, Manager, General Overview

Priscilla Wong – The lead for the Flash P&L development team for technology
Kate Schofield – Has worked with risk engine and is knowledgeable in the field.
Pete Spencer – Is working on revolutionizing P&L
Provost ad Interim John Orr – Electrical & Computer Engineering project advisor (WPI)

Approach:

Defined November 6th, 2007

1. Define the hardware used (5%)
    -Actual hardware specifications
2. Background research in computer engineering topics (35%)
    -Develop an understanding of the Cell processor that is practically a mini supercomputer
    -Understand other multi-core processors in order to give a good comparison
    -Floating point arithmetic should be looked into to understand the optimization process.
3. Analyze the current hardware (0%)
4. Create options for upgrades (0%)
    -Software optimization
    -Hardware upgrade
5. Estimate cost of switching to a hardware upgrade (0%)
6. Create plan for switching (0%)
7. Create hardware proposal (0%)
8. Celebrate. (0%)

Progress made:

Initial understand of the risk engine has been established. The risk engine is a server that accepts jobs from a user. The user specifies what type of job they want from the engine. In the case of the P&L, the user will specify that they want the Actual P&L. The risk engine will then pull in the information it needs to compute the Actual P&L. An estimate from Abhi places the risk engine at approximately 600 processors. This lends itself to an older technology that could be revolutionized by current technology options.

Research on the Cell processor is underway. Locating information on financial applications of the Cell processor has been relatively fruitful. The Cell processor would put Bank of America ahead of the curve since each of these tiny processors is like having a mini supercomputer. This would reduce the amount of cabinets necessary to do the same job as the risk engines current server. The Cell is designed to crunch numbers and is effectively designed to do the same amount of work as a couple of computers. Toshiba has tested the Cell by successfully running 48 MPEG-2 streams through it.

More information needs to be collected on the risk engine's hardware before an in-depth analysis can occur. Once the hardware is established; the power requirements, processing power, space requirement, and other vital aspects will be able to be compared to a smaller solution. This solution would include multi-core processors or the Cell processor. This would require multiple

computing racks; however this could yield some incredible results. The results of using the Cell or multi-core processors would enable the Actual P&L to be computed well before 8am so that the traders can have the Actual P&L before they even truly begin the day. This would be achieved by adding additional computing power to deal with the rigorous work that the Risk Engine has to undergo.

Direction:

A major focus is on getting the actual hardware specifications while looking for general algorithms that are executed for the calculation of P&L. These algorithms would only need to be a general look at what is going inside of the Risk Engine. This would enable us to avoid looking at the physical hard code and to only look at the higher level of mathematics that dictates how the risk engine operates.

Once the hardware has been established, this will lead to an in depth look at the power consumption, appropriate computing capabilities, and the space required for this computing apparatus. Background information of other hardware structures will need to be researched in order to develop a strong understanding of the problem at hand. Documents of using computers to compute financials will also need to be established for the long run.

A proposal for upgrade and optimization will be the ultimate goal of this project. If time permits, a test plan will also be issued. This test plan will be a way of measuring the feasibility of the new risk engine, the cost of the new risk engine, and the benefit of upgrading. This side project is not to take over the main project in any way and is solely for the benefit of the intern who has decided to undertake this task.

Appendix C – Proposed Solutions Document for Bank of America

Last Update: 29 November 2007

Main Project Status Report

Title: Profit and Loss Workflow Analysis
Name: Angelo Chandler
Name: Thomas MacDonald

Task:

To perform a gap analysis of the profit and loss workflow that currently occurs at Bank of America in London. This is to prepare a plan for optimization and how to condense the workflow so that the Actual P&L will be available at an earlier time than it currently is. The goal time is half past eight AM.

Contact List:

Abhishek Vashishth – Mentor, Manager, General Overview
Alex May – In change of the Flash P&L team in Middle Office
Dilruk Jayamanna – In charge of the Actual P&L team in Middle Office
Ben Cooper – Works on the trading floor, has solved certain P&L problems, extremely knowledgeable and helpful in almost all areas.
Priscilla Wong – The lead for the Flash P&L development team for technology
Kate Schofield – Has worked with Risk Engine and is knowledgeable in the field
Pete Spencer – Is working on revolutionizing P&L and has a lot of contact information
Phil Mitton – Expert on the East Bank Process

Approach:

Defined: Friday November 2nd, 2007

1. Collect colleague information (99%)
    - Set up a list of people to contact for certain information
    -Collect the problems (Still waiting for Dilruk)
    -Collect possible solutions to the problems (Still waiting for Dilruk)
2. System Level Analysis (95%)
    -Systems have been identified
    -Connections between systems have been identified
    -Interfaces have been identified
    -Clearer system level diagrams need to be created
3. Brief Description of Systems Involved (85%)
    -Preliminary system descriptions have been written
    -Review by Angelo and Thomas is imminent
    -Flash and Actual systems have been identified

-Risk Engine software manual has been identified
-A better understanding of how things are actually fed information

4. Look at all processes
-Manual (55%)
-Technical (95%)
-Who? (95%)
-What? (95%)
-Where? (95%)
-When? (95%)
-How? (70%)

5. Develop a strategy (75%)
-Fundamental understanding of events has been established
-A grip of what is going on has been established
-In-depth look at the reconciliation process is required

6. Proposition for Improvement (5%)
-Outline Areas for Improvement
-Research improvement areas
-Develop improvement plans and methods for each problem

7. Celebrate (0%)
-Celebration begins once the project has reached acceptable range

Flash P&L Progress

Main Drivers: Priscilla Wong and Alex May

Major Concerns:

- Partial Unwinds
- Cancelled / Removed Trades
- Duplication of new trades / NewBus tab
- Full Unwinds with an unsettled Fee
- Trades feeding the NewBus P&L Tab
- Timeliness of Running the Flash / System Performance
- Adjustment Tool requested
- Credit Spread Moves

Summary:

A partial unwind work around has been established by Alex May. The work around involves running the Flash P&L through the Risk Engine. This takes additional time and defeats the purpose of having a Flash. A partial solution has been designed for the Flash P&L. This solution will correctly add up the partial unwind fees for the day if the last action performed is a partial unwind. If the last action is not a partial unwind then Flash in unable to identify that there were any partial unwinds done on that day.

Flash is a program that runs on a desktop. Unfortunately, when this program is run on a computer it will use up most of the available computing resources and renders the computer unusable. A way of getting around this is to map smaller strategies but to increase the number of overall jobs. A bottleneck is offered when the computer has to pull the appropriate curves from Camden. There is interest in making the Flash a server-side program. Priscilla has mentioned a desire to even make the Flash a real-time application. However, Ben has noted that until the end of the day the Flash wouldn't change except for the new business. This makes sense as the profit and loss is established when the market closes for the day. New business can be affected by assignments and partial unwinds through the day.

An area for improvement is having the Flash report the three major categories of P&L kept as distinct areas while the others should be grouped together. The three major categories are credit spread, edit, and new business. Ben has suggested that all of the others besides these three be grouped together. Priscilla has also noticed that those three do seem to be the major players.

A .Net program written in C# is being developed to make viewing the Flash easier and more user friendly. Priscilla is currently looking for opinions on what should and shouldn't be enclosed inside of the program. This program would make dealing with the Flash more streamlined by offering the user capabilities of customizing what they would like to see.

A .Net program is also under development as an adjustment tool for fixing problems with the Flash. This comes from a request made by Alex May because there are times when he needs to change things manually and would like to be able to do it without haven't to bother the Flash development team. The .Net program appears to be a good step towards making the overall Flash process better.

Cancelled / removed trades are being addressed in the November 23rd release of the Flash. This makes Flash more robust by being able to ignore cancelled and removed trades that would not run in the Risk Engine. The trades coming through with the "Blot:####" ID that appear in the NewBus P&L are also being removed as these will cause errors in the Flash. Trades that are not uploaded to IRP/ADV should not appear in Flash. Trades that are assigned the wrong CDS curve in Flash will be getting an enhancement where they pull additional fields from eBlotter in order to identify the proper curve.

Proposed Solution for Flash P&L Problems:

.Net Flash Viewer

Description:

In development currently is a .Net Flash viewer. A function that should be present in this .Net Flash viewer should be the ability to be able to group P&L tabs together. This would

effectively sum up the P&L tabs that are being grouped together. The functionality would also have to be included to take the groups apart.

Reason:

Ben has requested the ability to group smaller P&L categories together. This is based on a desire to be able to quickly look at the areas that have the largest changes. This functionality would allow the traders to have a flexible tool that they can easily customize to their preferences. This would also increase the overall productivity of the team by enabling them to quickly see what they deem as important.

Server Based Flash

Description:

The server based Flash would be a server that has been specifically designed to run the Flash program. This would take the stress off of the computer that the program was originally run on while offering a great service. This service would include a standard place for finding Flash, being able to calculate Flash more efficiently, and to do slightly more complex calculations.

One aspect that should be addressed by this server is that it should be able to run certain tasks as if it was being run through Risk Engine. These specific tasks are the global new business, London-based P&L, and the unwind P&L. These tasks are currently run through Risk Engine to provide a Flash figure for these. If Flash had this built into it, it would enable the user to only have to go to one application. This one application would then gain significant value because it is able to compute everything that is necessary of the user.

This server would also be required to have enough computing power to keep the Flash very timely. Currently, when there are fixes to the Flash, the Flash will have to be run again after each fix. This takes time and should be addressed by the server option. The server should also focus on energy concerns and try to offer significant computing capabilities for cheap.

Normal Option for the Server:

Depending on the necessity of computational power the quad-core Xeon 5365 is capable of performing 24 GFLOPs for single and double precision calculations. This puts it towards the top of its class and would serve as an excellent basis for a server. Hewlett-Packard offers a server that is designed for this sort of work load. The server is called the HP ProLiant ML370 G5. It sports two Quad-Core Intel Xeon X5365 processors. Each of them is clocked at 3.0 GHz. However, there is a lot of flexibility with the processor choices. The Xeon 5400 series that is offered has an L2 cache of 12 MB instead of the Xeon 5300 series that has an L2 cache of 8 MB. The ram can be 1 GB, 2 GB, or 4 GB.

This is a 5U rack mount or it can be kept as a tower. This server would more than likely be adequate depending on how complex the Flash becomes. With the rack mountable version, more servers could be added to balance off the workload. This offers expandability but 5U is a significant amount of rack space when planning out a server cluster.

Unique Options for the Server:

If there is rack space that is available then Mercury Systems has a 1U Dual Cell-Based rack mounted server that provides 410 GFLOPs of computing power while only consuming 575W to 700W. The servers use a Linux-variant for the operating system, which has been known to be very stable.

If there is not rack space available then Mercury Systems offers the PowerBlock 200 System. This only provides 200 GFLOPs of performance. This performance factor will enable the Flash to have more processing power with less room being expended. It is however a rugged stand alone system. These Cell based options would require adaptation of the code, so that it is optimized for the Cell.

Reason:

The first reason is that the Flash process would be increased by not bogging down a workstation. This presents itself as a very desirable option while maintaining a cost effective nature. This will also be time effective as the Flash will then run much like the Risk Engine where a GUI sends a request for jobs. This would be ideal. Also, a centralization of how things are computed is extremely beneficial. Currently, Alex May's team is going through a lot of processes that should be encompassed in one program. It's very taxing on them as it requires a lot from people who are already hard-pressed for time.

Actual / Explain P&L Progress

Main Drivers: Ben Cooper and Dilruk Jayamanna

Figures Provided Below: Risk Engine System Diagram, Risk Engine Process Chart, Risk Engine Database Feeds

Major Concerns:

- Unwinds
- Assignments
- Middle Office and Front Office Reconciliation

Summary:

Unwinds and assignments can cause credit spreads to be reported incorrectly. Odin stores pure credit spread information and serves as the prediction for the Actual credit spread P&L.

Assignments will affect the New Business P&L and the Edit P&L by reporting some of the P&L inside of New Business P&L and some of it inside of Edit P&L. Assignments should rest solely inside of the New Business P&L. Edit P&L is for actual corrections in the P&L and should be reserved for such purposes. Unwinds will cause some of the P&L to go into credit spread while the other portion will be incorrectly reported in the edit P&L.

Currently, eBlotter doesn't exhibit many problems as there are three checks to make sure that the information is correct. A broker will put in a ticket and eBlotter will translate this ticket to a usable format. Then a trade assistant has to put in their ticket and it has to match the original. This is a vital part in the check and balance system seen in the trades.

Ben has also requested some sort of tool or process that makes dealing with the New Business P&L at of the end of the day quicker. Currently there is a manual check done by trader's assistants. This manual check should be automated because that would take the human element out of the equation. This would make things faster and would improve the overall accuracy of the New Business P&L. This is beneficial to increasing the accuracy of the overall New Business P&L.

Flash's purpose has been established quite clearly. Odin only does the first level deltas while Flash will do the second level derivative and give the convexity. This gives Flash the edge and is why it is compared with the Actual. Odin is also not aware of partial unwinds or assignments.

Proposed Solutions for Actual / Explain P&L Problems:

Odin Upgrade

Description:

Odin should receive an upgrade that would make it more like Flash and the Actual P&L computing devices. This upgrade would include functionality improvements by enabling Odin to deal with partial unwinds, full unwinds, and assignments. The upgrade would also include the ability to deal with convexity. Usually, the credit moves are small but when they are not, this has the ability to cause a large shift in P&L.

Reason:

This upgrade is necessary because the traders should be able to accurately predict their P&L. This would help the reconciliation process as the numbers would be much closer to the Flash and Actual P&L. Odin's job is to help the traders. Since traders cannot accurately predict their P&L, this effectively is a waste everyone's time. They should have the tools to be able to accurately predict their P&L. This would enable an overall efficiency boost because Odin would become a third source to build faith in the P&L system.

This would also increase faith in the Flash P&L. This is because Odin's marks would be pushed through a similar system. Odin would then enable the trader's to notice unwinds and assignments and to be able to adjust accordingly. By being able to see this, the traders will be able to look at the Flash P&L and see how far their marks from Odin are. The marks should much closer with these upgrades and will allow for a quick comparison.

East Bank Process Upgrade:

Description:

Currently, the East Bank Process is performed on two desktop computers that rest inside of middle office. This upgrade would move the East Bank Process to an SQL based server option. This SQL based server option would also be able to process the following two upgrade proposals. It would be optimal if the server allowed the user to run ad hoc queries. Security could be handled through Windows authentication processes. Also it would have to have intranet capabilities so that reports can be quickly grabbed off of the server.

Reason:

This upgrade would store the additional two program proposals. It would offer an increased computing apparatus to support the additional programs while offering a robust system for middle office. The East Bank process was also supposed to be a temporary fix. Its age is starting to show as there are now two computers that do the East Bank calculations. Condensing this to one robust computing apparatus will save power, offer more options, and will serve as a better option over time.

Flash & Actual Comparison Program

Description:

This program would be run after the East Bank Process. It would effectively take in the Flash P&L data and the Actual P&L data and compare them. This comparison would highlight the trades / traders that have discrepancies. These discrepancies would be highlighted in different colors due to their severity. A difference of 100,000 or greater would be red. From 99,999 to 25,000 would be yellow. 24,999 to 2,000 would be green. While 1,999 or less would not have a color associated with it. This would enable a reconciler to notice large differences immediately.

This highlighted file would not be optimal for making changes. That is why the original file that exits the East Bank Process would still be used for the actual reconciliation of numbers. The file that exits this comparison program would be delivered to front office and middle office.

This comparison program could also be used to highlight the differences between the non-reconciled and reconciled Actual P&L at the end of the day and distribute a file to the traders. This would let them see what changed since the initial report.

Reason:

This would effectively locate the errors immediately. It would make the reconciliation process much quicker as the numbers would be clearly identified as an error. Front office would also be given a file that they could use in combination with the Flash. They could quickly see the differences and be able to anticipate what the reconciled version will look like. This will effectively give front office a tool that enables them to be able to anticipate more accurately what really happened the day before. Middle office would still get the original file from the East Bank Process and the highlighted version. They would then go about the reconciliation process as they normally would.

The second functionality would enable traders to be ready and able to adjust their trades accordingly without having to wade through a lot of information. Vital differences would be highlighted so that the trader wouldn't have to wade through too much information. The entire report would have to be streamlined and customizable for the trader to be able to wade through the document with little effort. This would maximize the value of the overall report as the trader's time would be spent glancing at the report and focusing more on their trades.

Trigger Based Data Pull Program

Description:

This program would ultimately be fed by the Flash & Actual comparison program. It would take the areas that have discrepancies and pull up the according information to solve each problem. The reconciliation process would not be about finding the data or finding the discrepancies anymore as when this program is used in conjunction with the Flash & Actual comparison program all of that would be at the reconciler's fingertips. This would effectively pool the systems that act accordingly.

This program would basically take each of the discrepancy and add a button with the information of how the Flash computed the figure and how the Actual computed the figure. It would also give the original information in a user friendly display interface. This user friendly display interface would be based around giving an intuitive feeling of how to pull up information. Design of this interface would have to come from surveys of middle office. The most important information would have to be easily accessible with either a single click or a hot-key.

This program would take their information from the systems that report the necessary data such as Odin, eBlotter, Camden, Advantage, IRP, Loantrak, Alice, and Bloomberg. It would also have to keep track of the errors. This functionality could be increased by being able to generate weekly or monthly reports on where most of the errors come from. This would highlight areas of concern.

Reason:

This would allow middle office to have immediate access to the reconciliation information that they need from the associated systems. This would cut down the time searching through the files in order to figure out where the error occurred. This program could also store errors to develop charts and graphs that can be used for optimization purposes later. This would help locate common problems that the programmers would be able to compensate for the errors or fix their algorithms. This would further help the reconciliation process by enabling feedback on where discrepancies occur.

Process change for the P&L Workflow

Description:

The process change would involve employing the Flash & Actual comparison program. The first point that is important to note is that front office would be getting the highlighted file from the Flash & Actual comparison program. This would allow them to have enough information to make informed decisions about their trades by providing a reasonable range of information with notes on large discrepancies. At this time, the trader's would be able to also use Odin to compare any very large discrepancies. This gives three valuable sources of information.

Middle office would still get the East Bank Process's condensed excel file. They would, in addition, also receive the output of the Flash and Actual comparison program. This would enable middle office to see the differences clearly and would speed up the reconciliation process by identifying the problems. This should ultimately speed up the entire process. If this is used in conjunction with the trigger based data pull program it should cut the overall process time dramatically by offering most of the information necessary to reconcile the differences.

Reason:

The reconciliation process has been identified as the portion of the P&L workflow that requires the most time. Any additional tools to be able to resolve problems and reconcile the overall product faster would make it so that traders' can actually see what is happening with their investments. This would hopefully allow middle office to cut a few hours off the reconciliation process. Middle office's purpose for the P&L reconciliation is to serve as a check point and by giving front office the output of the Flash & Actual comparison program wouldn't void this. It would be giving front office a way of being able to deal with their trades and not depend solely on the Actual.

This does not bypass middle office as middle office would still be expected to provide the reconciled P&L in a timely fashion. The traders would understand that they do not have the full story of the P&L until middle office provided the reconciled version of the Actual. However, they would still have enough information to make educated choices.

Major Concern:

It has to be made absolutely clear to front office that the output of the Flash & Actual comparison program is a comparison between a non-reconciled Actual. A way to avoid this would to be name the Actual that is being compared something like “working” P&L or even non-reconciled P&L.

Data Logging Program

Description:

With the implementation of the comparison program, a data logging feature could be added. This data logging program would track all of the differences and pull up the reasons for the discrepancies. This program would have to be able to produce charts, graphs, standard deviations, and be able to monitor multiple sources. This would ultimately include a generalized report, that could be requested either daily, weekly, or monthly. This program could ultimately log the data for the non-reconciled Actual, the Flash, and the reconciled Actual P&L.

Reason:

With this additional feature, the Flash team can have data that details discrepancies and why the discrepancies exist. This’ll allow them to see visually where improvements can be made and how to approach the situation. Critical zones of improvement should become clear as they will show consistent errors. This feature shouldn’t be too hard to implement.

New Viewer for P&L

Description:

Instead of using Excel, the new Flash viewer could be used as a way of seeing the reconciled Actual in combination to being the new Flash viewer. This would allow for greater flexibility as a trader could easily shift in between the Flash and the reconciled Actual. It would also allow them to group similar tabs that they see fit. This could be used in conjunction with the comparison program and allow a trader to instantly see differences between the Flash, non-reconciled Actual, and the reconciled Actual.

Reason:

This would be using a resource that is already being developed and expanding on its data capabilities. This would ultimately allow traders to be able to access their information in a much timelier manner. It would also enjoy the benefits of being a reused piece of software that could have multiple uses instead of being specifically Flash devoted. This would add additional value to the tool.

Risk Engine Software Upgrade

Description:

Currently the Risk Engine reports partial unwinds, full unwinds, and assignments incorrectly. This upgrade would be a focus on being able to put each of these in the proper sections and not allow them to spread out like they currently do. Partial unwinds and full unwinds will cause a portion of the P&L to be reported in credit spread and edit P&L. This should be completely reported inside of the credit spread P&L section. Assignments currently report a portion of their information inside of the new business and edit P&L. This should be strictly reported inside of the new business P&L and not inside of edit P&L.

Reason:

This small upgrade would save a lot of trouble inside of the reconciliation process. It would address some of the key problems that arise during the reconciliation process. This would also eliminate one of the largest causes of error in the current report system

Overall benefit of all of these solutions

By combining all of these solutions, the tools are given to the traders, front office, and middle off that enable them to do their jobs in a more precise and timely fashion. This'll increase transparency inside the company by allowing everyone to understand what is happening during the entire P&L process. It will allow traders and front office to be able to have more reliable data and to be able to make more informed decisions.

Middle office's reconciliation job will become easier as they won't have to go to front office to look up errors or discrepancies immediately. This will allow front office to continue through the day and to be able to get more work done as they won't be inquired upon as much. Ultimately, technology will be helping front office and middle office by offering a better range of services.

The value of the overall plan will be greater than the sum of its parts because all of these parts benefit from the others. The beneficial relationship between these solutions is linked by the reliance on one another. This can be seen in the East Bank Process upgrade that includes the Flash & Actual comparison program and the trigged based data pull program. The relationship here shows that the East Bank Process's upgrade to a robust server would also provide the secondary function of being able to run the other two programs that'll help out middle office and the traders.

Authorship Page

Section	Main Author	Contributing Author
Executive Summary	Thomas MacDonald	
Introduction - Financial	Thomas MacDonald	
Introduction - Risk Engine Review	Thomas MacDonald	
Introduction - Econophysics	Angelo Chandler	
Financial Background	Thomas MacDonald	
Risk Engine Review Background	Thomas MacDonald	
Econophysics Background	Angelo Chandler	
General Approach Methodology	Thomas MacDonald	
Systems Engineering Methodology	Angelo Chandler	Thomas MacDonald
Flash P&L Methodology	Angelo Chandler	Thomas MacDonald
Econophysics Methodology	Angelo Chandler	
Risk Engine Review Methodology	Thomas MacDonald	
Conclusion P&L Worfklow Gap Analysis	Thomas MacDonald	
Conclusion Risk Engine Review	Thomas MacDonald	
Conclusion Econophysics	Angelo Chandler	
Recommendations and Future Work	Thomas MacDonald	Angelo Chandler
Appendix A - Terminology	Thomas MacDonald	
Appendix B - Midway Report to BofA	Thomas MacDonald	Angelo Chandler
Appendix C - Proposed Solutions for BofA	Thomas MacDonald	Angelo Chandler

Works Cited

"Arbitrage." Investopedia. 30 Oct. 2007 <http://www.investopedia.com/terms/a/arbitrage.asp>.

"Average Retail Price of Electricity, March 2007." Public Policy Institute of New York State, Inc. 27 Nov. 2007 <http://www.ppinys.org/reports/jtf/electricprices.html>.

Barney, Blaise. "Introduction to Parallel Computing." Lawrence Livermore National Laboratory. 22 June 2007. 22 Nov. 2007 <http://www.llnl.gov/computing/tutorials/parallel_comp/>.

Barney, Blaise. "Message Passing Overview." MHPCC. 3 July 1996. University of Hong Kong. 8 Nov. 2007 <http://www.hku.hk/cc/sp2/workshop/html/message_passing/message_passing.html>.

Bicer, Murat, Jonas Larsson, and Matt Sexton. "European Option Pricing with Monte Carlo Simulation on the Cell BE Processor." Mercury Computer Systems - White Paper.

Blachford, Nicholas. "Cell Archiecture Explained Version 2 - Part 1: Inside the Cell." Blachford. 2005. 21 Nov. 2007 <http://www.blachford.info/computer/Cell/Cell1_v2.html>.

Bock, Rudolf K. "Cholesky Decomposition." 7 Apr. 1998. 1 Nov. 2007 <http://rkb.home.cern.ch/rkb/AN16pp/node33.html#SECTION00033000000000000000>.

"Bond Ratings." Money-Zine. 24 Oct. 2007 <http://www.money-zine.com/Investing/Investing/Bond-Ratings/>.

Buede, Dennis M. Engineering Design of Systems - Models and Methods. (pp. 281). John Wiley & Sons. Online version available at: http://www.knovel.com/knovel2/Toc.jsp?BookID=1410&VerticalID=0

"Cell Architecture." University of North Texas. 21 Nov. 2007 <http://csrl.unt.edu/~kavi/CSCE6610/cell.ppt>.

Chacko, George, Anders Sjöman, Hideto Motohashi, and Vincent Dessan. Credit Derivatives. Upper Sandle River, New Jersey 07458: Wharton School, 2006. pp. 9-40. 227, 232-242

Chako, George, Peter Hecht, Vincent Dessain, and Monika Stachowiak. “Note on Bond Valuation and Returns.” Harvard Business School Note No. 205-008. Boston: Harvard Business School Publishing, 2004

Chaplin, Geoff. Credit Derivatives: Risk Management, Trading, & Investing. West Sussex, England: Wiley Finance, 2005. pp.3-4, 6, 143, 163-164

Cheng, Steve. "Black-Scholes PDE." Planet Math. 30 Oct. 2007 <http://planetmath.org/encyclopedia/BlackScholesPartialDifferentialEquation.html>.

Chen, Thomas, Ram Raghavan, Jason Dale, and Eiji Iwata. "Cell Broadband Engine Architecture and Its First Implementation." IBM. 29 Nov. 2005. 14 Nov. 2007 <http://www-128.ibm.com/developerworks/power/library/pa-cellperf/>.

Choudhry, Moorad. The Credit Default Swap Basis. New York: Bloomberg, 2006. pp. 2-8, 13-14, 37-45

"Collaterialized Debt Obligation." Risk Glossary. 29 Oct. 2007 <http://www.riskglossary.com/link/collateralized_debt_obligation.htm>.

"Credit Event." Creditflux. 25 Oct. 2007 <http://www.creditflux.com/glossary/credit+event.htm>.

Damodaran, Aswath. "Option Pricing Theory and Models." Damodaran Online (Investment Valuation: Second Edition). New York University Stern. 26 Oct. 2007 <http://pages.stern.nyu.edu/~adamodar/pdfiles/valn2ed/ch5.pdf>.

"Default Model." Riskglossary. 23 Oct. 2007 <http://www.riskglossary.com/link/default_model.htm>.

Dickson, John. "Rogue Trading?" Plus Magazine. Sept. 2001. 31 Oct. 2007 <http://plus.maths.org/issue16/features/derivatives/index.html>.

Dongarra, Jack J. "Performance of Various Computers Using Standard Linear Equations Software." Netlib. 20 Nov. 2007. University of Tennessee. 27 Nov. 2007 <http://www.netlib.org/benchmark/performance.pdf>.

Easton, John, Ingo Meents, Olaf Stephan, Horst Zisgen, and Sei Kato. "Porting Financial Markets Applications to the Cell Broadband Engine Architecture." IBM. June 2007. 20 Nov. 2007 <http://dl.alphaworks.ibm.com/technologies/cellsw/cellFMwhitepaper.pdf>.

Ellis, David. "Half-Life 2: the Orange Box - This Box Has Problems." 1up. 21 Nov. 2007. 2 Dec. 2007 <http://www.1up.com/do/previewPage?cId=3164539>.

Ercal, Fikrey. "2.2 Using Workstation Clusters." CSc 387. University of Missouri. 22 Nov. 2007 <http://web.umr.edu/~ercal/387/class.2-MPI.html>.

"Finance Corporate Debt Ratings Moody's and S&P." Blaha. 24 Oct. 2007 <http://www.blaha.net/Finance%20Corporate%20Debt%20Ratings.htm>.

Fink, Peter. "Pricing Complex Options Using a Simple Monte Carlo Simulation." John's Finance Page. Comp. John Janusz. 30 Oct. 2007 <http://members.allstream.net/~johnjaz/montemain.htm>.

Fisher, Susie. "Alice 4.0." Gmwiki. 2 Aug. 2007. Bank of America. 25 Nov. 2007 <http://gmwiki.bankofamerica.com/display/GSPBAT/Alice+4.0>.

"Folding@Home on the PS3." Standford University. 22 Nov. 2007 <http://folding.stanford.edu/English/FAQ-PS3>.

Foster, Ian. "1.3 a Parallel Programming Model." Argonne National Laboratory. 1995. 8 Nov. 2007 <http://www-unix.mcs.anl.gov/dbpp/text/node9.html>.

Gibbons, Phillip B. "What Good are Shared-Memory Models?" Proceedings of the 1996 ICPP Workshop on Challenges for Parallel Processing (1996): 103-114. IEEEXplore. George C. Gordon Library, London. 19 Nov. 2007. Keyword: shared memory model.

Goldberg, David. "What Every Computer Scientist Should Know About Floating-Point Arithmetic." Numerical Computation Guide. Mar. 1991. 12 Nov. 2007 <http://docs.sun.com/source/806-3568/ncg_goldberg.html>.

Gropp, William, Rusty Lusk, Debbie Swider, and Rajeev Thakur. "The Message Passing Model." Argonne National Laboratory. 22 Nov. 2007 <http://www-unix.mcs.anl.gov/mpi/tutorial/mpibasics/sld003.htm>.

Gschwind, Michael, H P. Hofstee, Brian Flachs, Martin Hopkins, Yukio Watanaba, and Takeshi Yamazaki. "Synergistic Processing in Cell's Multicore Architecture." Penn State. 20 Nov. 2007 <http://citeseer.ist.psu.edu/cache/papers/cs2/624/http:zSzzSzwww.research.ibm.comzSzpeoplezSzmzSzmikegzSzpaperszSz2006_ieeemicro.pdf/gschwind06synergistic.pdf>.

Heakal, Reem. "What is a Corporate Credit Rating?" Investopedia. 22 Oct. 2003. 23 Oct. 2007 <http://www.investopedia.com/articles/03/102203.asp>.

Hollasch, Steve. "IEEE Standard 754 - Floating Point Numbers." 24 Feb. 2005. 23 Nov. 2007 <http://steve.hollasch.net/cgindex/coding/ieeefloat.html>.

"HP ProLiant DL360 G4p 2.8GHz Intel Xeon." Bizrate. 27 Nov. 2007 <http://www.bizrate.com/desktopcomputers/hp-proliant-dl360-g4p-2.8-ghz-intel-xeon--pid469987842/compareprices.html>.

"HP Proliant DL360 G5 Entry." Cost Central. 27 Nov. 2007 <http://www.costcentral.com/proddetail/HP_ProLiant_DL360_G5_Entry/416559001/K23560/>.

"HP ProLiant DL360 Generation 4p (G4p)." Hewlett-Packard. 27 Nov. 2007 <http://h18004.www1.hp.com/products/quickspecs/12166_div/12166_div.pdf>.

"IEEE 754 Floating Point." Randelshofer Online. 23 Nov. 2007 <http://www.randelshofer.ch/fhw/gri/float.html>.

"Interest Rate Swaps Pricing." Derivative One. 29 Oct. 2007 <http://www.derivativesone.com/kb/interest_rate_swaps.aspx>.

"Introduction to Parallel Computing." Department of Computer Science and Information Systems. University of Limerick. 8 Nov. 2007 <http://www.csis.ul.ie/Modules/CS4838/tutorialsforweb/Introduction%20to%20Parallel%20Computing.htm>.

Investorguide Staff. "Stock Basics - Common and Preferred Stock." Investor Guide. 1 Nov. 2007 <http://www.investorguide.com/igu-article-818-stock-basics-common-and-preferred-stock.html>.

Jacobs, Jason. "Xeon 5150 Vs. E5320." Techware Labs. 29 Mar. 2007. 7 Nov. 2007 <http://www.techwarelabs.com/reviews/processors/xeoncomparison/>.

Janusz, John. "Derivative Pricing Theories: Welcome." John's Finance Page. 26 Oct. 2007 <http://members.allstream.net/~johnjaz/bsm.htm>.

Janusz, John. "Binomial Pricing Theory." John's Finance Page. 30 Oct. 2007 <http://members.allstream.net/~johnjaz/bt.htm>.

Janusz, John. "Monte Carlo Simulation." John's Finance Page. 30 Oct. 2007 <http://members.allstream.net/~johnjaz/mont.htm>.

Kahle, J A., M N. Day, H P. Hofstee, C R. Johns, T R. Maeurer, and D Shippy. "Introduction to the Cell Multiprocessor." IBM. 7 Sept. 2005. 16 Nov. 2007 <http://www.research.ibm.com/journal/rd/494/kahle.html>.

Keable, Crispin. "17th Machine Evaluation Workshop." STFC Computational Science and Engineering Department. 6 Dec. 2006. 21 Nov. 2007 <http://www.cse.scitech.ac.uk/disco/mew17/talks/Keable_IBM_MEW17.pdf>.

Kindratenko, Volodymyr. "High-Performance Computing on FPGAs." National Center for Computational Sciences. 7 Sept. 2007. Univesity of Illinois At Urbana-Champaign. 23 Nov. 2007 <http://www.ccs.ornl.gov/workshops/FallCreek07/presentations/kindratenko.pdf>.

Koch, Geoff. "Discovering Multi-Core: Extending the Benefits of Moore's Law." Intel. July 2005. Michigan State University. 19 Nov. 2007 <http://www.intel.com/technology/magazine/computing/multi-core-0705.pdf>.

Kothari, Vinod. "ISDA Credit Event Definitions." Credit Derivatives. 25 Oct. 2007 <http://www.credit-deriv.com/isdadefinitions.htm>.

Krowne, Aaron. "Cholesky Decomposition." Planet Math. 4 June 2006. 1 Nov. 2007 <http://planetmath.org/?op=getobj&from=objects&id=1287>.

Kurzak, Jakub. "New Approaches to Numerical Linear Algebra on the CELL Processor." Innovative Computing Laboratory. 25 Oct. 2006. University of Tennessee. 23 Nov. 2007 <http://www.cs.utk.edu/~dongarra/cell2006/cell-slides/03-Jakub-Kurzak.pdf>.

Kyrnin, Mark. "Dual-Core Processors." About. 7 Nov. 2007 <http://compreviews.about.com/od/cpus/a/dualcore.htm>.

"Long Call." Option Trading Tips. 2005. 31 Oct. 2007 <http://www.optiontradingtips.com/strategies/long-call-option.html>.

"Long Put Option." Option Trading Tips. 2005. 31 Oct. 2007 <http://www.optiontradingtips.com/strategies/long-put-option.html>.

"Long Put / Put Purchase / Buy Put Option." Trade Soft. 31 Oct. 2007 <http://www.trader-soft.com/option-trading/option-strategies/put-long.html>.

Mercury Systems. Telephone interview. 15 Dec. 2007.

Mutschler, Ann S. "Intel Creates 2nd 65nm Manufacturing Process." Electronic News. 20 Sept. 2005. 19 Nov. 2007 <http://www.edn.com/index.asp?layout=article&articleid=CA6258496>.

"NOVA Online | Trillion Dollar Bet | the Formula That Shook the World." PBS. Nov. 2000. 26 Oct. 2007 <http://www.pbs.org/wgbh/nova/stockmarket/formula.html>.

"NOVA Online | Trillion Dollar Bet | Site Map." PBS. 26 Oct. 2007 <http://www.pbs.org/wgbh/nova/stockmarket/formulaoptions.html#options>.

"NOVA | Transcripts | Trillion Dollar Bet." PBS. 8 Feb. 2000. 26 Oct. 2007 <http://www.pbs.org/wgbh/nova/transcripts/2704stockmarket.html>.

"Off-Balance-Sheet Financing." Investopedia. 29 Oct. 2007 <http://www.investopedia.com/terms/o/obsf.asp>.

Ogrinz, Michael, and Kathiresan Solaiyappan. "Camden Home." Bank of America.Com. 8 Nov. 2007. Bank of America. 22 Nov. 2007 <http://gmwiki.bankofamerica.com/display/camden/Camden+Home>.

O'Sullivan, Bryan. "What is a Thread?" Teideal Glic DeisbhéAlach. 9 June 2005. 21 Nov. 2007 <http://www.serpentine.com/blog/threads-faq/what-is-a-thread/>.

Page, M. "IBM's CELL Processor: Preview to Greatness?" PCStats. 15 May 2005. 14 Nov. 2007 <http://www.pcstats.com/articleview.cfm?articleid=1727&page=3>.

Parlett, Beresford N. The Symmetric Eigenvalue Problem. SIAM, 1998. Google Book Search. 13 Nov. 2007 <http://books.google.com/books?id=uaYXlkHVPpEC&pg=PA31&lpg=PA31&dq=guard+digit&source=web&ots=06PkdJ-EeD&sig=dpRosxIj1U0HitNkMgl6sJ0gae0#PPA31,M1>.

"Pentium 4 Xeon Cluster 2.2/2.4/2.8 GHz." Top500 Supercomputing Sites. 22 Nov. 2007 <http://top500.org/system/6159>.

"Planning Considerations for Multicore Processor Technology." Dell. May 2005. 7 Nov. 2007 <http://www.dell.com/downloads/global/power/ps2q05-20050103-Fruehe.pdf>.

"ProLiant DL300 Series." Hewlett Packard. 27 Nov. 2007 <http://h10010.www1.hp.com/wwpc/ca/en/sm/WF04a/12146350-12146352-12146476-12146476-12146498.html>.

Reporting Group, Risk Engine. "Risk Engine User Guide." Risk Engine. 3 Oct. 2007. Bank of America. 16 Nov. 2007 <http://crpscsgap105/RiskEngine/servlet/FileRetriever?filename=%5C%5CCRPNYCDEV22%5Cp4store%5Cvob_adv_riskengine_docs%5CUser+Guides%5CRisk+Engine-User+Guide.pdf>.

Rubash, Kevin. "Black - Scholes -- Option Pricing Models." Bradley University. 26 Oct. 2007 <http://bradley.bradley.edu/~arr/bsm/pg04.html>.

Rubinstein, Mark. "Chapter 4. Binomial Option Pricing Model." In the Money. 1998. 26 Oct. 2007 <http://www.in-the-money.com/presentation/sld063.htm>.

Schurr, Stephen. "Bond Ratings." TheStreet. 16 Aug. 2000. 25 Oct. 2007 <http://www.thestreet.com/basics/gettingstarted/1036369.html>.

Scott, Michael L. "Shared Memory Models." Apr. 2002. University of Rochester. 8 Nov. 2007 <http://www.cs.rochester.edu/research/cashmere/2002-04_pmodels_talk.pdf>.

"SGI Altix ICE 8200, Xeon Quad Core 3.0 GHz." Top500 Supercomputing Sites. 22 Nov. 2007 <http://www.top500.org/system/8966>.

Shamlin, David. "Threads Unraveled: a Parallel Processing Primer." SAS Institute Inc. 22 Nov. 2007 <http://www2.sas.com/proceedings/sugi29/217-29.pdf>.

Shimpi, Amand L. "Cool Runnings: the Intel Bobsled Team." AnandTech. 30 July 2001. 29 Nov. 2007 <http://www.anandtech.com/showdoc.aspx?i=1509&p=3>.

Silverman, Dennis. "Solution of the Black-Scholes Equation Using the Green's Function of the Diffusion Equation." 20 May 1999. University of California, Irvine. 1 Nov. 2007 <http://www.physics.uci.edu/~silverma/bseqn/bs.pdf>.

"Sony Global - Press Release - IBM, Sony, Sony Computer Entertainment Inc. and Toshiba Unveil Cell Processor." Sony. 29 Nov. 2004. 14 Nov. 2007 <http://www.sony.net/SonyInfo/News/Press_Archive/200411/04-1129E/>.

"Sony/Toshiba/IBM (STI) CELL Processor." Spring 2007. University of Tennessee. 23 Nov. 2007 <http://www.cs.utk.edu/~dongarra/WEB-PAGES/SPRING-2007/Lect04b.pdf>.

Stokes, Jon. "Introducing the IBM/Sony/Toshiba Cell Processor - Part I: the SIMD Processing Units." Ars Technica. 7 Feb. 2005. 14 Nov. 2007 <http://arstechnica.com/articles/paedia/cpu/cell-1.ars>.

Sun, C. "Intel Pentium 4 3.0 GHz 800MHZ FSB Processor Review." PCStats. 11 June 2003. 22 Nov. 2007 <http://www.pcstats.com/articleview.cfm?articleid=1403&page=4>.

"The JP Morgan Guide to Credit Derivatives." Investing in Bonds. 24 Oct. 2007 <http://www.investinginbonds.com/assets/files/Intro_to_Credit_Derivatives.pdf>. p. 44

Team Latte. "All About the Cholesky Matrix." Risk Latte. 9 Aug. 2005. 1 Nov. 2007 <http://www.risklatte.com/portfolioEngineering/article20050809.php>.

"The Linpack Benchmark." Top 500 Supercomputing Sites. 27 Nov. 2007 <http://www.top500.org/project/linpack>.

"The Message Passing Model." Hewlitt-Packard. Dec. 2003. 22 Nov. 2007 <http://docs.hp.com/en/B6060-96013/ch01s01.html>.

"Theta." Risk Glossary. 26 Oct. 2007 <http://www.riskglossary.com/link/theta.htm>.

"Total Return." Investopedia. 29 Oct. 2007 <http://www.investopedia.com/terms/t/totalreturn.asp>.

"Tranches." Investopedia. 29 Oct. 2007 <http://www.investopedia.com/terms/t/tranches.asp>.

"Types of Bonds - Education Center." Yahoo! 23 Oct. 2007 <http://finance.yahoo.com/education/bond/article/101197/Types_of_Bonds>.

"Ubisoft Acknowledges Assassin's Creed Freezes on PS3." Gamekiq. 2 Dec. 2007 <http://www.gamekiq.com/wordpress/2007/11/18/ubisoft-acknowledges-assassins-creed-freezes-on-ps3/>.

"Unwind Financial Defintion of Unwind." Free Online Dictionary. 25 Oct. 2007 <http://financial-dictionary.thefreedictionary.com/unwind>.

Vashishth, Abhishek. Personal interview. 30 Oct. 2007

"Vega." Risk Glossary. 26 Oct. 2007 <http://www.riskglossary.com/link/vega.htm>.

"What is a Bond? - Education Center." Yahoo! 23 Oct. 2007 <
http://finance.yahoo.com/education/bond/article/101194/What_is_a_Bond>.

Williams, Samuel, John Shalf, Leonid Oliker, Shoaib Kamil, Parry Husbands, and Katherine Yelick. "Scientific Computing Kernels on the Cell Processor." 2007. Berkeley. 22 Nov. 2007 <http://bebop.cs.berkeley.edu/pubs/williams2007-scicomp-cell.pdf>.

Williams, Samuel, John Shalf, Leonid Oliker, Parry Husbands, Shoaib Kamil, and Katherine Yelick. "The Potential of the Cell Processor for Scientific Computing." Lawrence Berkley National Laboratory. 2005. University of California. 22 Nov. 2007 <http://repositories.cdlib.org/cgi/viewcontent.cgi?article=4262&context=lbnl>.

Wystup, Uwe. "Vanilla Options - Chapter 1." Math Finance. 26 Oct. 2007 <http://www.mathfinance.de/FXRiskBook/chap-1.pdf>.

Yongzeng, and Jerome Spaingler. "Applications of Monte Carlo/Quasi-Monte Carlo Methods in Finance: Option Pricing." Smart Quant. 30 Oct. 2007 <http://www.smartquant.com/references/MonteCarlo/mc6.pdf>.