

DISTRIBUTED SYSTEMS

Instructor: Prasad Dewan (FB 150, dewan@unc.edu)



COURSE HOME PAGE



Comp 734: Distributed Systems

Course Overview

This course will provide an implementation-oriented study of distributed systems. Some of the topics covered will include inter-process communication, group communication, synchronization, remote procedure call, peer to peer and centralized sessions, fire-walls, causal broadcast, atomic broadcast, scalability, fault tolerance, replication, and transactions/concurrency control. These are foundational concepts, which are becoming particularly relevant with the emerging areas of cloud computing and distributed games. These concepts will be introduced as layers in a general distributed infrastructure. Your projects will implement new layers and provide alternative implementations of some of the existing layers. When implementing a layer, you will act both as an application programmer, using abstractions of the layers below, and a systems programmer, defining abstractions for the layers above. The number of lines of code required by each layer will be relatively small; however the compositions of these layers will be complex.

The main difference between this course and a distributed systems course is that it will address the design and implementation

<http://www.cs.unc.edu/~dewan/734/current/index.html>



LECTURES AND ASSIGNMENTS

Schedule (Tentative)

	Unit (Start Date)	Slides	Chapters	Assignment
	Introduction	PowerPoint 2007		
	Threads and Thread Coordination (Read on your own)	PowerPoint 2007		
	Java Non-Blocking Socket Channel I/O			Distributed Non-Blocking Halloween Simulation
	Java Remote Method Invocation			Distributed RMI-based Halloween Simulation
	Sync replicated Objects			Replicated Sync-based Halloween Simulation
No book				
	GIPC byte communication (Oct 9, 10)	PowerPoint 2007		Socket-based GIPC
PPT slides and sometimes word doc				
	RPC driver (Oct 24, 26, 31, Nov 2, 7)	PowerPoint 2007		
Current assignment is on the web - start working ASAP on it				
				Extendible Multi-Platform Serialization
Outline of other assignments given				
				Synchronous Receive, Procedure and Function Call
	GIPC P2P	PowerPoint 2007		



SOFTWARE

Downloads

Beau Halloween Simulation (Library, keep it compressed)	beau_project.zip
Coupled Halloween Simulations (Eclipse project, uncompress and link to libraries)	CoupledTrickOrTreat.zip
ObjectEditor (Library)	oeall17.jar
GIPC	

Software to be continuously updated



GRADE DISTRIBUTION

Exams (Two midterms, no final)	40%
Assignments (Home work)	60%
Fudge Factor (Class participation, other factors)	10%



GETTING HELP

Can discuss solutions with each other at a high level

Not at the code level

Sharing of code is honor code violation

Can help each other with debugging as long as it does not lead to code sharing

Assignments may contain solution in English (read only if stuck)



PIAZZA

Getting Help and Class Discussion

We will be using Piazza for class discussion and getting help. The system is highly catered to getting you help fast and efficiently from classmates, the TA, and myself. Rather than emailing questions to the teaching staff, I encourage you to post your questions on Piazza. If you do not get a response within a day or two on Piazza, please send mail to help401@cs.unc.edu. But try Piazza first. Do not send mail to an individual instructor, as that can overwhelm him - such mail will be ignored.

Before posing a question, please check if this question has been asked before. This will reduce post clutter and reduce our burden. Repeat questions will be ignored by the instructors.

Piazza allows anyone to respond. So if you see a question that you think you can respond to, please do so, as that will reduce our burden and help you "teach" your fellow students.

This will be a form of class participation that will be noted when I allocate my fudge points!

Hope it works well

If you have any problems or feedback for the developers, email team@piazza.com.

Find our class page at: <https://piazza.com/unc/fall2013/comp734>



DISTRIBUTED PROGRAM?

A program “involving” multiple computers

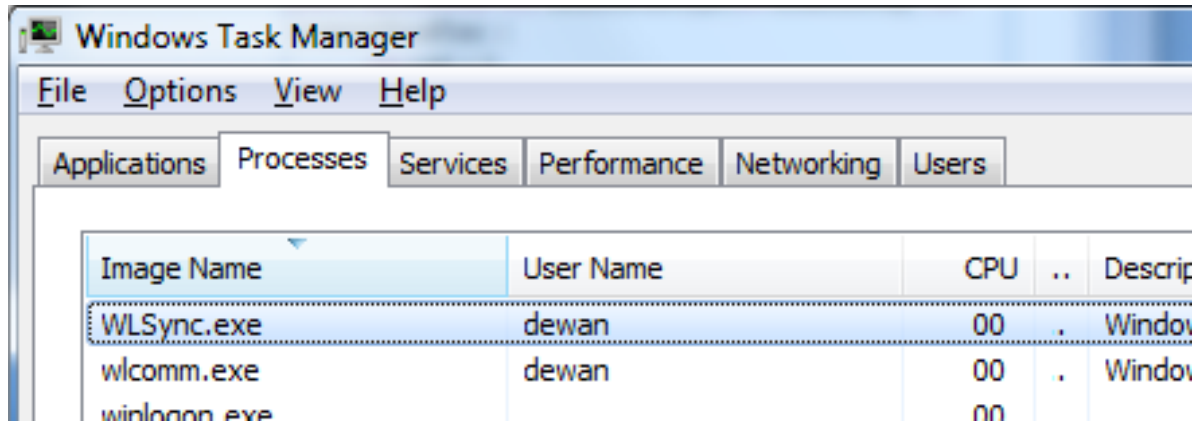
Specific computers must be bound at run time

→ Program can run on a single computer

Definition involves processes



PROGRAM VS. PROCESS VS. THREAD




Program

```
public class AP2PAliceTOTSimulation {  
    public static String ID = "9100";  
    public static String NAME = "Alice";  
    public static int USER_NUMBER = 0;  
    public static void main (String[] args) {  
        Tracer.showInfo(true);  
        AP2POTSessionClientCreator.createP2PDelay
```

Process is execution instance of program, associated with program and memory

Same program can result in multiple processes

Execution instance

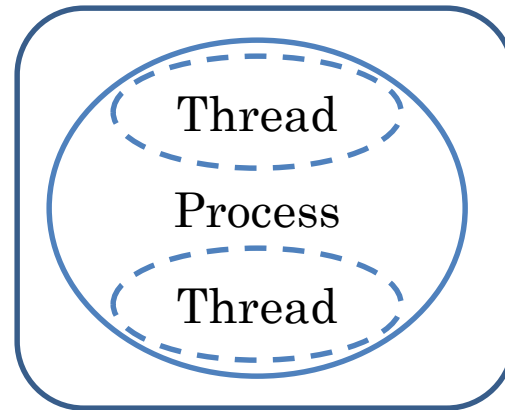
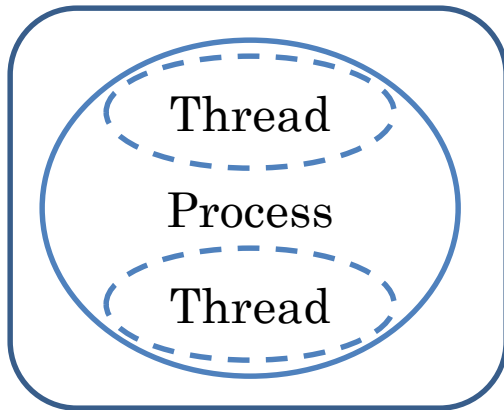


Processes are independent activities that can interleave or execute concurrently


Thread is also an independent activity, but within a process, associated with a process and a stack



DISTRIBUTION OF PROCESSES/THREADS




Different processes can execute on different (distributed) computers



A single process executes on one machine



DISTRIBUTED PROGRAM



Exec

Connected process pair : Some computation of a process can be influenced by or influence computation of the other process

Connected process group: each process is coupled to at least one other process in the group

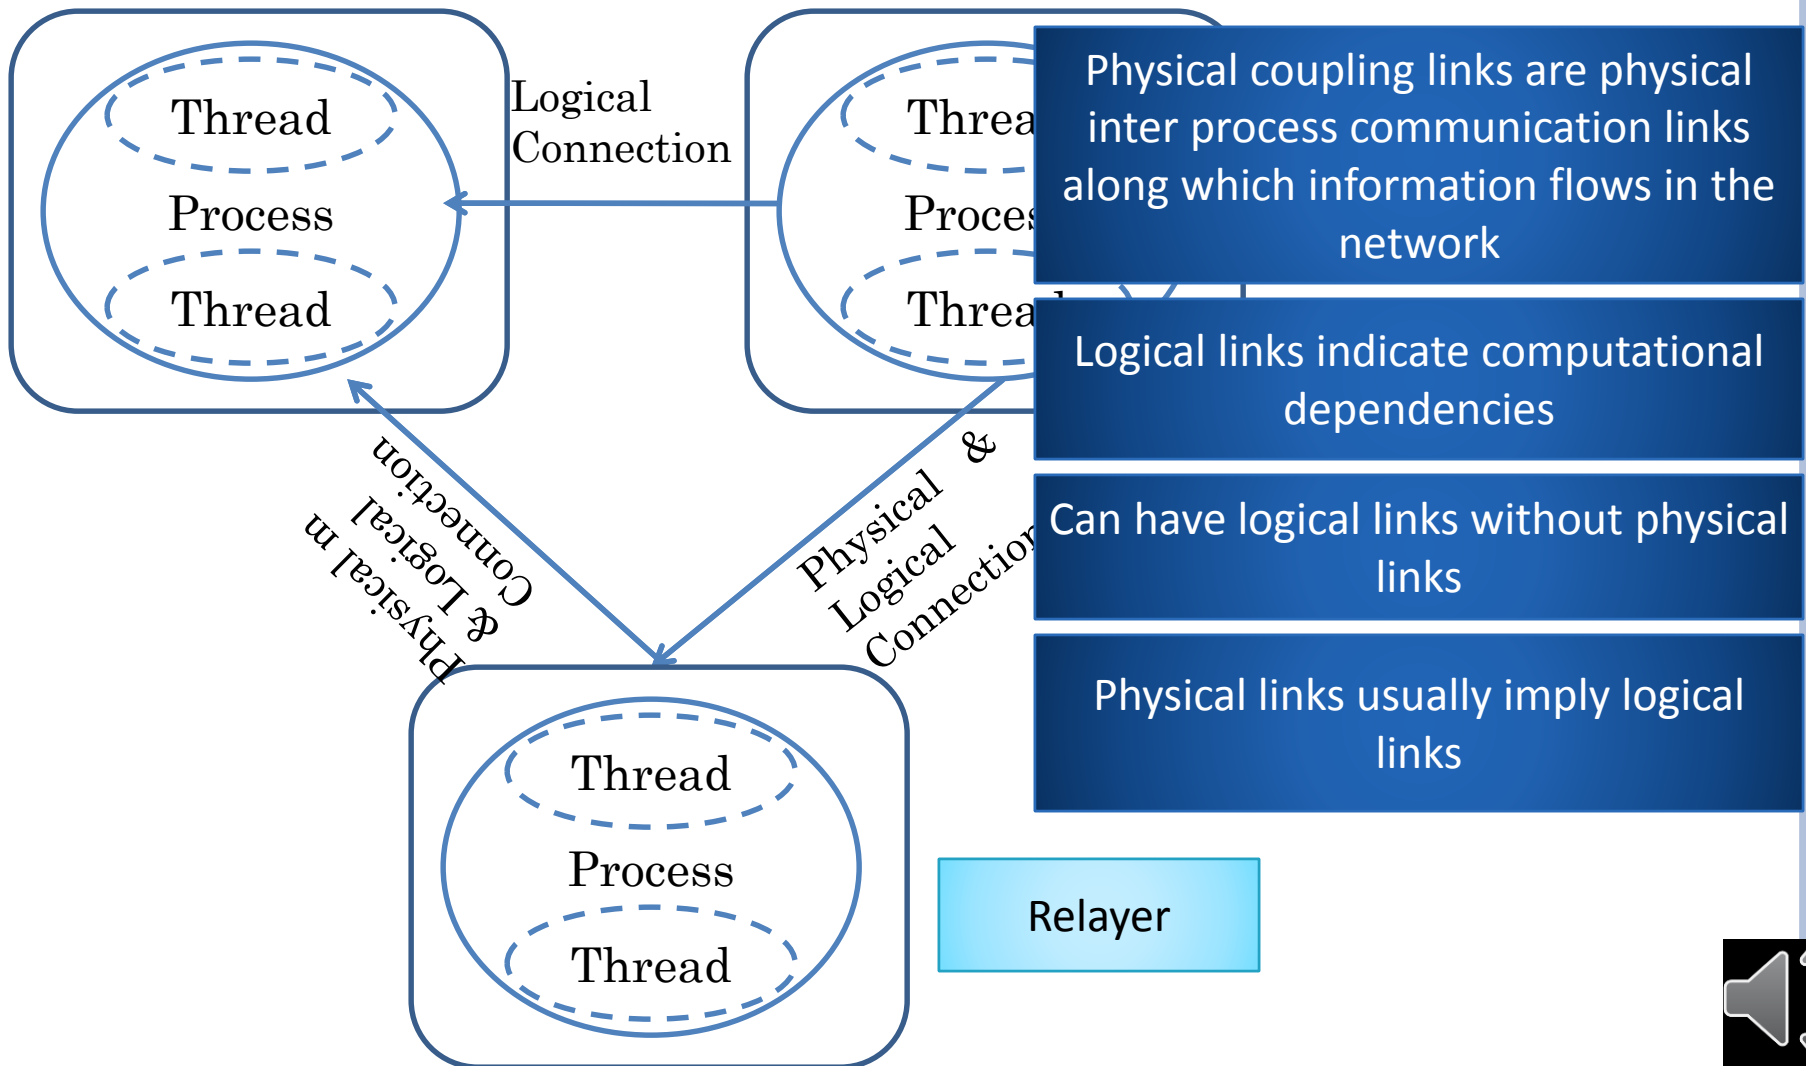
```
public  
    public static String ID = "9100";  
    public static String NAME = "Alice";  
    public static int USER_NUMBER = 0;  
    public static void main (String[] args) {  
        Tracer.showInfo(true);  
        AP2PTOTSessionsClientCreator.createP2P()  
    }  
}
```

Graph crated by creating pair-wise dependency links is not partitioned—every node reachable from every other node

```
public static String ID = "9100";  
public static String NAME = "Alice";  
public static int USER_NUMBER = 0;  
    main (String[] args) {  
        Tracer.showInfo(true);  
        AP2PTOTSessionsClientCreator.createP2P()  
    }  
}
```



LOGICAL VS. PHYSICAL INTER PROCESS CONNECTION LINKS



DISTRIBUTED APPLICATIONS

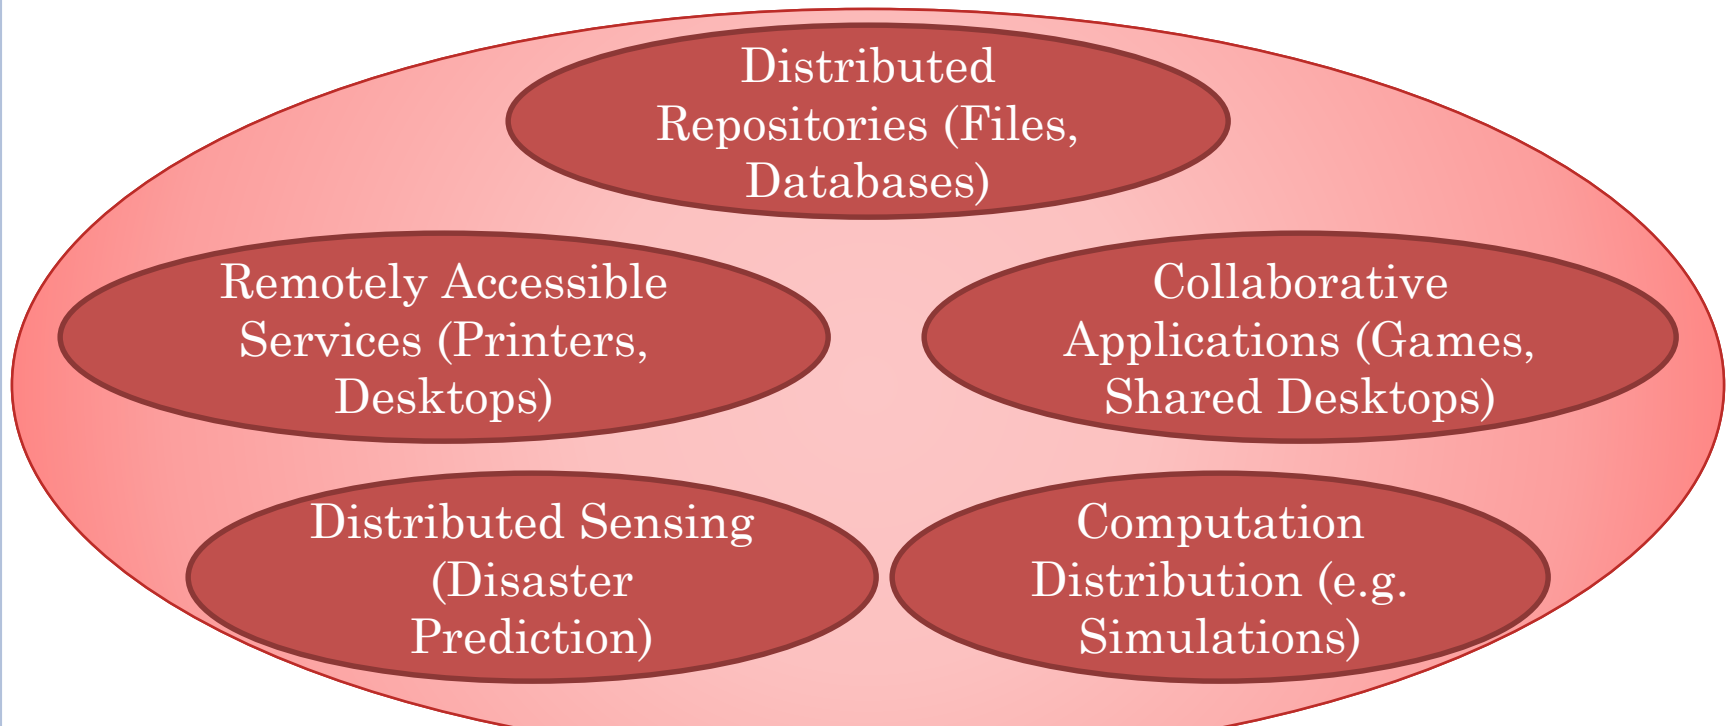
Distributed applications?

Non distributed applications?

In today's world, what is or should not be distributed?



SOME DISTRIBUTED DOMAINS



Full courses on some of these areas, with concepts specific to them
(Distributed Databases, Collaborative Applications)

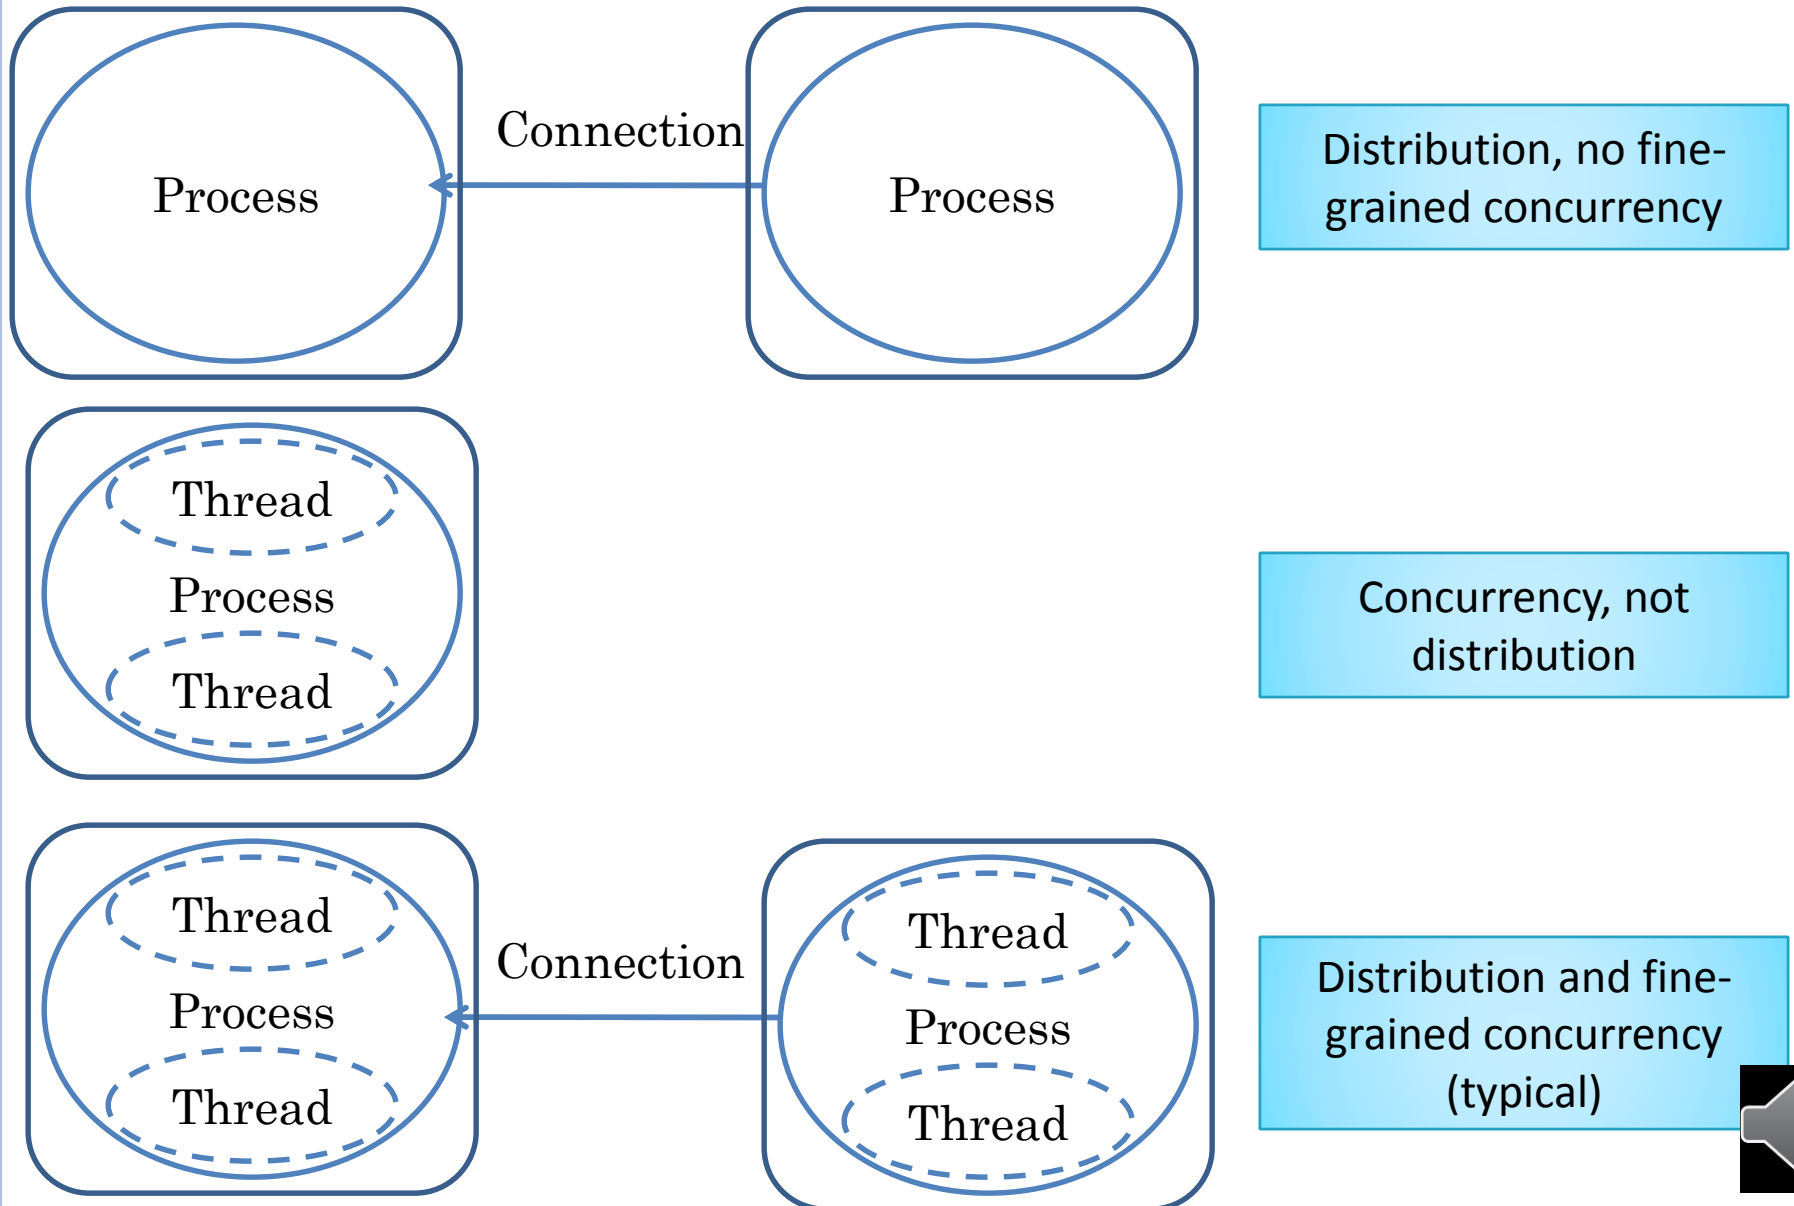
Will look at domain-independent concepts at the intersection of them

Will not take an application-centric view

Fundamental Issues?



DISTRIBUTION VS. CONCURRENCY



NON-DISTRIBUTED VS. DISTRIBUTED PROGRAM

Non-Distributed

Creates a single process logically and physically unconnected to any other process

Must deal with sequential and possibly concurrency issues

Distributed

Creates a pair or larger group of connected processes

Must also deal with distribution and usually concurrency issues



SYSTEMS VIEWPOINT

System

Computer abstractions to implement some class of programs

Operating System

Processes, Files, Memory Management , Threads...,

Database Management System

Query Language, Transactions, ...

Programming Languages

Arrays, Loops, Classes, ...

Distributed Systems

Data Communication, Remote Procedure Call (RPC), ...

Byte/object communication consists of byte/object of exchange

RPC assumes communication consists of procedure requests and return value responses



DISTRIBUTED SYSTEMS

Study of design and/or implementation of computer abstractions for developing distributed programs

Why distributed systems?

Why systems?

Alternatives to understand how to program some domain of applications?

Non distributed programs?



ALTERNATIVES TO UNDERSTANDING

Programming: Abstraction use

Programming: Use of a specific set of non distributed abstractions (e.g. , functional, MATLAB programming)

Distributed Programming : Use of a set of distributed abstractions (e.g. Socket/RPC Programming)

Systems: Abstraction design and/or implementation

Design and implementation of non distribution abstractions (Object-Oriented vs. Functional Languages, Compilers/Interpreters)

Design and implementation of distributed system abstractions (e.g. Data Communication /RPC Design and/or Implementation)

Theory: Models and algorithms

Non distributed model and algorithms (Turing Machines, HeapSort,)

Distributed Models and Algorithms(e.g. ? Phase commit, Group Comm. Model)



RATIONALE

Abstraction Design vs. Implémentation

Abstraction design linked to implementation: Designs are done of only efficiently implementable abstractions

Abstractions vs. Theory (Models, Algorithms)

Abstractions are implemented operational models and have (the more) practical algorithms in them

Abstraction Design & Implementation vs. Use

Maturity with design and implementation issues allows you to better understand the semantics of a specific abstraction.

Abstract implementations require advanced programming/ software engineering techniques– “you cant really program if you have not written a compiler”



TEACHING ABSTRACTION DESIGN & IMPLEMENTATION?

Lectures address design; assignments, implementation (e.g. Implement a PL interpreter in another PL)

Implementations can be complex and need instruction

Lectures give high-level pseudo code for complex algorithms; assignments full implementation (e.g. compilers)


Pain/gain ratio high, semester barely enough time for compiler

Lectures discuss code for a system of abstractions : assignments extend/modify this code

Code must be understandable and ideally also elegant



THE XINU APPROACH TO TEACHING OS



Layering

Interrupt Management

Thread Communication

Thread Synchronization

Thread Management

Approach not used in distributed computing

Need distributed system layers

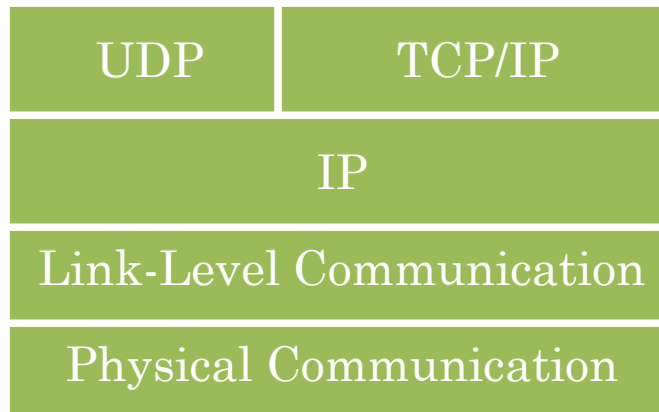
Reuse of previous layers keeps code short (and hence presentable in class)

Can unravel a system in stages to a class

Layering good for software engineering as well as pedagogical reasons



LAYERS EXIST IN NETWORKING

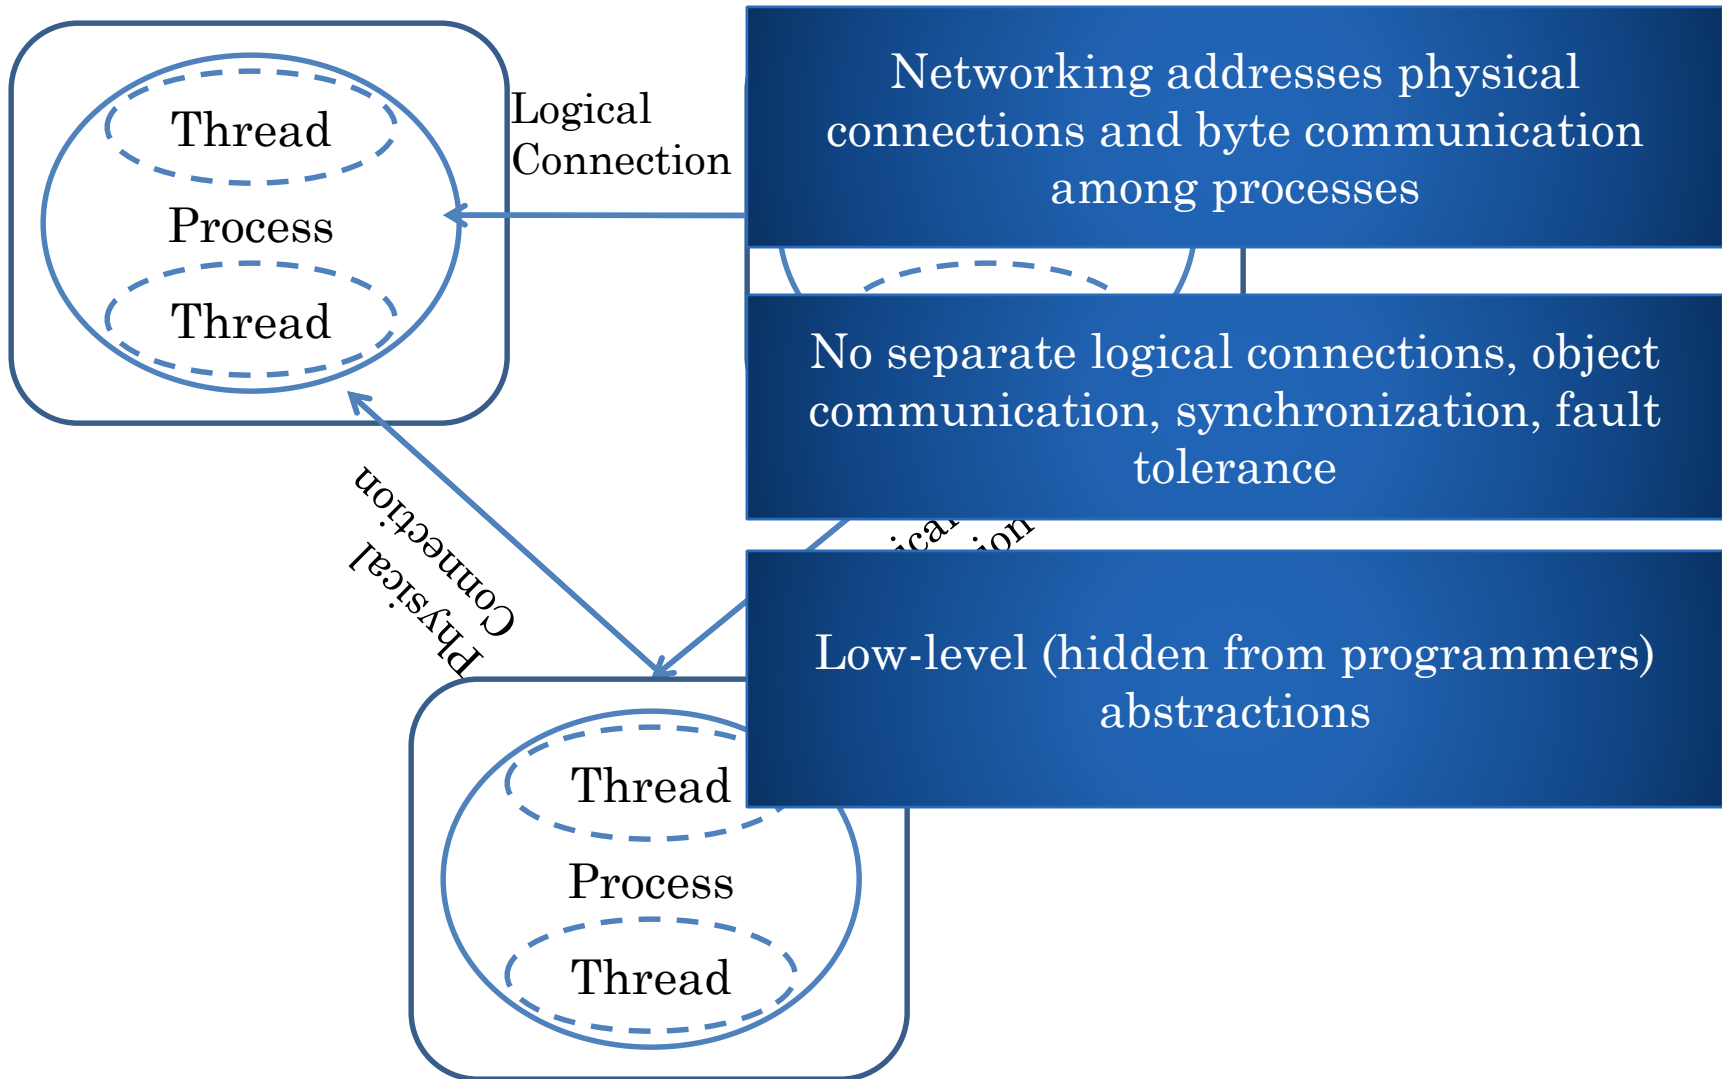


Physical communication in networking involves machines and used hardware machine addresses

Physical communication in distributed systems is between processes and indicates routing of information among processes



DISTRIBUTED VS. NETWORK LAYERS



DISTRIBUTED SYSTEM VS. NETWORKING ABSTRACTIONS

Programming Language
Abstractions

Assembly Language
Abstractions

Distributed Abstractions

Networked Abstractions

Distributed Abstractions

OS Byte Communication API

Just as programming language abstractions are built on top of assembly language abstractions

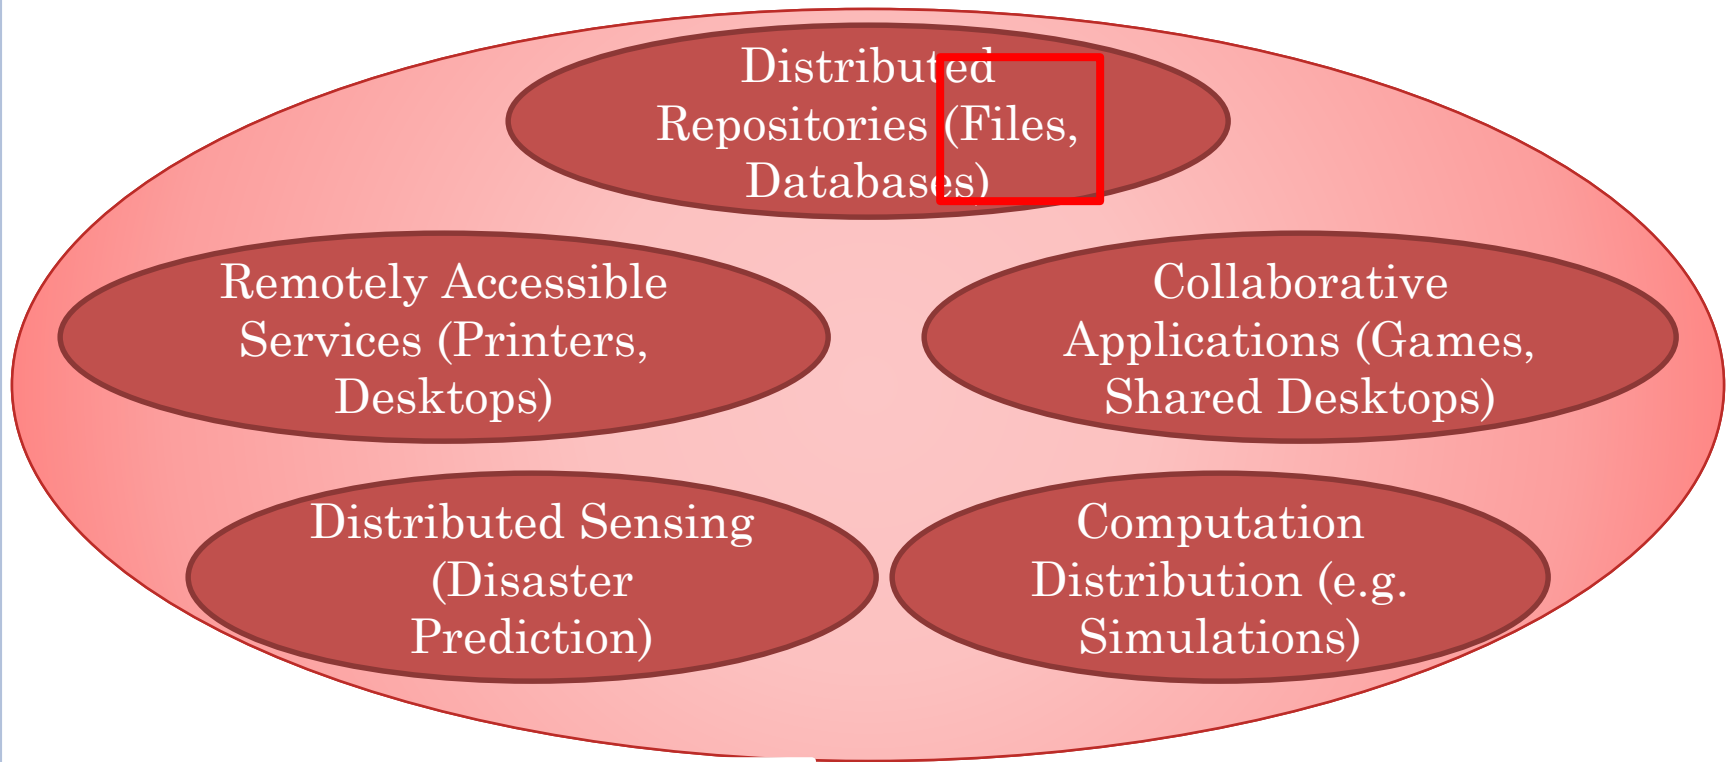
Distributed system abstractions are built on top of networked abstractions

Knowledge of assembly/networked abstractions important to implement PL/distributed abstractions

Byte communication APIs, close to networked abstractions, is provided by operating systems (e.g. sockets), which hide networking abstractions



DOMAIN INDEPENDENT?



Distributed Abstractions

OS Byte Communication API

Will look at domain-independent concepts at the intersection of them

Even though OS abstractions developed to build distributed OS (file systems), they are by definition domain-independent



LANGUAGE VS. OS ABSTRACTIONS

Both operating systems and programming languages provide domain-independent abstractions

Operating systems support processes and language-independent abstractions for accessing protected info and sharing information among processes (files, IPC)

Programming languages must provide fine-grained abstractions needed within a process

They also provide an interface to OS abstractions through libraries or language constructs

They can also extend the OS abstractions (e.g. typed files)



LANGUAGE VS. OS, DISTRIBUTED ABSTRACTIONS

Byte communication is all that operating systems provide

Non distributed programming languages such as C provide only OS abstractions

Distributed programming languages such as Java provide a richer variety of abstractions

Java provides threads and reflection, making it easy to implement our own replacements and extensions of Java abstractions

Will use Java as implementation language

To extend and replace Java abstractions/layers, knowledge of them useful

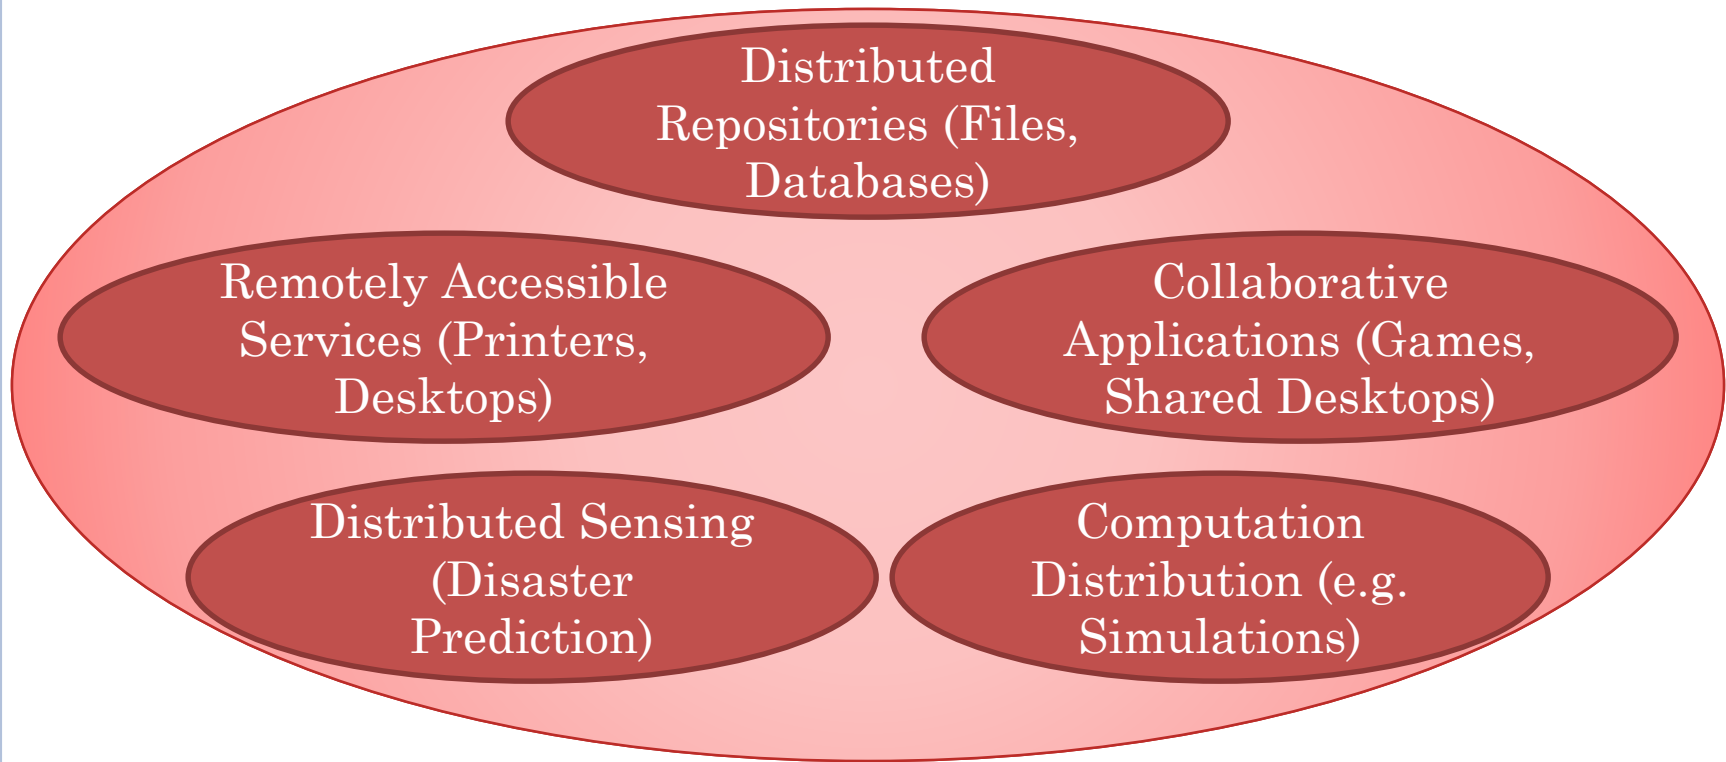


OBJECTIVES

At the end of the course you will



DISTRIBUTED COMPUTING



Internet/Cloud computing increasing relevance of the fundamental concepts



PRACTICAL RELEVANCE

For distributed applications, likely to use the code you implemented than existing abstractions

Can send objects over NIO socket channels

Existing Java RPC does not work on Android devices, but the one you implement will

Use Sync, which apparently is the basis of some new Mobile platforms

Will implement many abstractions not part of standard Java



SOFTWARE ENGINEERING PRINCIPLES

Classes

Existing classes will be used, inherited but not modified directly

Interfaces

Alternative implementations will create new classes implementing existing interfaces

Factories and
Abstract factories

These will allow easy switching between different implementations

Generics

Implementation rather than use of generics to unite buffer and object communication

Will be both a distributed computing and software engineering course



RELEVANCE TO OS

Inter-process communication key to design of new OS's, even non distributed OS

Extensive use of bounded buffers

Will study and use thread synchronization in depth

Will study how distributed OS are implemented

Will gain understanding of fundamental OS concepts except memory management



INTRODUCTION TO SYSTEMS

Systems: Abstraction design and implementation

Design and implementation of non distribution abstractions (Object-Oriented vs. Functional Languages, Compilers/Interpreters)

Design and implementation of distributed system abstractions (e.g. Data Communication /RPC Design and/or Implementation)

Distributed systems covers concepts from many fields

