

# Análise e Modelagem de Desempenho de Sistemas de Computação

Profa. Jussara M. Almeida  
1º Semestre de 2011

# Problemas, problemas, problemas

- Um servidor de banco de dados tem uma CPU e 10 discos. Transações para execução no banco de dados chegam a uma taxa de 2 transações por segundo (2 tps). Durante a sua execução, uma transação alterna processamento com acesso a um dos discos, possivelmente mais de uma vez.

Qual o tempo de resposta médio de uma transação? Qual a taxa de processamento (throughput) do servidor?

O que precisa ser feito???

- Medição em sistema real: se ele estiver em operação (pode ser muito intrusivo, necessita ferramental)
- Simulação: precisa de tanto? Com que nível de detalhe?
- Modelagem analítica: precisa de inputs - quais?

# Problemas, problemas, problemas

- Suponha que uma investigação mais detalhada de um log de um servidor de dados revela que transações individuais submetidas ao servidor têm características diferentes, por exemplo:
  - Tempos de processamento muito heterogêneos
    - Consultas triviais X médias X complexas
  - Categorias diferentes: transações online X processamento (noturno) em batch
  - SLAs (acordos de nível de serviço) diferentes
    - Transações online têm que ser atendidas em menos de 1 segundo, em média, enquanto processos batch não têm requisitos estritos

Qual o tempo de resposta médio de uma transação do tipo X?

Qual o throughput the transações do tipo X?

# Problemas, problemas, problemas

- Suponha agora que o mesmo servidor seja usado como parte de uma aplicação cliente/servidor. Estações de trabalho clientes são conectadas ao servidor de banco de dados através de uma rede LAN. Clientes trabalham de forma independente e alternam entre processamento local ("thinking") e "waiting" por uma resposta do servidor.

Qual o tempo de resposta médio de uma transação de um determinado tipo, incluindo tempo de transmissão na LAN? Qual o throughput do servidor?

# Problemas, problemas, problemas

- Suponha agora que o servidor de banco de dados faça parte de um serviço Web utilizado por aplicações de terceiros. Suponha ainda que os gerentes de IT desejam prover uma garantia no tempo de resposta médio das transações para seus clientes (não pode ser superior a 4 segundos, em média). Para tanto, um sistema de controle de admissão é adicionado ao sistema de forma a limitar o número de transações em processamento simultâneo a no máximo 10 .

A atual configuração do servidor irá atender a necessidade dos gerentes?

Se não, o que terá melhor relação custo-benefício: comprar um disco a mais e fazer balanceamento de carga ou comprar uma CPU 2x mais rápida

# Problemas, problemas, problemas

- O servidor de banco de dados faz parte de um serviço Web, composto de um servidor HTTP (front-end) e um servidor de aplicação. Os gerentes do serviço querem garantir que a probabilidade de uma requisição ao **serviço** tenha um tempo de resposta superior a 5 segundos seja no máximo 0.05:

$$\text{Prob}(R > 5) = 0.05$$

A atual configuração do servidor irá atender a necessidade dos gerentes?

Se não, o que terá melhor relação custo-benefício: comprar um disco a mais e fazer balanceamento de carga ou comprar uma CPU 2x mais rápida

# Problemas, problemas, problemas

Facing sharp competition, consider an automotive parts distribution company that has decided to develop and implement a new call center. The goals of the new call center system are to: 1) foster better relationships with customers, creating customer loyalty and ensuring quality service, 2) improve efficiency and service performance, and 3) identify and explore new sales opportunities.

The IT team has decided to create an integrated customer relationship management system. The proposed new architecture consists of computer telephony integration software that controls inbound calls and routes these calls to service representatives integrated with application software.

# Exemplo de problema que vamos abordar

The call center will be staffed by service representatives who handle customer orders, returns, queries by phone, fax, and email.

Each customer service representative will have immediate access to the following basic functions: 1) historical tracking of all customer contacts, 2) a single view of customer information, 3) records of resolution of past problems, and 4) help functions. Since customers are less tolerant of long calls, the managers require that the system meets sub-second response times on all functions and be operational 24 hours a day; 7 days a week.



# Problemas, problemas, problemas

The company wants to consider various scenarios when designing its call center. The IT team is planning to design, build, test, and deploy the new call center applications in 12 months. Before rollout, management wants assurances that the performance requirements are satisfied. Questions that they might ask:

- Is the system design able to meet the subsecond response time for all functions? Is that feasible?
- What will be the impact of doubling the number of system representatives in the next year?
- Can acceptable performance levels be maintained after integrating the system with the mainframe-based inventory application?
- Is the system capacity adequate to handle up to 1000 calls in the busiest hour and yet preserve the subsecond response time goal?
- How do failures in the database server affect the 24x7 availability goal?
- What is the impact of starting to offer Web-based self-service to customers?

Problemas,  
problemas  
problemas

NYTimes,  
Apr 19, 2007

Justin Sullivan/Getty Images

Douglas Philips scanned his BlackBerry for e-mail messages in San Francisco Wednesday. More than five million users of the wireless devices were affected by an unexplained nationwide loss of service for 10 hours.

By **BRAD STONE**

Published: April 19, 2007

Where were you when the BlackBerrys went out?



Fred R. Conrad/The New York Times

Ilana Raz checked her BlackBerry outside Grand Central Terminal.

On Tuesday night at 8 p.m. Eastern time, technical problems cut off more than five million [BlackBerry](#) users in the United States from their cherished wireless e-mail. Service was restored 10 long, data-starved hours later.

The BlackBerry blackout was grueling to many — and revealed just how professionally and emotionally dependent so many people had become on their pocket-size electronic lifelines.

Stuart Gold was in Phoenix on a business trip when the service went down. Mr. Gold, the marketing director for Omniture, a software firm, noticed ominous red X's next to his outgoing e-mails.

E-MAIL

PRINT

REPRINTS

SAVE

SHARE

ARTICLE TOOLS  
SPONSORED BY

*Waitress*

Part of the problem, though, could be the service's rapid growth: R.I.M. says it has added three million subscribers in the last 12 months, for a total of eight million, in part because of the popularity of its superslim BlackBerry Pearl.

Problemas,  
problemas,  
problemas

NYTimes,  
Feb 9, 2008

The New York Times

## Media & Advertising

WORLD U.S. N.Y. / REGION BUSINESS TECHNOLOGY SCIENCE HEALTH SPORTS OPINION  
MEDIA & ADVERTISING WORLD BUSINESS SMALL BUSINESS YOUR MONEY DEALBOOK MARKETS



### Nielsen Tells TV Clients It Is Working on Ending Delays in Ratings

By [BILL CARTER](#)  
Published: February 9, 2008

With many of its television clients, including the broadcast networks, expressing increasing frustration with delays in reporting ratings for programs, the Nielsen Company sent a letter on Friday to its customers expressing disappointment with its performance and promising to do better — though not for a while.

E-MAIL

PRINT

REPRINTS

SAVE

SHARE

→ Nielsen now plans to upgrade its computer servers, which are at the company technology base in Oldsmar, Fla. But that may take some time. In its letter, the company said, “We expect to see a material reduction in the number of data delays beginning in mid-March with continuous improvement every 20 business days thereafter for the next four months.”

→ He and other Nielsen critics want to know why the company did not foresee the vast new demands and plan accordingly.

→ Ms. Herkovic conceded the point. “I don’t think we did as good a job as we should have in anticipating these spikes,” she said.

# Problemas, Problemas, Problemas

## Michael Jackson Invented Pop, Beat the Rap, Crashed Google

BY CLAY DILLOW Fri Jun 26, 2009 at 12:00 PM



The news of Michael Jackson's passing yesterday shocked the world, but perhaps none were as stunned as Google and Twitter. Upon rumors of the pop icon's death in Los Angeles, fans flocked to the Web for confirmation, resulting in an avalanche of searches, Tweets, and page views that crashed Google and brought Twitter to a standstill as the number of Tweets per second doubled in an instant.

The tremendous volume of queries for the term "Michael Jackson" was interpreted by Google as an automated attack, triggering a collapse of the site that lasted for about half an hour. More than one-third of the top 100 Google searches among U.S. users were reportedly M.J.-related during the height of the frenzy.

The



Meanwhile Twitter managed to remain online, but it might as well have crashed. According to Twitter analysis tool [Twist](#), at one point yesterday afternoon 22.6% of messages posted to Twitter were related to the King of Pop's passing as users barraged the site with questions and rumors, and later fond farewells. At one point, nine of the top ten topics on Twitter involved Jackson's death. The surging volume caused the site to grind to a halt temporarily, and Twitter Search was reportedly running about twenty minutes behind real time. This is the same Twitter the U.S. State Department asked to [remain functional](#) last week so protesters in Iran would be able to

# Problemas, Problemas, Problemas

## The PayPal outage cost its users between 7 and 32 million USD

Posted in [Main](#) on August 4th, 2009 by Pingdom

As you may know, PayPal suffered from downtime yesterday. Overall, the problems lasted approximately 4.5 hours before being fully resolved. Since a significant number of e-commerce sites and online services handle some or all of their transactions through PayPal, how much money did the PayPal outage end up costing its users?



According to eBay (which owns PayPal), about \$2,000 in payments flow through PayPal's systems every second. This means that PayPal processes about **\$7.2 million in payments every hour** for its users.

On its official blog [PayPal has stated](#) that the service was completely down, globally, for about one hour. So the outage cost PayPal users *at least* around \$7 million. But PayPal also admits that it took an additional 3.5 hours after that before the service was fully restored, something we also need to take into account.

**This means that the outage and following service problems cost Paypal users somewhere between \$7 million and \$32 million in lost payments.**

Exactly where the actual cost landed between those two extremes is hard to tell, but judging by the comments we have seen in blogs and on Twitter, the problems were significant even after the initial one-hour service outage.



# Problemas, Problemas, Problemas



 EVENTS

 GIGAOM.TV

 PRO RESEARCH

HOMEAPPLEBROADBANDCLEANTECHCLOUDCOLLABORATIONMOBILEVIDEO

 NewTeeVee

Search

## CNN: Inauguration P2P Stream a Success, Despite Backlash

By Janko Roettgers | Feb. 7, 2009, 12:01am PT | 16 Comments

 Tweet 0

 Share

 Like







Sometimes it's tough to be the winner. CNN [attracted by far the largest online audience](#) for its web stream of President Obama's inauguration, delivering a total of 25 million streams on Jan. 20th. It did so by utilizing [Octoshape's P2P plug-in](#) for Adobe's Flash player, a fact that was widely known within the industry but only received very [little coverage](#). Until yesterday, that is, when [Windows Secrets](#) suddenly came out [blasting](#) the network for "deceptive marketing," alleging that CNN hijacked its users' upload bandwidth. Others quickly jumped on the bandwagon, [going so far as to suggest](#) that "the most trusted name in news just stole your computer."

I got in touch with CNN and Octoshape to get both sides of the story, which are, unsurprisingly, slightly different. Both companies reject the claim that they misled viewers. But Mike Wise, technical adviser of the R&D group of CNN parent company Turner Broadcasting System, had an even more important point to make: CNN used P2P, he told me, because it had to.

The inauguration is now seen as one of if not the largest live video event in the history of the Internet, and the folks at CNN believe that traditional CDNs simply weren't up to the task. "Turner anticipated that the resulting stresses on the various CDNs and the Internet in general would make it extremely difficult if not impossible to serve such an unprecedented audience without the use of P2P technologies." Wise told me in an email.

 **Skytap**

Free Trial

CLOUD FREEDOM  
Develop, test, share,  
collaborate, & relax.



Let the MPAA Speak, There's Nothing To Be Scared Of | TorrentFreak  
<http://t.co/rND73pJ>

### JANKO'S POSTS

Netflix tests new web UI for kids. First look

Cord Cutters: Checking out the Boxee iPad app

### GIGAOM TV



Play video

James Yu and Tikhon Bernstein

# Problemas, Problemas, Problemas

## Amy Winehouse death gets 20 million Twitter mentions

Lucy Carne | The Sunday Telegraph | July 24, 2011 1:34PM | 25 comments

Share

Recommend

Send

28 people recommend this. Be the first of your friends.

10 retweet

Share



The body of dead English pop star Amy Winehouse is carried from her house in north London.

Picture: AFP Source: The Sunday Telegraph


< Prev 1 of 3 Next >

### AMY Winehouse's tragic death triggered an instant outpouring of grief that swept the world.

Within seconds of the confirmation that the British five-time Grammy award-winning singer had died in her London home, fans and celebrities took to Twitter to express their shock and sorrow.

In a wave of mourning, Winehouse became the highest searched name on the Internet and was mentioned by 20 million people on Twitter within minutes of the news breaking of her death.

Celebrities, many of whom were friends with the singer, expressed their immense sadness at her sudden death, believed to be from a drug overdose.




# Principais Desafios

- Complexidade e diversidade de aplicações, serviços, plataformas e infraestruturas







# Principais Desafios

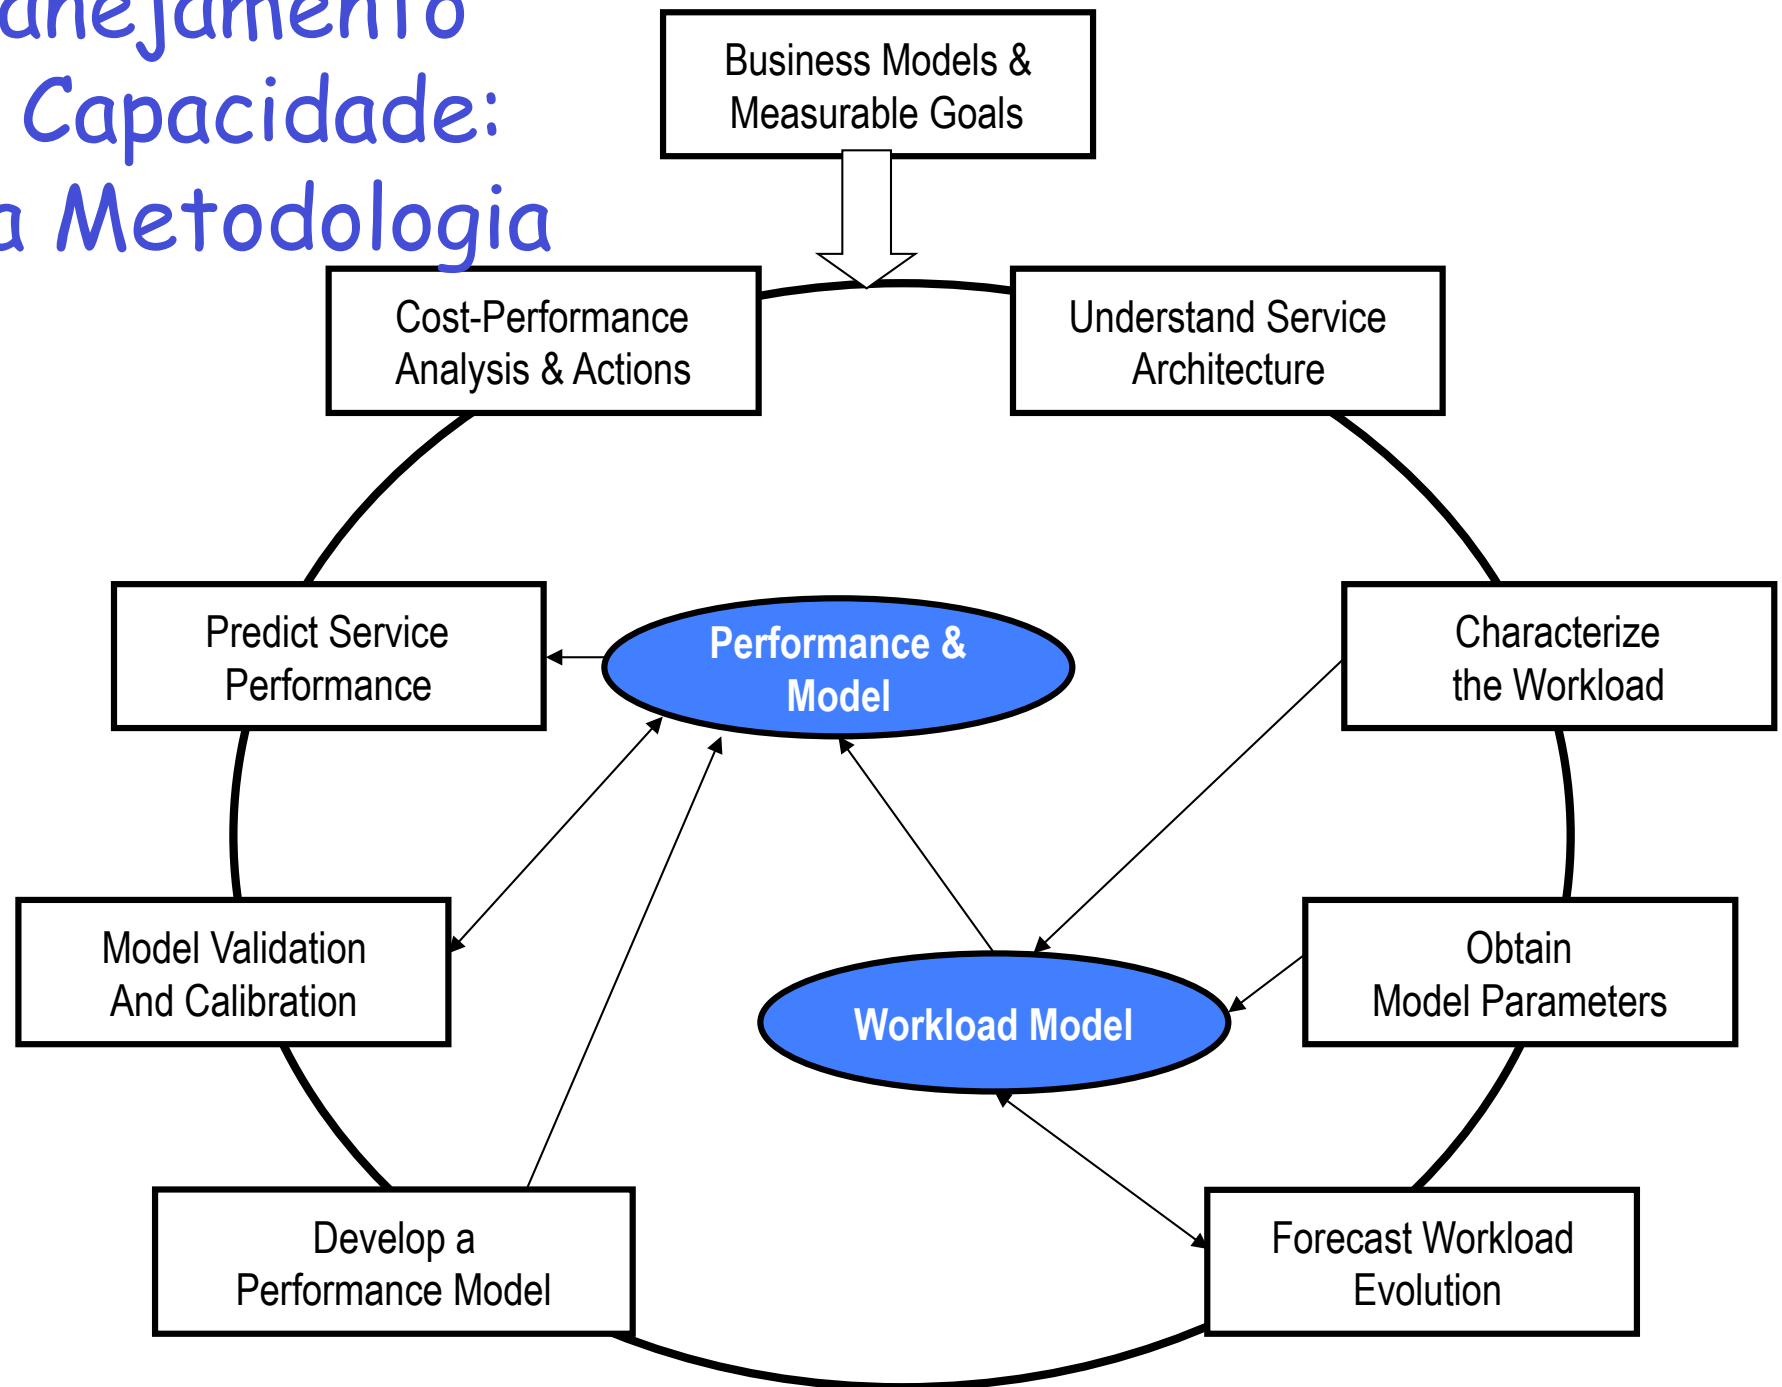
- Grandes volumes de dados:
  - 2 bilhões de usuários da Internet
  - + 500 milhões de usuários do Facebook
  - + 200 milhões de usuários do Twitter
  - + 100 milhões de tweets por dia
  - + 2 bilhões de vídeos vistos por dia no YouTube
- Dados raramente disponíveis publicamente
- Coleta, processamento e armazenamento:
  - Amostragem
  - Representatividade

# Principais Desafios

- Padrões heterogêneos e dinâmicos
  - Evolução temporal
  - Alta variabilidade e mudanças drásticas (flash crowds)
  - Interpretação sujeita a alto grau de subjetividade



# Planejamento de Capacidade: Uma Metodologia



# A disciplina

- Objetivo: aprender técnicas que permitam responder perguntas como as mostradas com baixo custo e precisão aceitável.

Good models abstract out some of the complexities of real systems while retaining what is essential to meet the model goals

- Foco:
  - Modelagem analítica de desempenho: fórmulas **simples**, que fornecem resultados **rápidos** e com **precisão adequada** para vários **problemas práticos**.
  - Caracterização e modelagem de cargas de trabalho

# A disciplina

- Técnicas:
  - Modelagem de Desempenho
    - Leis fundamentais
    - Análise de valores limites
    - Análise de valores médios
  - Modelagem de Cargas
    - Caracterização, benchmarks
- Discussão de vários artigos atuais
  - Apresentações/seminários por alunos
- Avaliação (tentativa):
  - Projeto : 35 pontos
  - 2 provas : 50 pontos
  - Participação (seminários, discussão, listas de exercícios): 15 pontos