



SE252:Lecture 1, 6 Jan 2015

Introduction to Cloud Computing

Yogesh Simmhan





SE252

Introduction to Cloud
Computing, *Jan 2014*



Who am I?

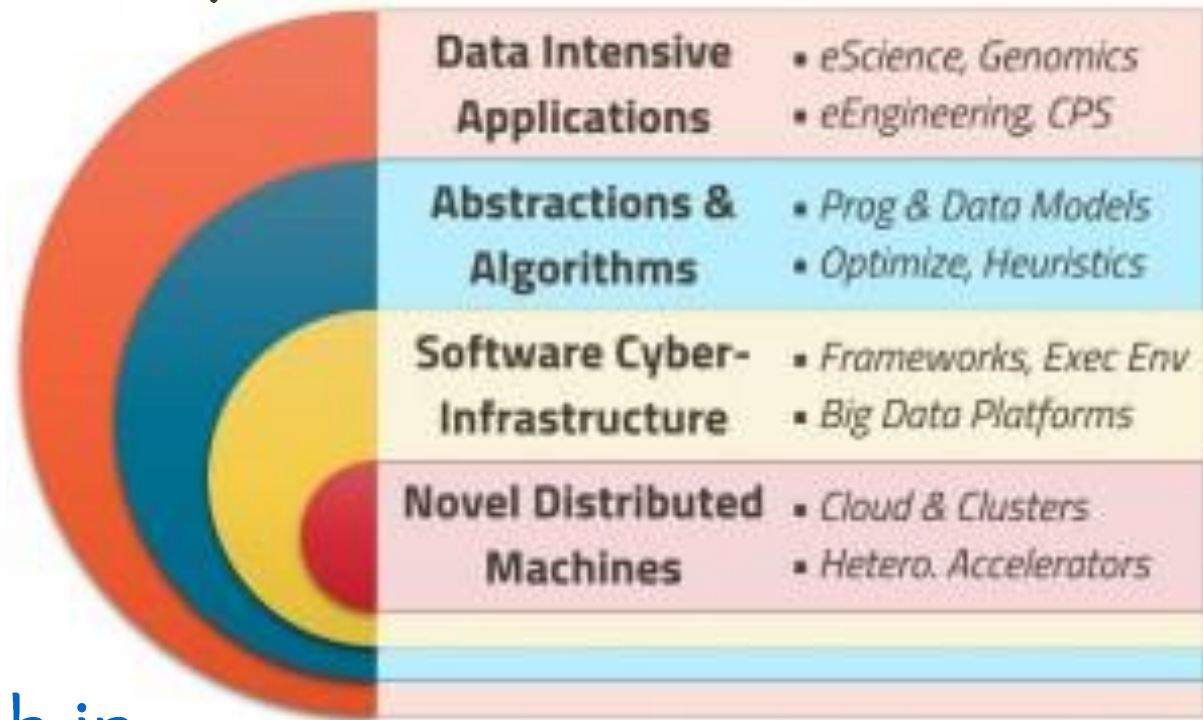
- “Yogesh”
 - ...or Dr. Simmhan
 - ~~Not Dr. Yogesh, “Sir”, other honorifications 😊~~
- Assistant Professor, SERC
- Former:
 - Research Assistant Professor, USC, Los Angeles
 - Postdoc, Microsoft Research, San Francisco
 - Indiana University graduate



The DREAM:Lab

- **Distributed Research on Emerging Applications and Machines**
 - Applied distributed systems research

Open for New Students!



- www.dream-lab.in



Who are you?

- Name
- Department
- Degree & Year
- Have you signed up?



What this course is

1. Cloud computing as a technology
 - How to use Cloud tools, APIs, SDKs?
 - *Access to Amazon AWS IaaS Cloud*
2. Cloud computing as a distributed systems environment
 - Why cloud computing works?
 - How to design applications for Clouds?
3. Cloud computing as a research topic
 - What are the gaps and emerging ideas?



This course is NOT for you if:

- You have not taken algorithms/data structures/OS/networking UG courses
- You're not comfortable with programming
 - Network programming
 - ✓ Java, Python
 - .NET
 - ✗ ~~C/C++~~
 - *No TA, so you should be able to resolve programming problems on your own!*



Intended Learning Outcomes (ILO)

■ What you need to be *able to do* after *learning* from the course...expectations

- 1. Remembering** Define, describe, draw, find, identify, label, list, match, name, quote, recall, recite, tell, write
- 2. Understanding** Classify, compare, exemplify, conclude, demonstrate, discuss, explain, identify, illustrate, interpret, paraphrase, predict, report
- 3. Applying** Apply, change, choose, compute, dramatize, implement, interview, prepare, produce, role play, select, show, transfer, use
- 4. Analysing** Analyse, characterize, classify, compare, contrast, debate, deconstruct, deduce, differentiate, discriminate, distinguish, examine, organize, outline, relate, research, separate, structure
- 5. Evaluating** Appraise, argue, assess, choose, conclude, critique, decide, evaluate, judge, justify, predict, prioritize, prove, rank, rate, select, monitor
- 6. Creating** Construct, design, develop, generate, hypothesise, invent, plan, produce, compose, create, make, perform, plan, produce

Source: Anderson and Krathwohl (2001); Biggs and Tang (2007)



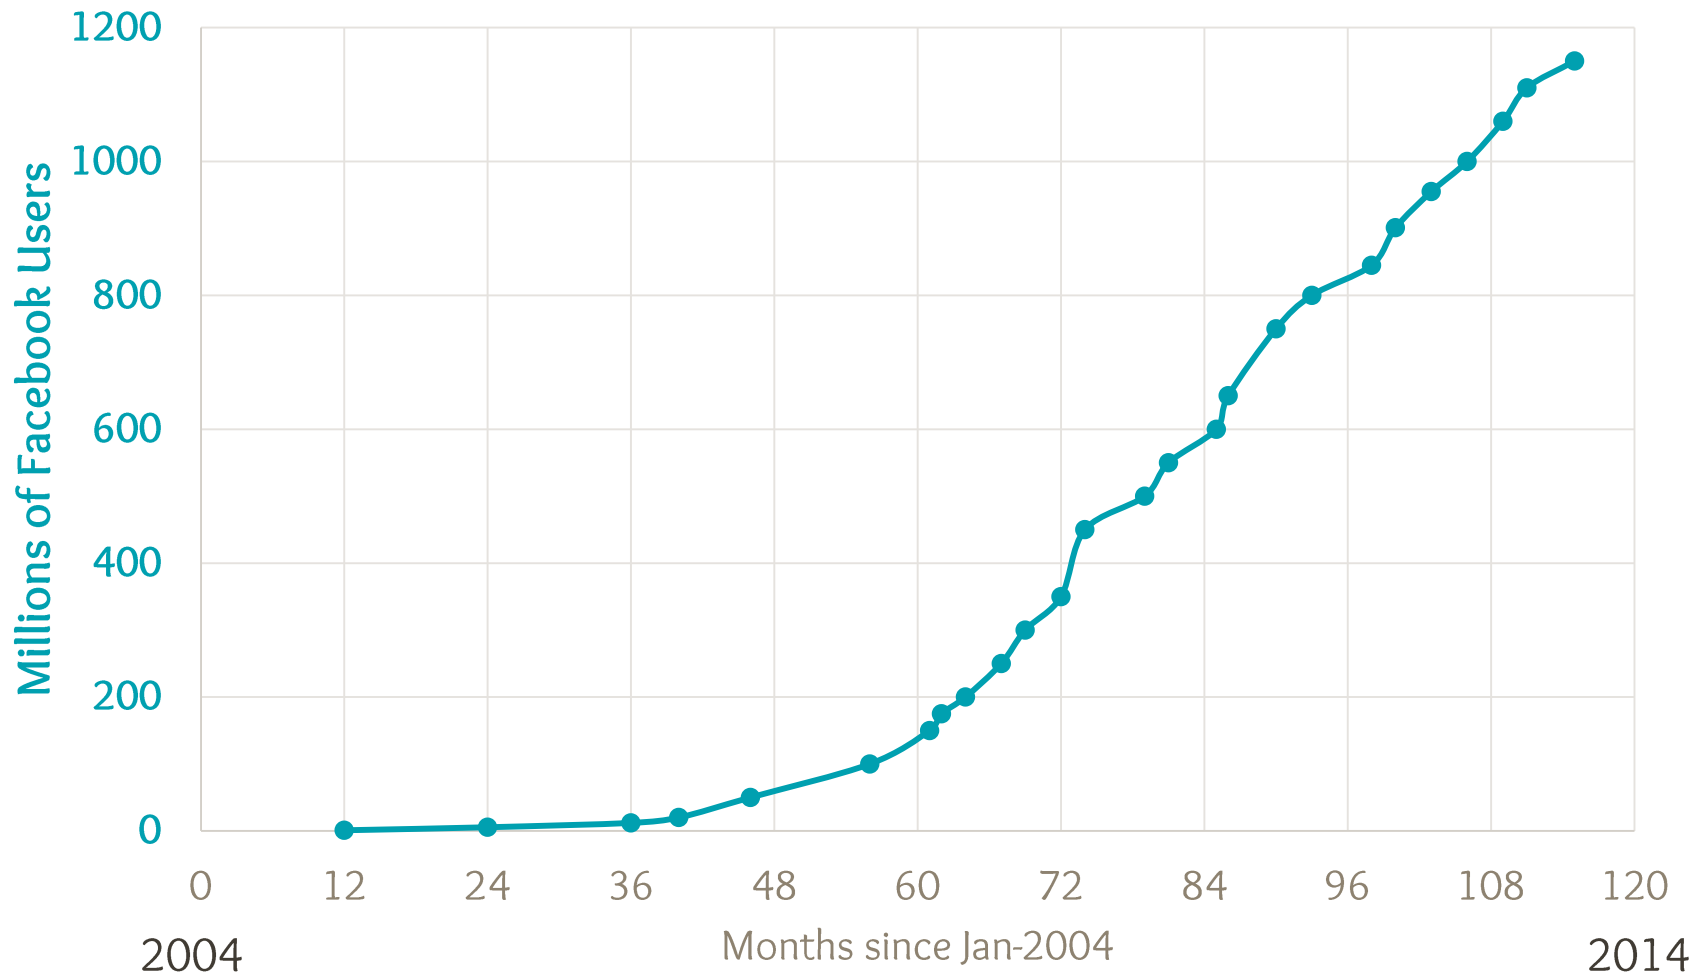
Facebook: A *Canonical* Cloud Based Application



ILO1: Parallel and Distributed Systems Context

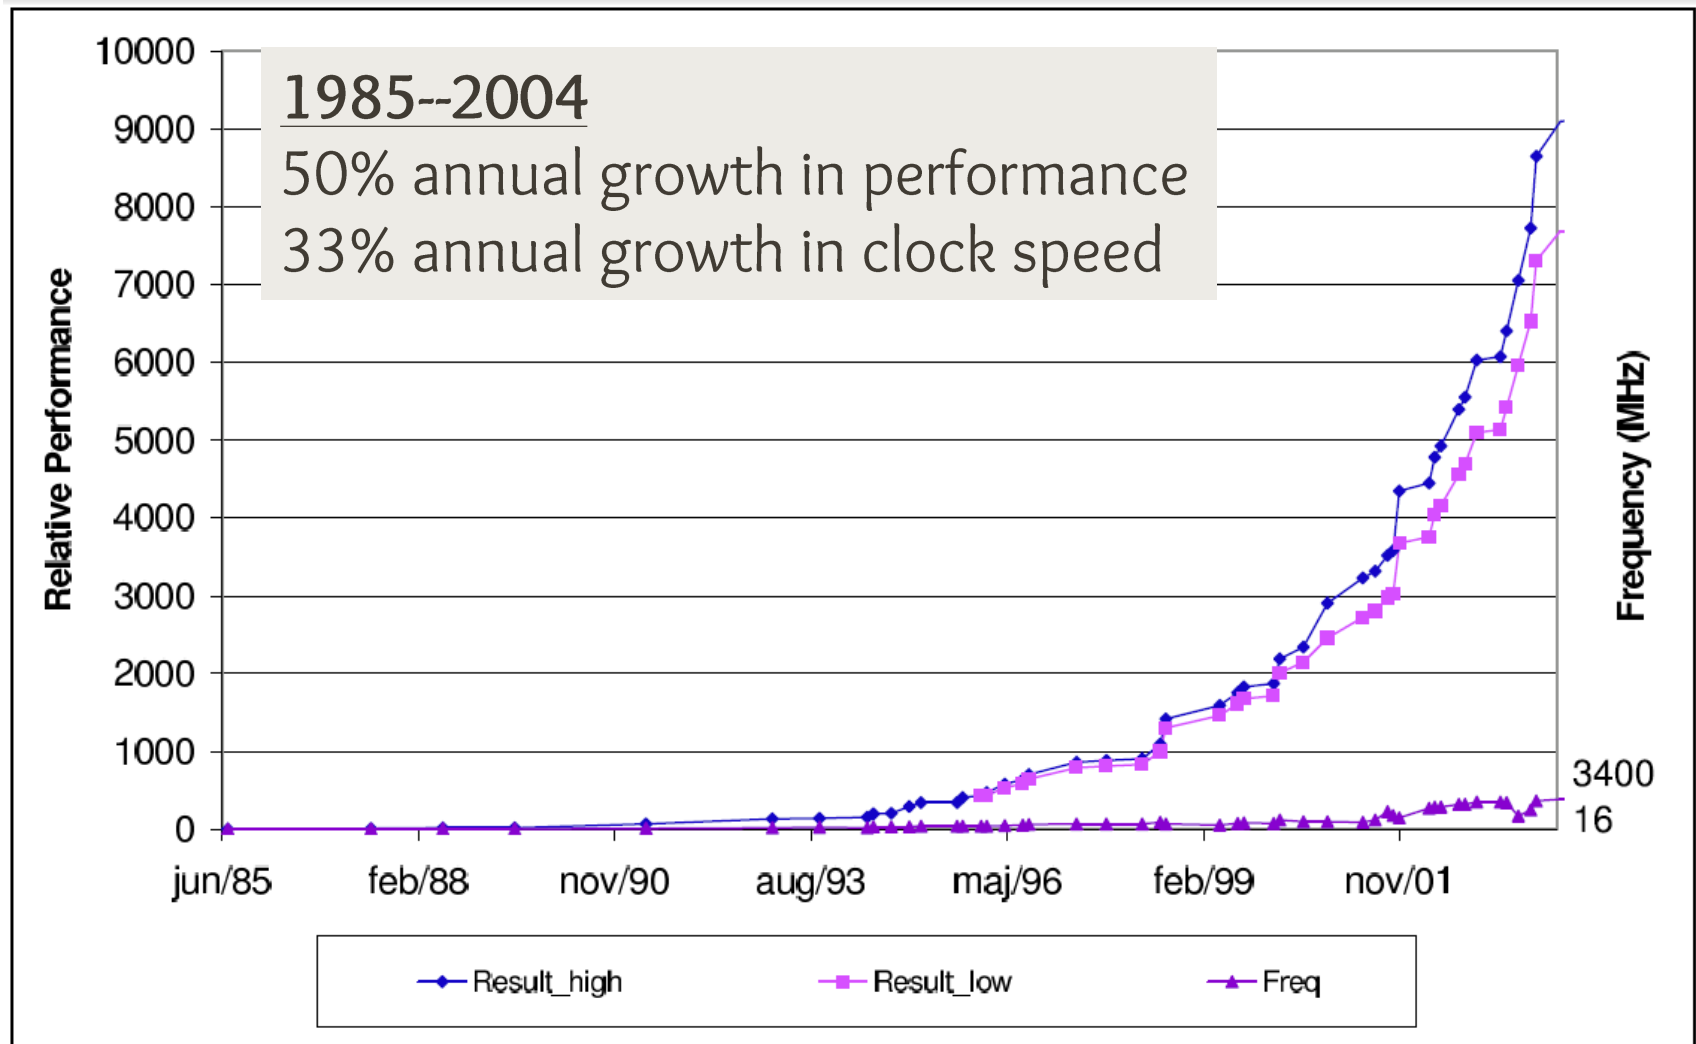


Growth of Facebook





Single CPU Performance (Scale **UP**)

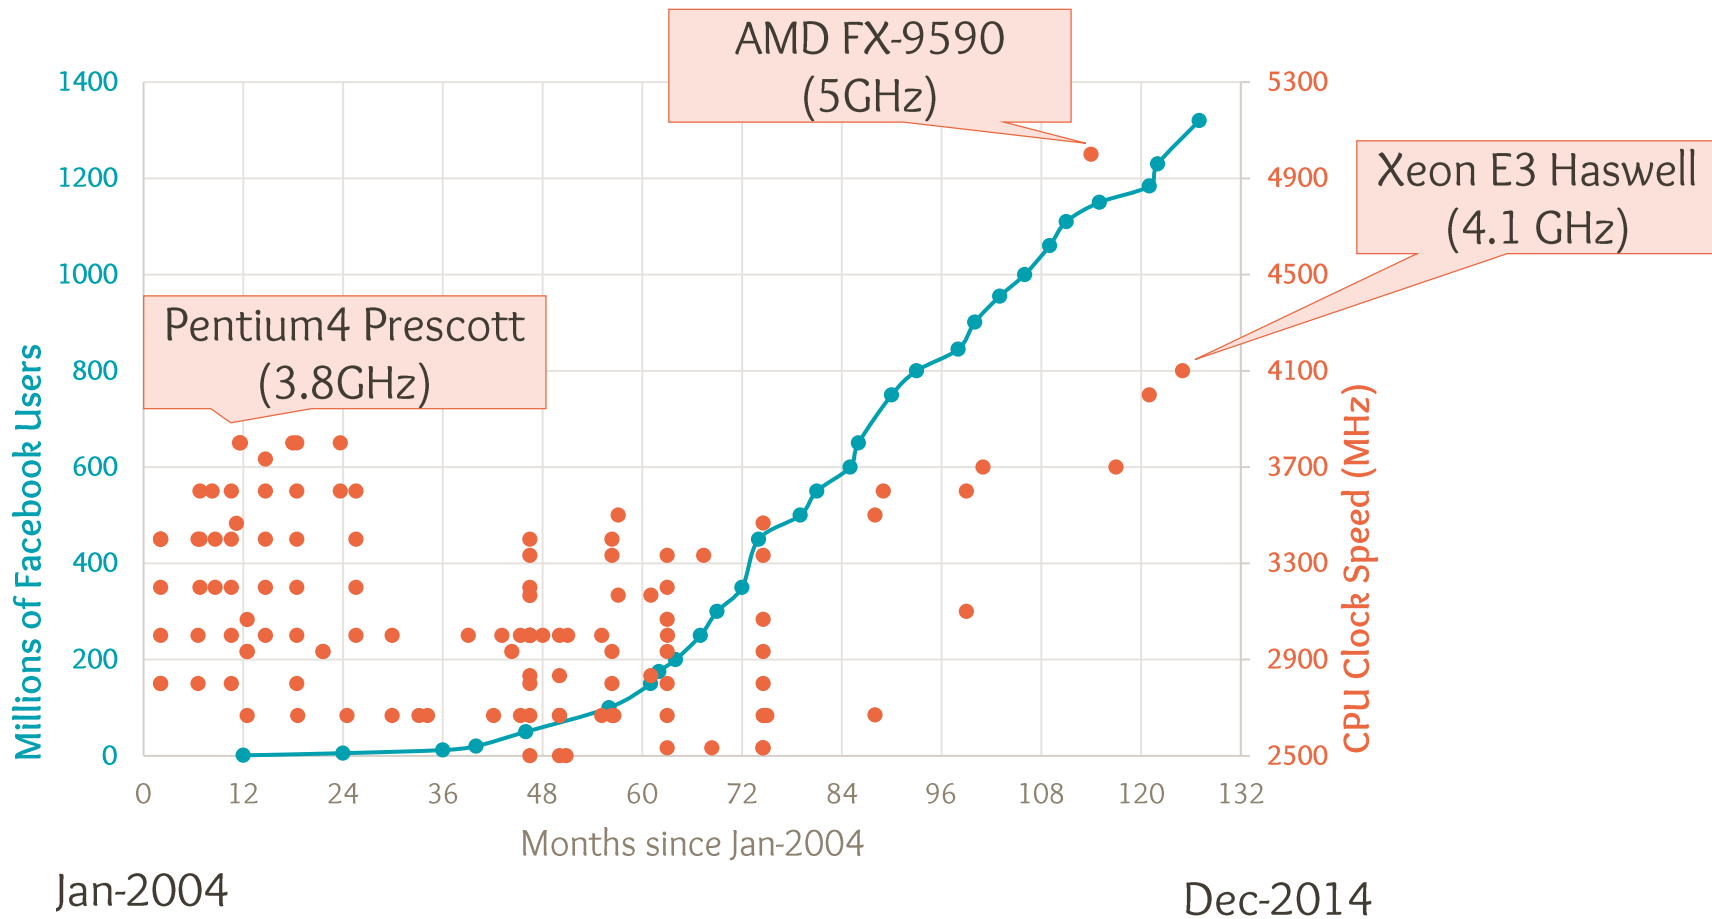


Computer performance 1985-2004, normalized to SPEC CINT89.

"An In-Depth Look at Computer Performance Growth", Ekman, et al, 2004

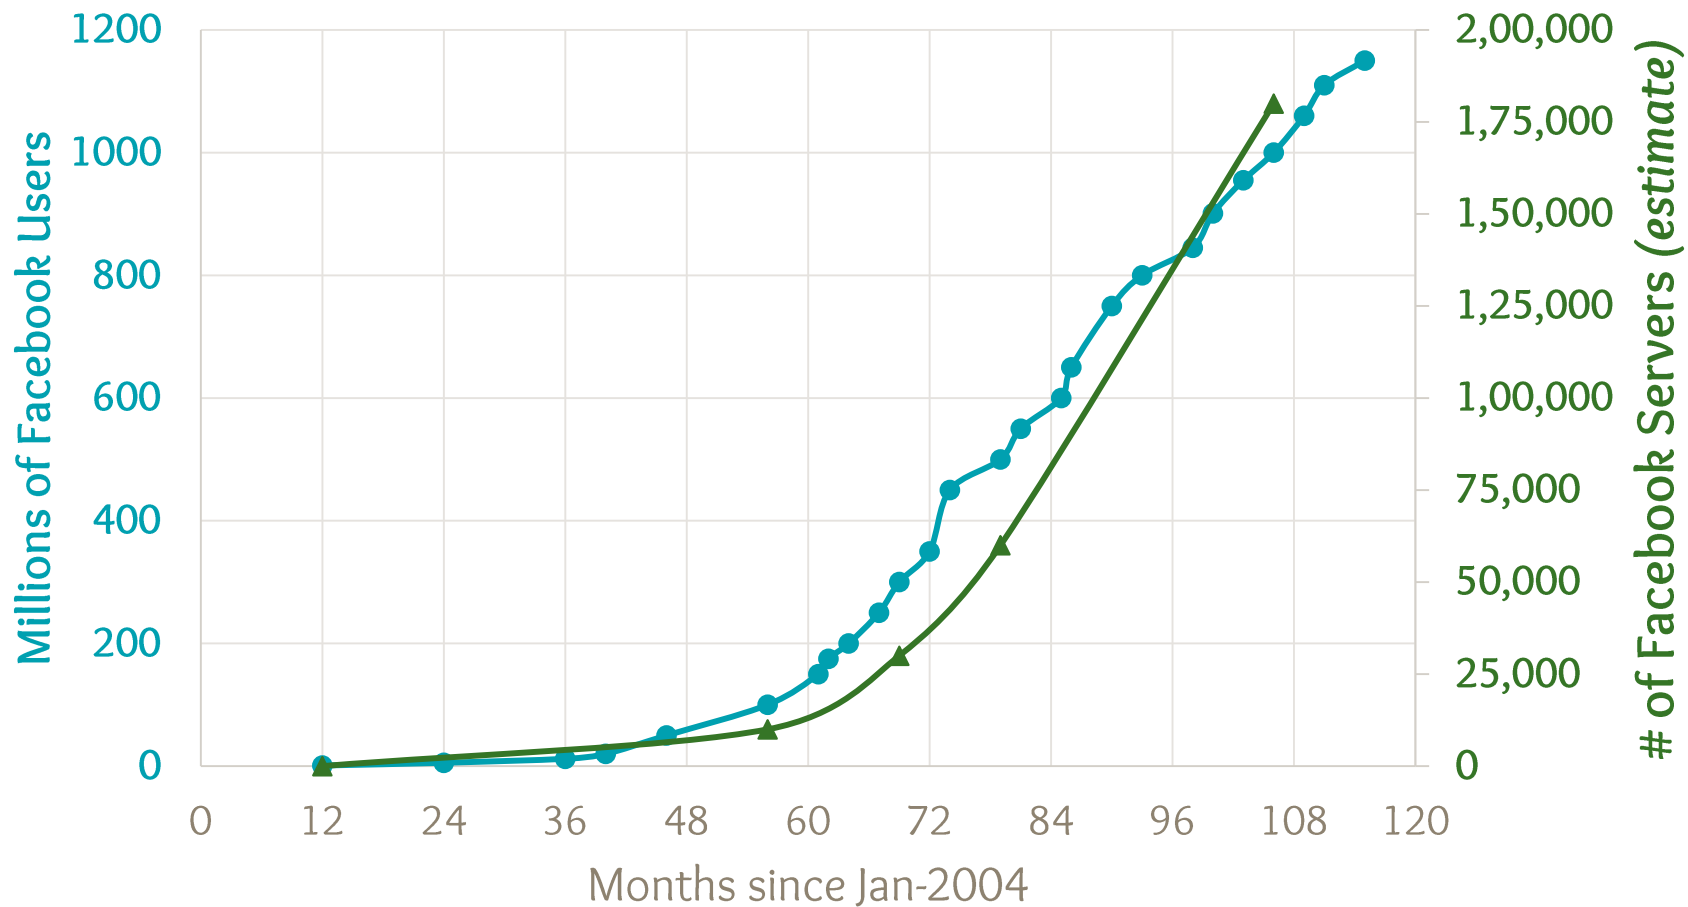


Facebook & Clock Speed Growth





Facebook & Scale **OUT**





Scale **Up** vs Scale **Out**

- Understand parallel computing paradigms
 - Shared memory vs. Distributed memory
 - Parallel, Distributed, Cloud, ...
 - Concurrency from micro to macro scale
 - Performance, scalability, efficiency
- Taxonomy, Pros & Cons, When to pick what, ...
- **ILO1: Parallel and Distributed Systems Context**



ILO2: Cloud Virtualization, Abstractions and Enabling Technologies



My Facebook "Groups"

Facebook interface showing the profile of **Yogesh Simmhan**. The page includes a search bar, a status update field, and a navigation menu. The **GROUPS** section is highlighted with a red rounded rectangle, listing:

- d2icenter
- AID (Association for India's ...)
- Groups at SRM University
- Create Group...



My Browser



"Facebook"
Server



Client-Server Interactions

- Request-Response model
- Often ask for a static content
 - https://www.facebook.com/images/loaders/indicator_blue_medium.gif 
 - <https://fbstatic-a.akamaihd.net/rsrc.php/v2/yK/r/Rf0y2zAXKrZ.gif> 
 - Served by facebook or other Content Distribution Networks
- Sometimes perform a “Remote Procedure Call”
 - **Get** the list of **groups** that **I** belong to
- Serialization, deserialization





developers.facebook.com/tools/explorer

Facebook Graph Explorer

GET

Learn more about the Graph API syntax.

Node: 6839096

- id
- name
- groups →

+ Search for a field

+ Search for a field

Request: User ID & Fields to "GET"

Response: List of group names

RPC using RESTFull
Request Response
using HTTP
Standards

```
{
  "id": "6839096",
  "name": "Yogesh Simmhan",
  "groups": {
    "data": [
      {
        "name": "AID Bangalore",
        "unread": 7,
        "bookmark_order": 2,
        "id": "536069206450363"
      },
      {
        "name": "d2icenter",
        "unread": 13,
        "bookmark_order": 1,
        "id": "268551846502548"
      },
      {
        "name": "AID (Association for India's Development) ",
        "administrator": true,
        "bookmark_order": 3,
        "id": "4848588894"
      }
    ]
  },
  "paging": {
    "next": "https://graph.facebook.com/v2.2/6839096/groups?icon_size=16&limit=5000&offset=5"
  }
}
```

D
E
M
O

<https://graph.facebook.com/v2.2/6839096?fields=id%2Cname%2Cgroups&format=json&method=get>



Service Oriented Architecture (SOA)

- Allows standard way for clients & services to
 - Exchange **data Structures** over the network
 - » Serialization & Deserialization
 - Invoke **remote methods** over the network
 - Defining QoS, discovering services, etc.
- SOA is an enabling *technology* for Clouds
 - E.g. **SOAP**, **REST**, **Thrift***, **Protocol Buffers**, etc.
 - RPC *concept* has existed since 1980's



SOA & Virtualization

- All Cloud operations can be performed as web service calls by application
 - Storage services, VM Management services,...
- Virtualization is the other key enabler
- ILO2: Cloud Virtualization, Abstractions and Enabling Technologies



ILO3: Algorithms and Programming Patterns for Cloud Applications



Facebook Webpage over the Years

http://mashable.com/2011/02/04/facebook-7th-birthday/

[Welcome to Thefacebook]

Thefacebook is an online directory that connects people through social networks at colleges.

We have recently opened up Thefacebook at the following schools:

Alabama • Appalachian State • Cooper Union • CSU Sacramento
 CUNY Hunter • Denison • Furman • Gonzaga • IUPUI • Kutztown
 Marymount • Marymount Manhattan • Monmouth • Montclair
 MS State • Muhlenberg • N. Arizona • North Dakota • Portland
 Rose-Hulman • Sacred Heart • Sarah Lawrence • Seton Hall
 St. Johns • St. Lawrence • SUNY New Paltz • SUNY Oswego
 Texas State • Toronto • Tulsa • University of the Pacific • Utah • Washington State
 Willamette • Wisconsin Milwaukee • WPI • Xavier

For a complete list of supported schools, click [here](#).

Your facebook is limited to your own college or university.

You can use Thefacebook to:

- Search for people at your school
- Find out who is in your classes
- Look up your friends' friends
- See a visualization of your social network

2004 1 million users

Brian Moore's Profile

Account Info:
 Name: Brian Moore
 Member Since: May 21, 2005
 Last Update: July 19, 2005

Basic Info:
 School: Puget Sound '09
 Status: Student
 Sex: Male
 Residence: Todd 311
 Birthday: 09/02/1986
 Home Town: Shorewood, WI 53111
 High School: Shorewood HS '05

Information:
 Email: bmoore@up.edu
 Surname: Doctabu
 Mobile: 414.702.7426
 Websites: <http://www.doctabu.com>
<http://www.livejournal.com/users/doctabu>
<http://www.flickr.com/photos/doctabu>

Personal Info:
 Looking For: Friendship
 Interested In: Women
 Relationship Status: In a Relationship with Rachel Buehle (Tiny Tyes Day Care)
 Political Views: Very Liberal
 Interests: Film, Graphic Design, Video Editing, Computers, Bowling, Dancing, Acting, Singing, Listening to Decent Music, Sleeping, Being Crazy
 Favorite Music: Beck, The Beatles, They Might Be Giants, Phoenix, Paul Simon, Outkast, Avianus G, Redwood, Red Hot Chili Peppers, Stevie

2005 5.5 million users

facebook home search browse invite help

Welcome Mark!

News Feed

David-Ezra Shamash uploaded a mobile photo. 2:21pm

Mobile Uploads
 2 photos by David-Ezra Shamash

James Wang created a group. 8:05pm

Chris Hughes joined the group Stop Nuclear Proliferation, At Home & Abroad. 7:41pm

Jon Hay and Meredith Loth ended their relationship. 6:02pm

Aditya Agarwal created a group. 1:42pm

Paul Oakenfold is washed up. Mark + Evonne
 info: Former Oakenfold fans.

Chris Hughes joined the group Feingold for President 2008. 8:22am

Daniel Corson went from being "married" to "in a relationship". 2:07pm

Katie Geminder uploaded a mobile photo. 12:49pm

Launch! 1 photo by Katie Geminder
 Location: HQ

Kevin M. Bombino is no longer in a complicated relationship. 12:34pm

2006 12 million users

facebook Profile edit Friends Networks Inbox home account privacy logout

Search

Matthew David Colebourne
 Update my status...

Networks: London
 Sex: Male
 Interested In: Women
 Relationship Status: It's Complicated
 Birthday: June 26
 Hometown: Wokingham, England
 Religious Views: Atheist

Mini-Feed

Information

Education and Work

qubox X

ccComment User Conversations X

Comments of Blackblade

After the Launch managing a business, managing a micro business, independent professional, web entrepreneur community, corporate branding, small business consultant, small business coach, small business expert, small business marketing, small business help, set up a business blog, how to get more clients, small business blogging 3 Steps of a Savvy Entrepreneur

I agree completely... particularly on the last point.

The beauty of web-based businesses, in particular, is the quick and relatively easy access to a huge market space (nationally and internationally) that was previously much harder to reach.

However, that means that any individual's ability to "predict" the market is much reduced... I'm sure, for example, that Skype had no idea they would be so successful in Roland prior to the event.

So, I'm always concerned about anyone who suggests that they know exactly what will work and what will not. I think that you need invention to produce new business offerings but then evolution to polish those offerings to meet the users needs. That means you have to learn from your users, other experts and the whole community.

2007 50 million users

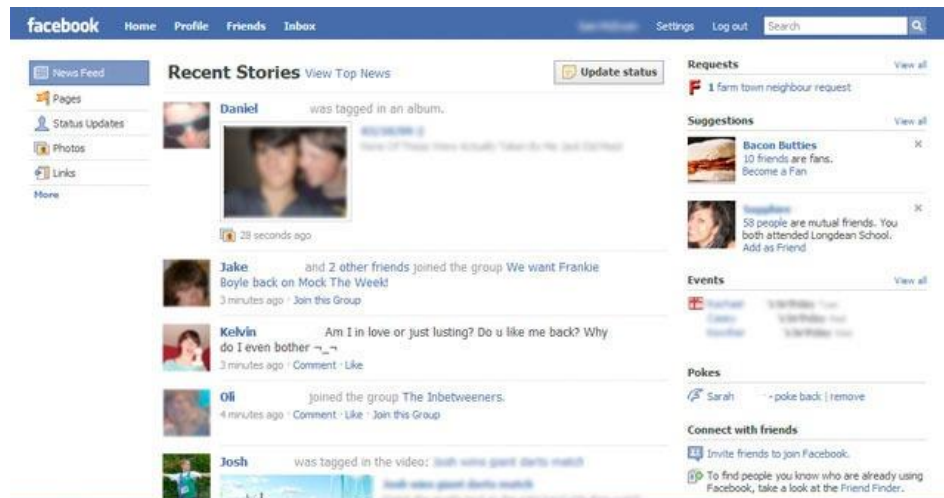


Facebook Webpage over the Years

http://mashable.com/2011/02/04/facebook-7th-birthday/



2008 100 million users



2009 350 million users



2010 500 million users



2011 Still growing



Facebook's BigPipe

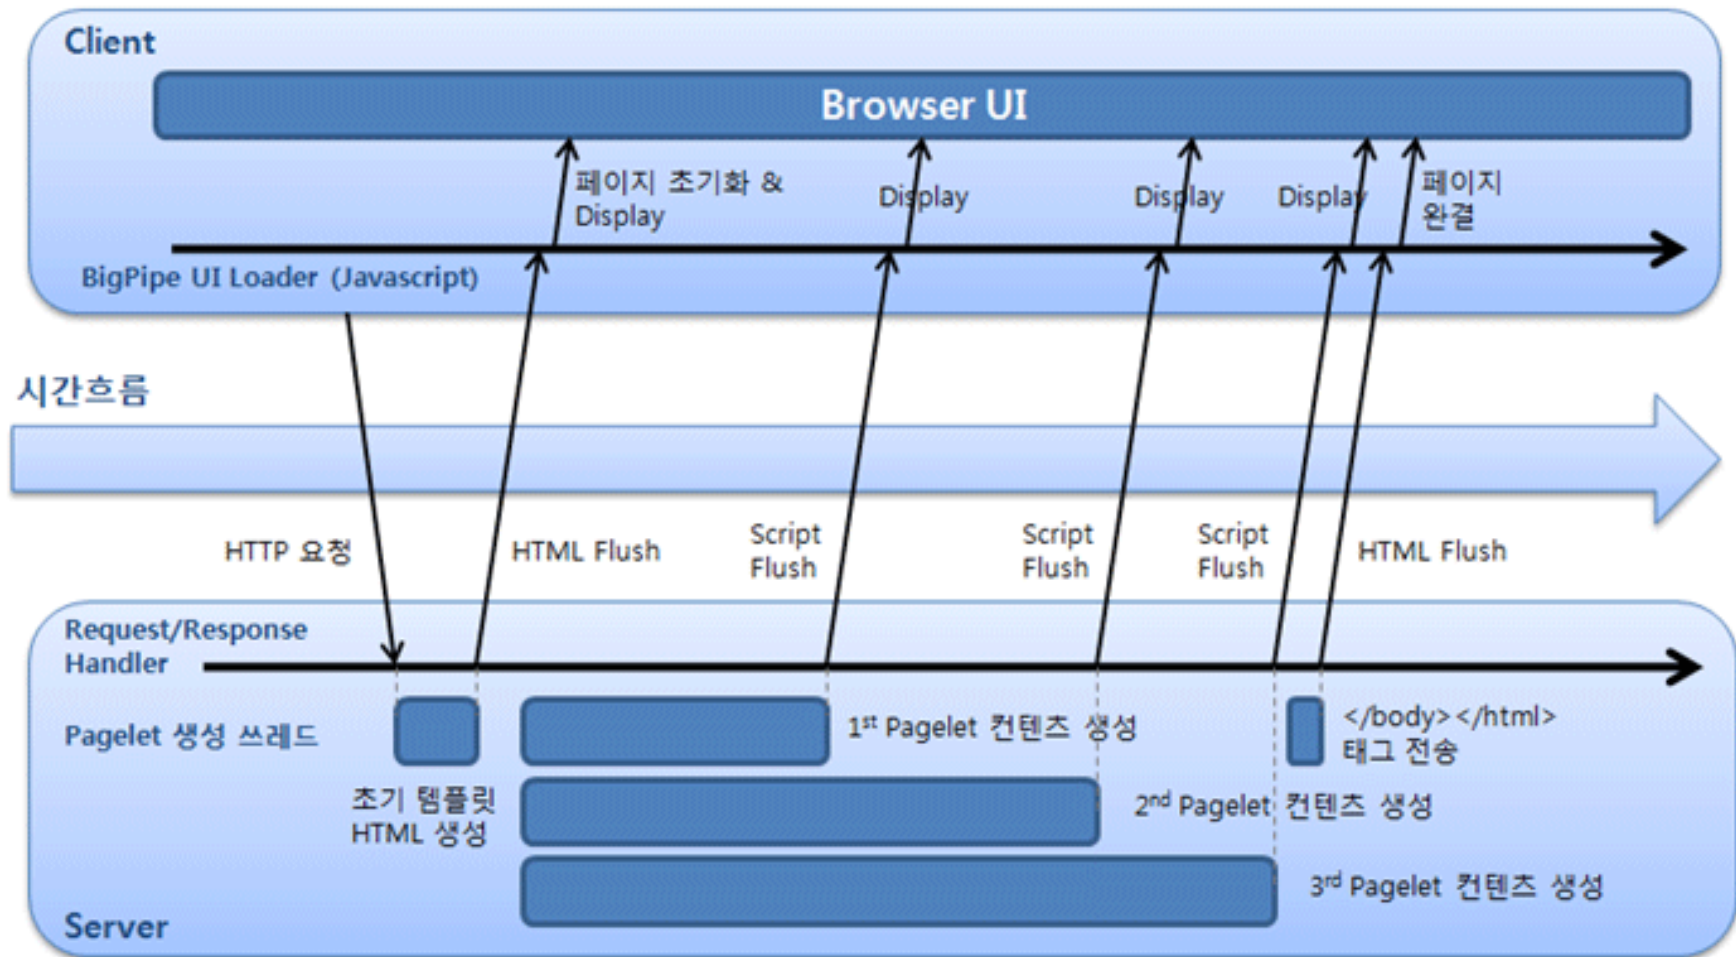
https://www.facebook.com/note.php?note_id=389414033919

The screenshot displays a Facebook news feed for user Changhao Jiang. The interface is divided into several sections:

- Header:** Facebook logo, search bar, and navigation links for Home, Profile, and Account.
- Left Sidebar:**
 - User profile for Changhao Jiang (Edit My Profile).
 - News Feed (selected).
 - Messages, Events, Photos, Friends (10).
 - Applications, Games, Ads and Pages, Groups (5), Local Picks.
 - Friends Online list: Christopher Palow, Eugene Letuchy, Hong Yan, Makinde Adeagbo, Mark Zuckerberg, Roger Erdong Chen, Wayne Kao, Ke Yang, Venkat Venkataramani, Xin Qi.
- News Feed:**
 - Placeholder: "What's on your mind?"
 - Paul Saab** (3 hours ago): "it's out! <http://github.com/facebook/flashcache> facebook's flashcache at master - GitHub". Includes a link to the GitHub repository and a comment from Mike Schroepfer.
 - Ben Maurer** (about an hour ago): "It'd be interesting to use ARC (http://en.wikipedia.org/wiki/Adaptive_replacement_cache) to automatically tune the dirty threshold. The ARC concept would need to be extended a bit (since a missed read is a bit different than a missed opportunity for write coalescing). - Why did you choose FIFO as the policy for dealing with dirtying?". Includes a comment input field.
 - Steven Grimm** (8 hours ago): "Shanghai's Back on Top of the World - TIME". Includes a link to the TIME article and a comment input field.
 - Facebook Engineering**: "Cultural Learnings from Russia: Engineering Version". Includes a link to the article and a comment input field.
- Right Sidebar:**
 - Requests:** 9 friend requests, 1 friend suggestion, 5 group invitations, 21 other requests.
 - Sponsored:** "Try Facebook Ads" advertisement with a "Create an Ad" button and text: "Reach the exact audience you want with Facebook's customizable targeting. Click here to learn more about advertising on Facebook."
 - Events:** Qiaozhu Mei's birthday (Today), Jingshu Huang's birthday (Monday), Xiong Zhang's birthday (Monday), Julie Tung's birthday (Tuesday).
 - Get Connected:** "Who's on Facebook?", "Who's not on Facebook?", "Who's here because of you?", "Connect on the go".
- Bottom:** Chat (23) notification.



Pipelined vs Single Shot Page Loading

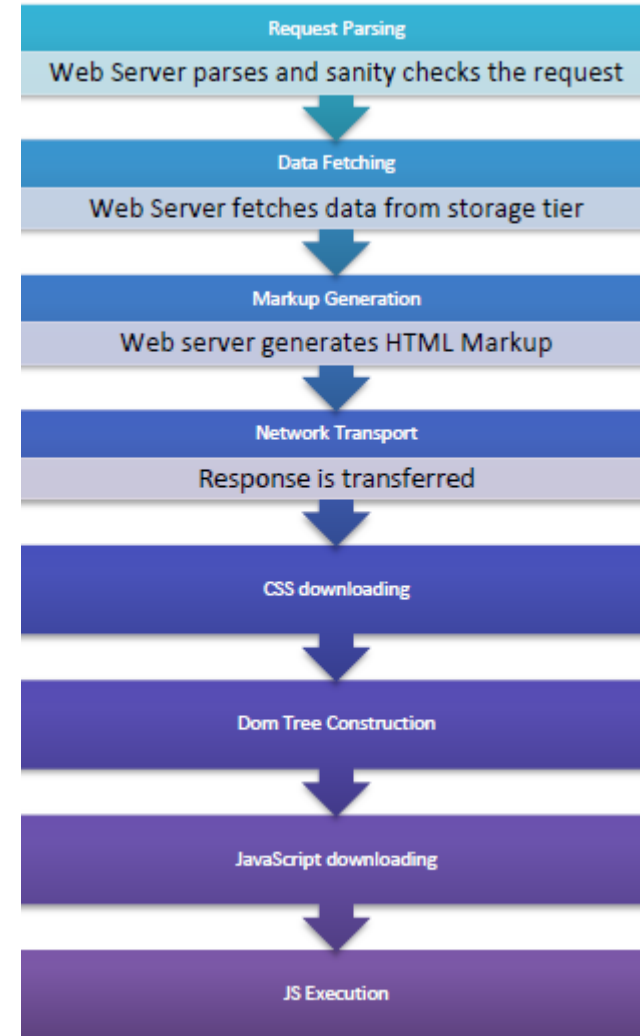


DEMO



Data Parallel Execution

- Each *pagelet* is independent unit of rendering
- Same set of tasks executed for each pagelet (data), *in parallel*
 - CSS D/L, JS D/L, Render
- This is also a task “graph”
 - Sequence of tasks that execute one after the other





Data & Task Parallel Models

- Data parallel model is common in Clouds
 - E.g. Map-Reduce, Giraph
 - Helps exploit independent units of data on multiple machines/processors/threads
- Task parallel model help in composition
 - Allows tasks to operate concurrently on same/different data
- ILO3: Algorithms and Programming Patterns for Cloud Applications

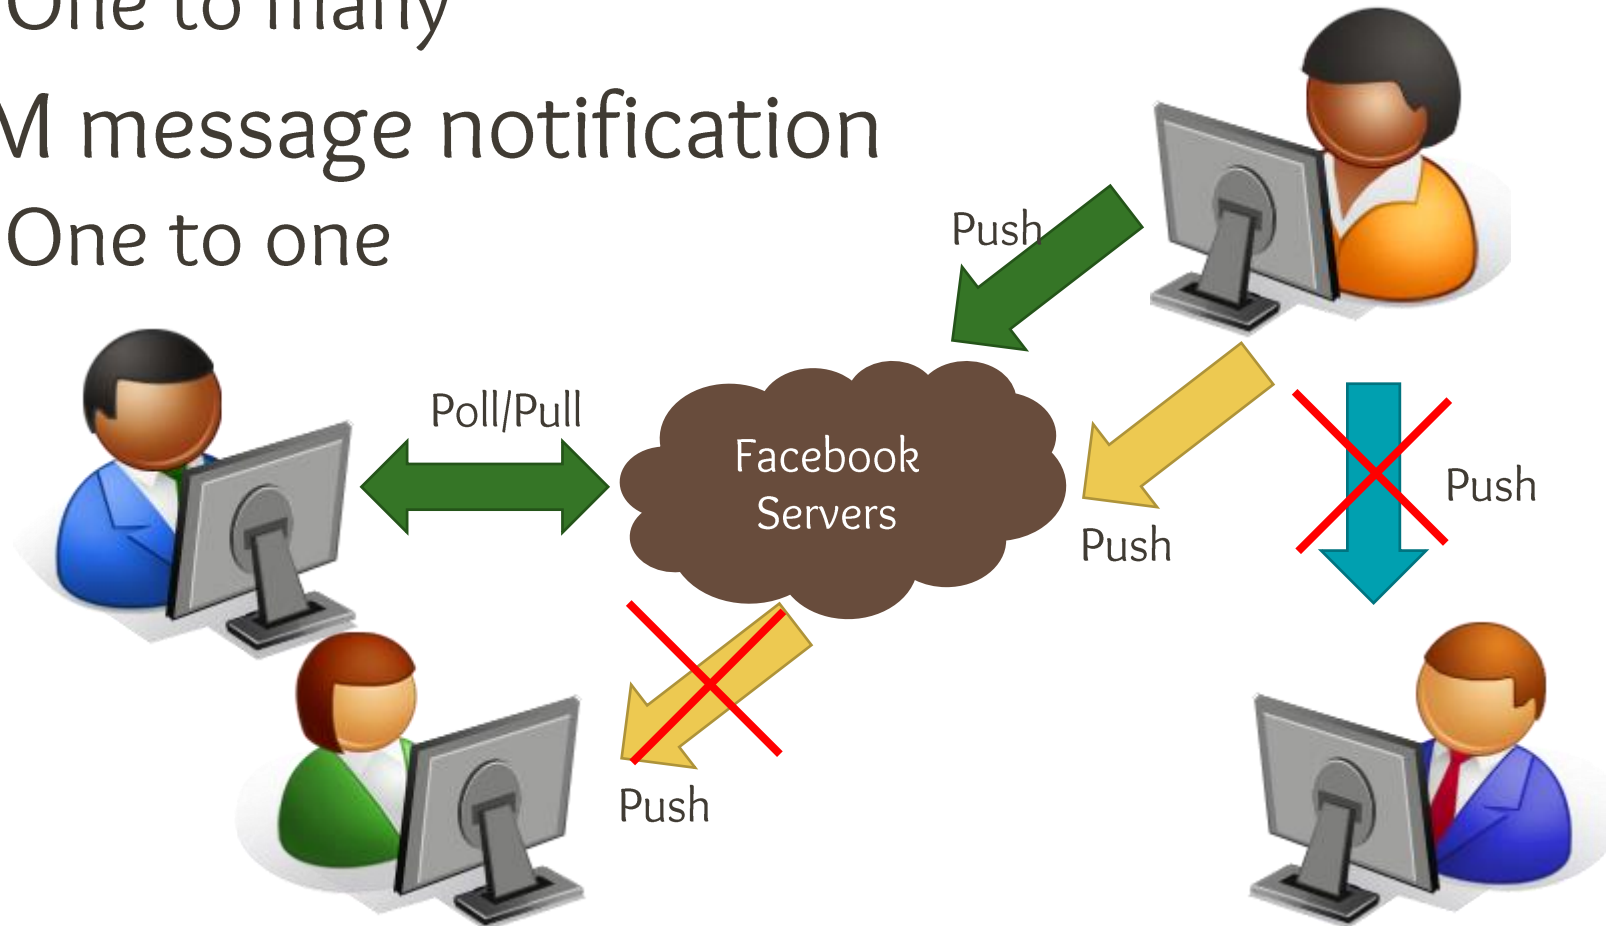


ILO 4: Application Execution Models on Clouds



Facebook (Browser) Chat

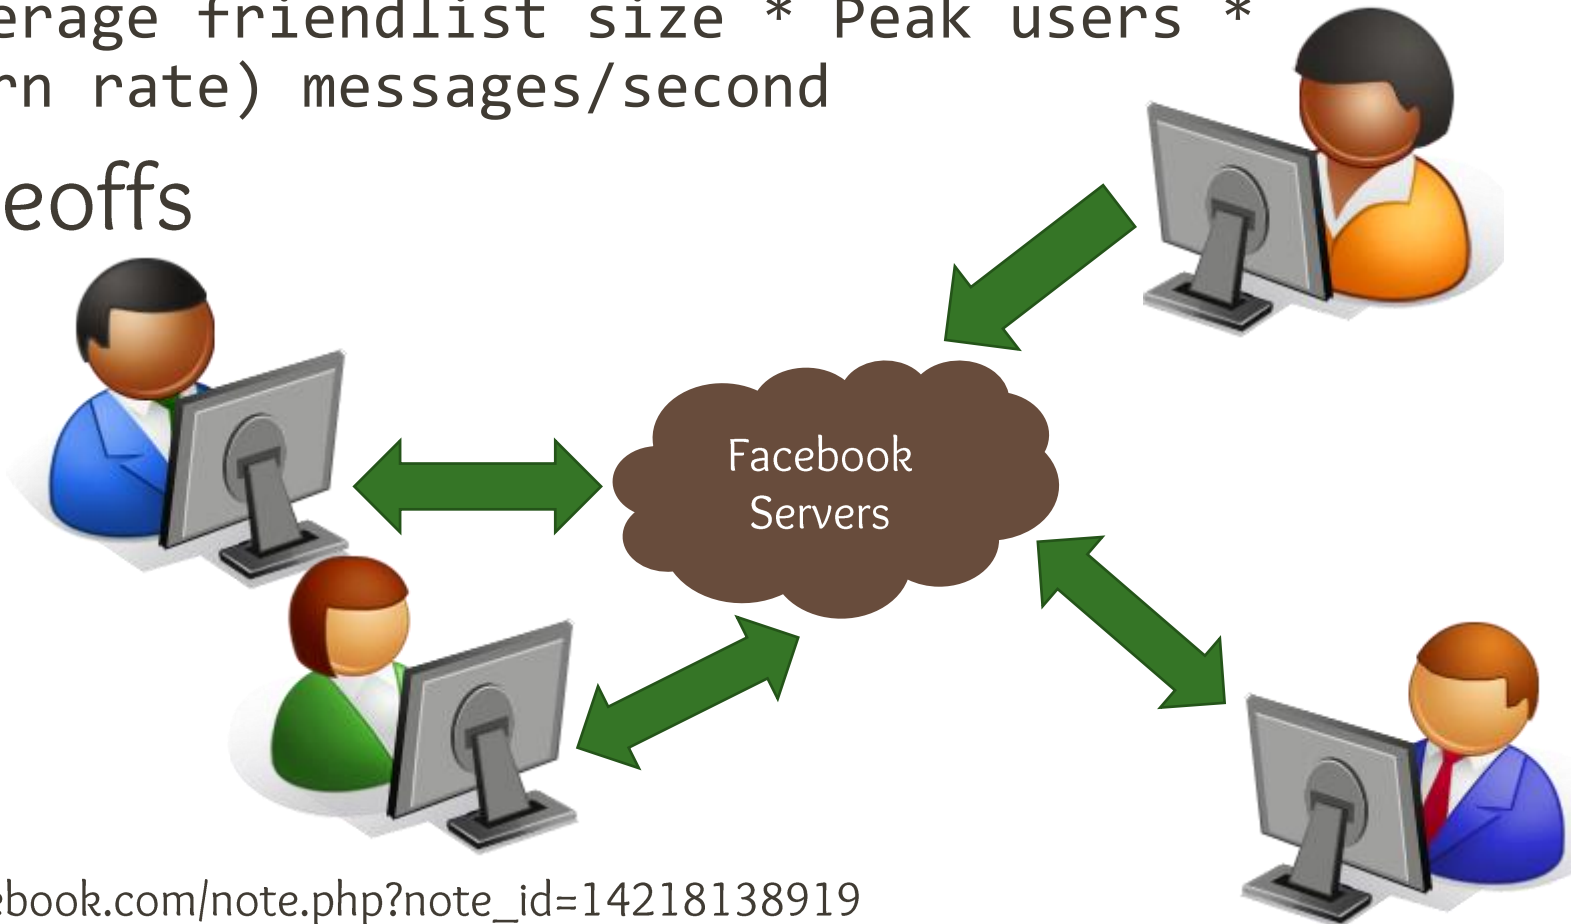
- Online/offline status
 - One to many
- IM message notification
 - One to one





Synchronous vs Asynchronous

- Immediately or eventually delivered?
 - Cost of sync goes up with # of users
 - (Average friendlist size * Peak users * Churn rate) messages/second
- Tradeoffs





Programming Cloud Apps

- Coordinate execution across VMs, clients
 - Synchronization of activities
 - Storage-based vs. In-memory operations
 - Trade-off between guarantees
-
- ILO 4: Application Execution Models on Clouds

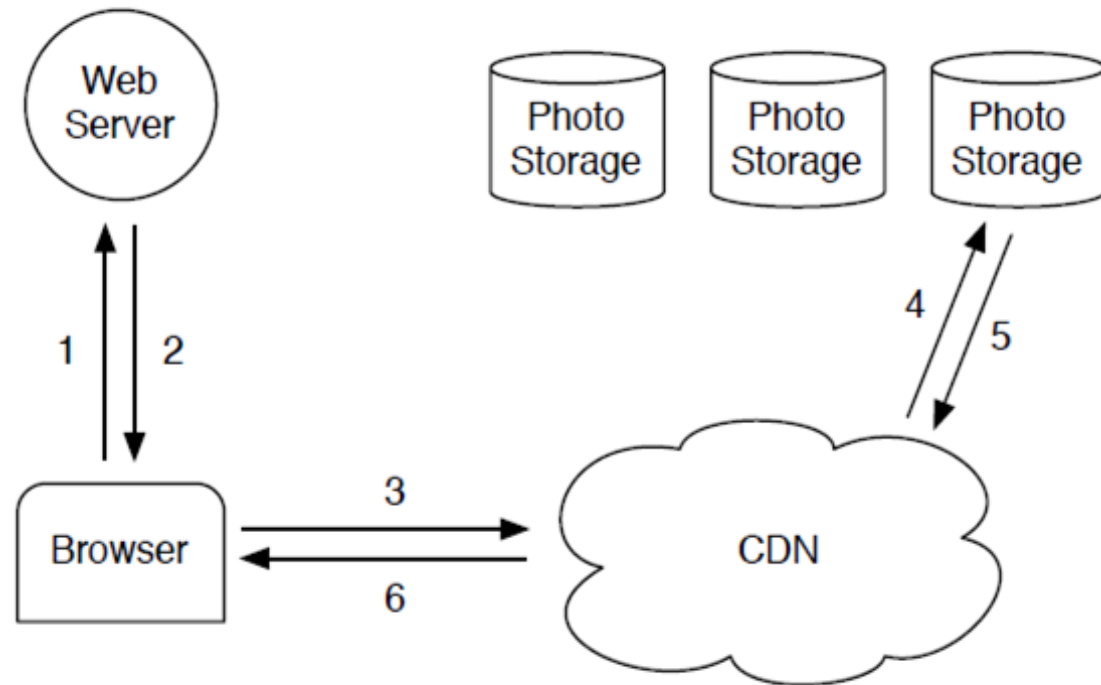


IL05: Performance, scalability & consistency on Clouds



Data Locality for Facebook Photos

- Content Distribution Networks (CDN)
 - Host “Hot” data, Spatially close to the client
 - E.g. Akamai
- Low latency
- Freshness?
- Consistency?

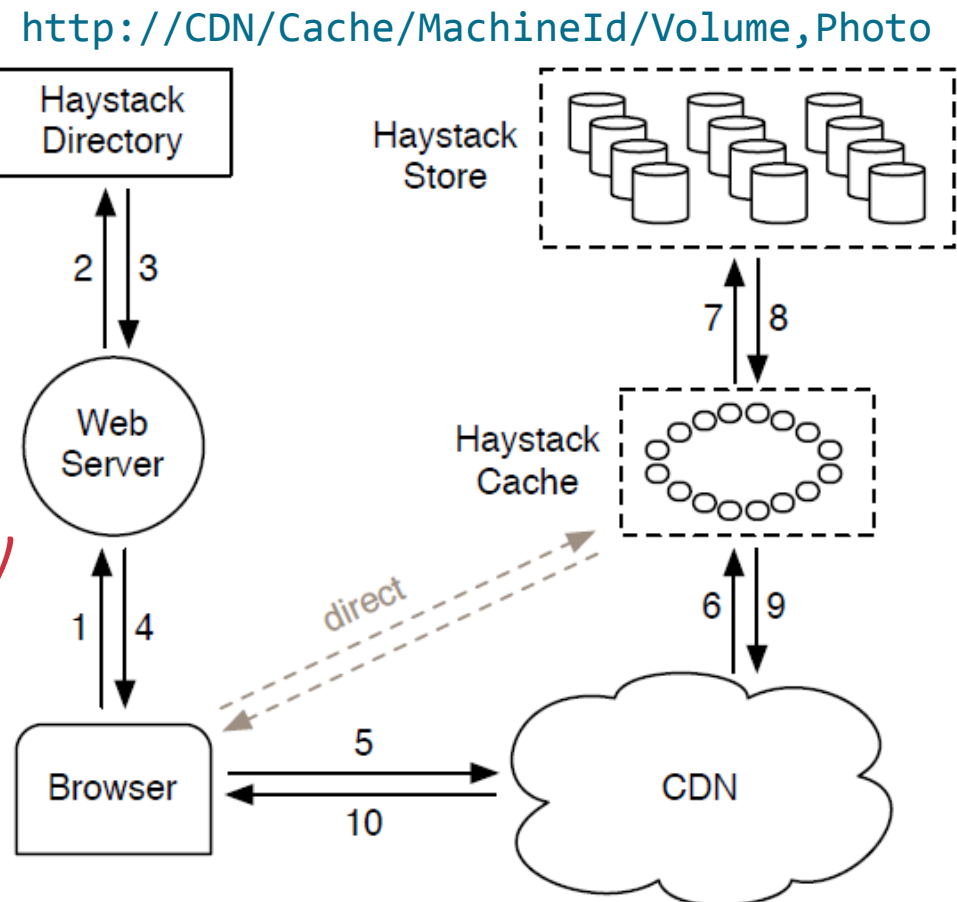




Performance, Availability, Consistency, Scalability, ...

- CDNs are costly for long tail
 - And poor cache hits
- Haystack
 - Incremental URLs
 - Better caching
 - Low latency response

- ILO5: Performance, scalability & consistency on Clouds





Bottomline

Cloud is ~~a revolution~~ **an important evolution**

- Yet another distributed system
- Novel tools, technologies, applications
- Builds on fundamental concepts
 - Distributed systems
 - Programming models & algorithms
- Different applications, reliability, performance trade-offs



Teaching and Learning Activities (TLA)

- **How** we'll go about meeting the ILO's
 1. Lectures
 2. Guest Lectures
 3. Homework
 4. Research Reading & Summarization
 5. Project
 6. Exam



Assessment (*1000 point scale*)

- **30% Homework** Three homework assignments (*100 points each*).
- **10% Research Summary** Reading and summary report on one research paper (*100 points*).
- **30% Project** One ungraded and two graded midterm (*100 points*) and final (*200 points*) project assignments.
- **25% Exams** One Mid-term (*100 points*) and one Final (*150 points*) exam.
- **5% Participation** Participation (i.e. not just "attendance") in classroom discussions and online forum for the course (*50 points*).



Academic Integrity

- Students must uphold [IISc's Academic Integrity guidelines](#). While these are common sense, it is helpful to review them since failure to follow them will lead to sanctions and penalties.
- **This includes a reduced or failing grade in the course. Severe cases of academic violations will be reported to the Institute and may lead to an expulsion.**



Academic Integrity

- Learning takes place both within and outside the class. Hence, **discussions** between students and reference to **online material** is encouraged as part of the course to achieve the intended learning objectives.
- However, while you may learn from any valid source, you **must form your own ideas** and **complete problems and assignments by yourself**.
- *All works submitted by the student as part of their academic assessment must be their own!*



Academic Integrity

- **Plagiarism:** *Verbatim* reproduction of material from external sources (web pages, books, papers, etc.) is not acceptable.
- If you are *paraphrasing* external content (or even your own prior work) or were otherwise *influenced* by them while completing your assignments, projects or exams, you must clearly acknowledge them.
- *When in doubt, add a citation!*



Academic Integrity

- **Cheating:** While *you may discuss* lecture topics and broad outlines of homework problems and projects with others, you cannot collaborate in completing the assignments, copy someone else's solution or falsify results.
- You cannot use notes or unauthorized resources during exams, or copy from others.
- The narrow exception to collaboration is between team-mates when competing the project, and even there, *the contribution of each team member* for each project assignment should be clearly documented.



Academic Integrity

- **Classroom Behaviour:** Ensure that the course atmosphere, both in the class, outside and on the online forum, is conducive for learning.
- *Participate* in discussions but *do not dominate* or be abusive. There are no “stupid” questions 😊 Be considerate of your fellow students and avoid disruptive behaviour.



Resources

- **Textbook:** *Distributed and Cloud Computing: From Parallel Processing to the Internet of Things*, Kai Hwang, Jack Dongarra and Geoffrey Fox, Morgan Kaufmann, 2011 ([Tata Book House](#))
 - Select sections from Part I and II (Ch. 1-6)
- **Official Mailing List:** se252.jan15@mailman.serc.iisc.in
 - Everyone MUST be on the list. All *official course announcements* will be posted on it.



How to use the mailing list

- Unmoderated
- Open for everyone to discuss
- Can suggest resources to others
- ***Do not post partial/full solutions!***
- Can request for official clarifications
- Official announcements will be posted
 - Your responsibility to check regularly



Teaching & Office Hours

- **Lecture:** TTh 2PM-330PM, SERC 202
 - *If you miss the lecture, its your responsibility to catch up from slides & talking to other students!*
- **Office Hours:** Fri, 4-5PM or by appointment (*send email*), SERC 411
 - May have project discussions at this time, if necessary



Tentative Schedule (*Subject to change...See webpage*)

- <http://www.serc.iisc.ernet.in/~simmhan/SE252/>