

Future Trends in Computing

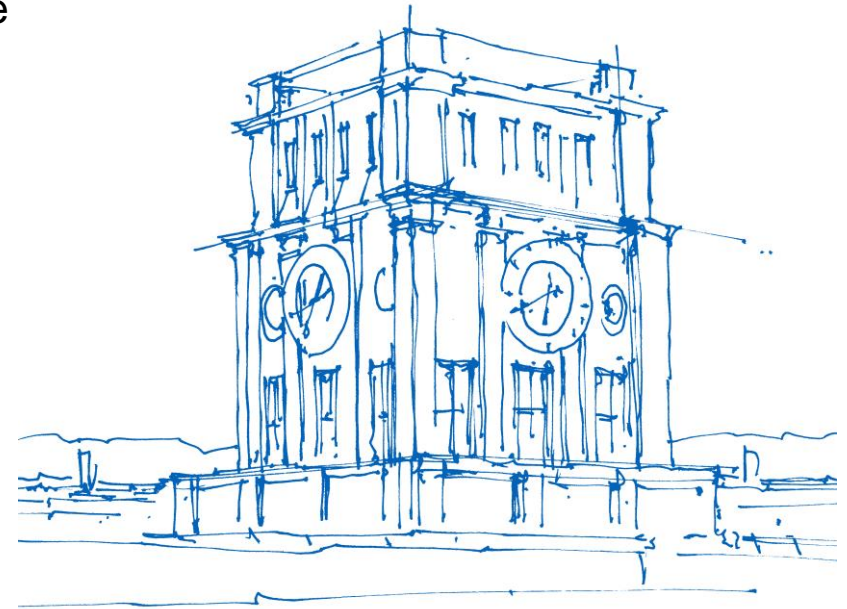
X10 Introduction

Alexander Pöppel, Emily Ao Mo-Hellenbrand

Technical University of Munich

Department of Informatics

Chair of Scientific Computing in Computer Science



Uhrenturm der TUM

Schedule

Week	Event
1 st (now)	X10 introduction
2 nd – 9 th	Biweekly lecture + Assignments
10 th – 13 th	Larger Assignment (Project Phase)
14 th	Final Presentation

Lecture Phase:

- Lecture (2nd, 4th, 6th, 8th week) and Office Hours alternating
- Time: Mon 14:00-15:30
- Room: MI 02.07.023

Goal:

- Implement parallel solver for Shallow Water Equations
- Get acquainted with novel programming concept for parallel architectures

X10 Intro Days Agenda:

Day 1

10:00 Setup & Hello world

11:00 Part 1: Basic syntax introduction

Exercise: sequential Quick Sort

13:00 Lunch Break

14:00 Part 2: Parallel features

Exercise: Parallel Monte Carlo Pi

Exercise: Parallel Heat Simulation

Day 2

10:00 Part 3: PGAS features

Exercise: Distributed Monte Carlo Pi

Exercise: Distributed Sort

12:00 Lunch Break

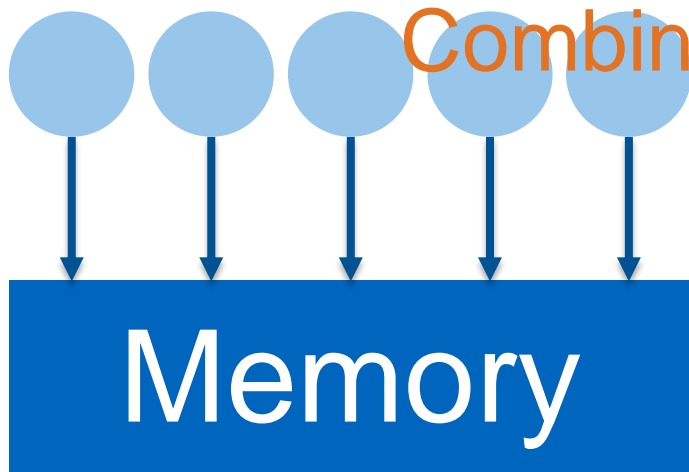
13:00 Part 4: Advanced features

Exercise: Distributed Heat Simulation

Introduction

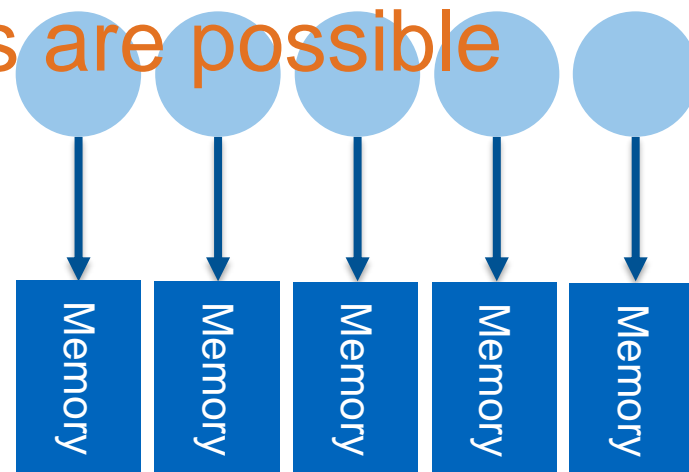
Shared memory parallelism

All threads in application share the same memory address space (One process).
Communication is done using shared objects.



Distributed Memory Parallelism

One process per concurrently executed unit.
Potentially executed on different machines,
with communication using active messages.



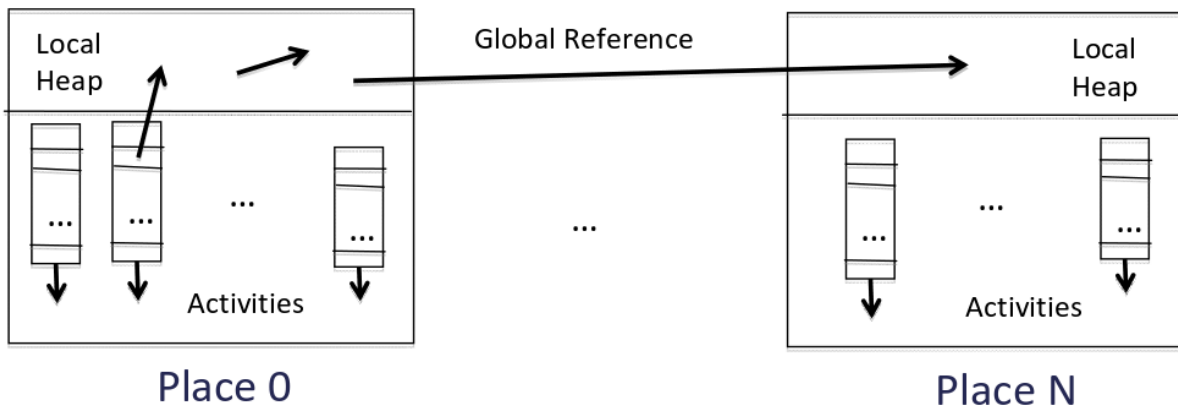
Combinations are possible

Introduction

Partitioned Global Address Space

Idea: Use one large address space (potentially spanning multiple machines) and partition it into smaller address spaces.

- Every object resides in one partition
- Copies of objects may be sent over the network
- Referenced to objects in other address spaces may be held.



X10

- X10 is such an (A)PGAS-oriented language
- Developed by IBM, started as a research project in ~2006
- Borrows elements from Java, Scala and others...
- Object-oriented, Functional Features
- Parallel:
 - Scales across multiple nodes in a cluster (Distributed Memory)
 - Able to utilize intra-cluster parallelism. (Shared Memory)
 - Can use heterogeneous resources (CUDA)

X10 – Get it!

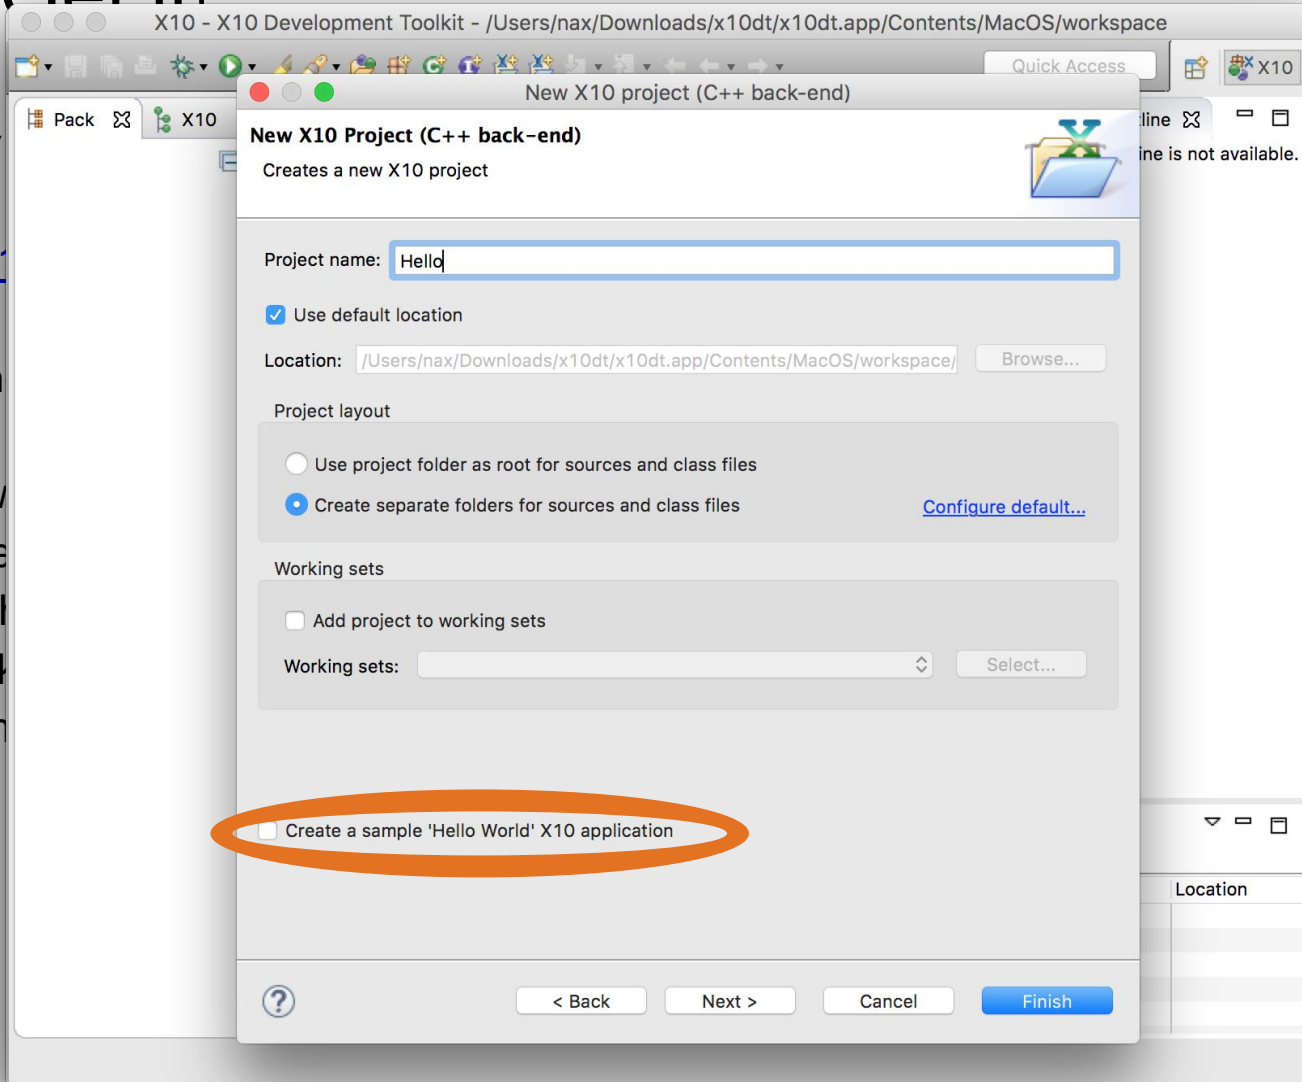
X10 is freely

<http://www.x10.org>

Download the

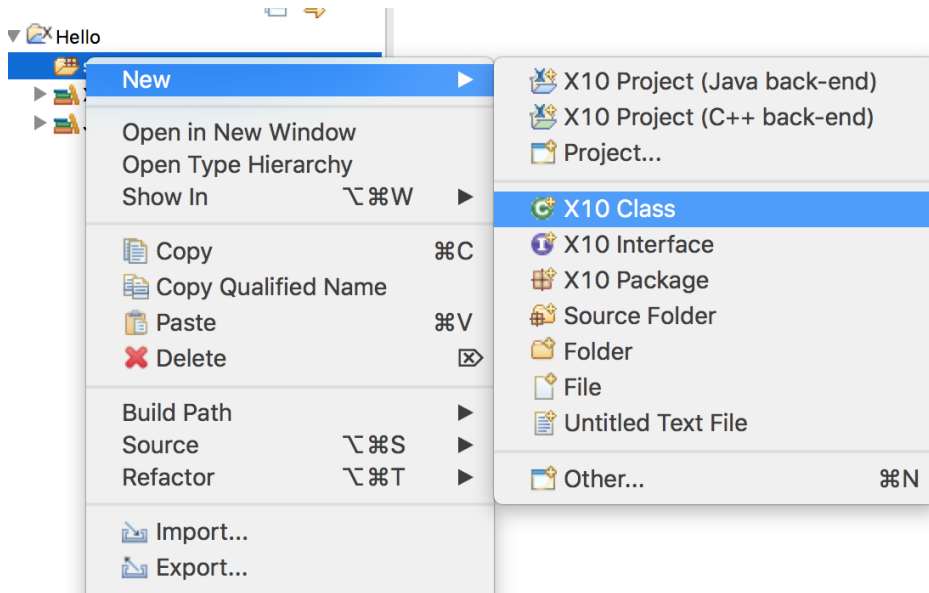
Create a new

- There are
- Choose the
- Both back
mileage n



, your

X10 – Hello World



```
public class HelloWorld {
```

```
    public static def main(Rail[String]) {  
        Console.OUT.println("Hallo Lab");  
    }  
}
```

```
}
```