

Introduction to Distributed Computing

(with hands-on Squeak demo for K-12)

Supercomputing 2005

Randy Heiland
Scientific Data Analysis Lab
heiland@indiana.edu



pervasivet^echnologylabs

AT INDIANA UNIVERSITY

What is Distributed Computing?

- ◆ **Wikipedia: “the coordinated use of physically distributed computers”**


We’re going to assume you know what a computer is, so we’ll try to explain what is meant by “physically distributed” and “coordinated use”.



“physically distributed” computers

Q: What distinguishes two (or more) coordinated computers being:

- ◆ **within a room?**
- ◆ **within a building?**
- ◆ **within a State | continent | planet?**



Answer: The Network!

- ◆ “coordinated” computers implies some sort of communication – electronic, optical, or telepathic 😊
- ◆ Communication between two computers separated by an ocean will (probably) take longer than two computers within a room
- ◆ The underlying application – the computation – will determine how important communication needs to be



Communication models

- ◆ Tightly-coupled
 - ◆ processors need to exchange data frequently
- ◆ Loosely-coupled
 - ◆ processors exchange data infrequently
- ◆ Uncoupled
 - ◆ processors work independently



Decomposing a problem

- ◆ How do we decompose (partition, break-up) a problem so that we can take advantage of distributed computing?
- ◆ Two basic strategies:
 - ◆ Data decomposition
 - Have each computer perform the same calculation, but just on its particular chunk of data
 - ◆ Task decomposition
 - Have each computer perform different tasks/calculations (perhaps on very same data)



Parallel vs. Distributed vs. Grid

Although an oversimplification, let's describe these 3 computing paradigms as follows:

- ◆ Parallel computing – in single machine room; very fast network; homogeneous computers
- ◆ Distributed computing – “spatially” larger than Parallel; distributed memory; heterogeneous computers
- ◆ Grid computing – “spatially” larger than Dist' d; may involve dist' d parallel computers



Let's look at some
“successful” examples of
distributed computing
projects...



From the very large – SETI@home

- ◆ Search for Extraterrestrial Intelligence
- ◆ Started 1999



What is SETI@home?

SETI@home is a scientific experiment that uses Internet-connected computers in the Search for Extraterrestrial Intelligence (SETI). You can participate by running a free program that downloads and analyzes radio telescope data.

PARTICIPATE

Rules and policies
Create account
Download
BOINC
BOINC Wiki
Donate

ABOUT

Technical news
Server status
Bookstore
Science newsletters
Science links
Sponsors

COMMUNITY

Message boards
Questions & answers
Profiles
Teams
Web sites & IRC
Porting & optimization
Pictures & music
Translation

YOUR ACCOUNT

Login
Preferences
Certificate

STATISTICS

Top participants
Top computers
Top teams

Site search:



To the very small




Folding@home

distributed computing



Our goal: to understand protein folding, protein aggregation, and related diseases

What are proteins and why do they "fold"? **Proteins** are biology's workhorses -- its "**nanomachines**." Before proteins can carry out their biochemical function, they remarkably assemble themselves, or "**fold**." The process of protein folding, while critical and fundamental to virtually all of biology, remains a mystery. Moreover, perhaps not surprisingly, when proteins do not fold correctly (i.e. "misfold"), there can be serious effects, including many well known **diseases**, such as Alzheimer's, Mad Cow (BSE), CJD, ALS, Huntington's, Parkinson's disease, and many cancers and cancer-related syndromes.



Results from Folding@Home



For the mathematically-inclined

<http://www.mersenne.org/>



GIMPS
The Great Internet Mersenne Prime Search
Finding **8** World Record Primes!

$2^P - 1$
MAY BE
PRIME!



Previous 3 successful examples

- ◆ Not to trivialize them, but they have been successful, in part, because each has a communication model which is basically uncoupled, i.e. each processor is able to independently work on a chunk of data
- ◆ Life isn't so easy with a tightly-coupled application



HOW do processors communicate?

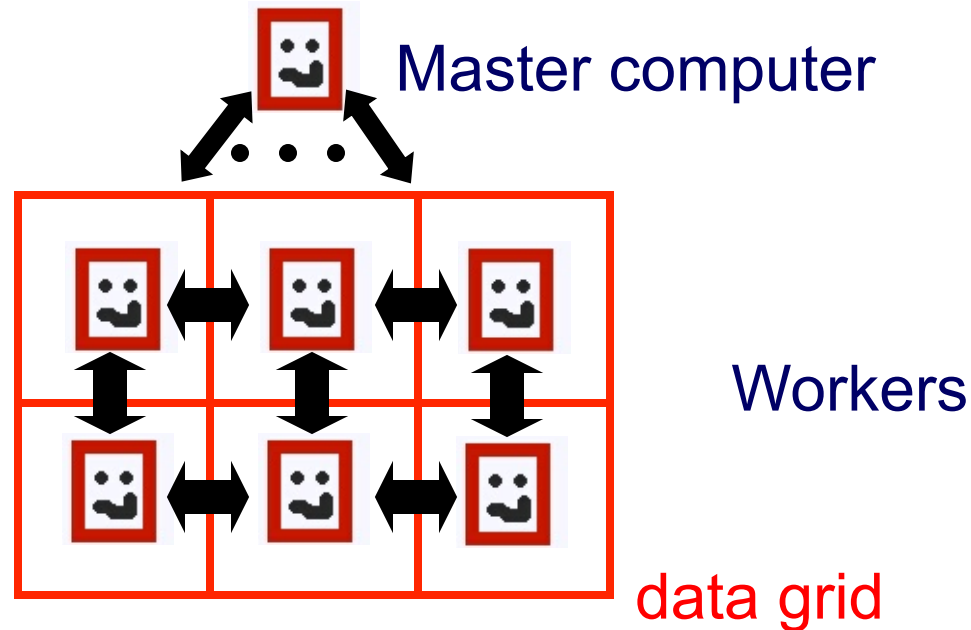
- ◆ Message-Passing Interface (MPI) – de facto standard for computational science applications.

Two widely used, open source implementations:

- ◆ MPICH
- ◆ LAM/MPI



Communication patterns



One common pattern is found in applications involving data decomposition. Here each processor handles its section of data, communicating with its neighbors, as well as the “master”. There are many other patterns.



Advanced Topics

- ◆ Distributed Algorithms, Nancy A. Lynch
MIT OpenCourseWare → Electrical
Engineering and Computer Science



Enough lecture – let's play!

- ◆ Let's explore some simulated parallel/distributed computing concepts (and have some fun) using a freely available software package called Squeak.



Targeting K-12: squeakland.org

- ◆ Free/open source
- ◆ “Media authoring” tool
- ◆ Fun, interactive
- ◆ Math & science concepts through simulations
- ◆ Alan Kay – “father of the personal computer”




If you don't have Squeak installed, you can get it from squeakland.org



Startup screen:



Click on Navigator tab to open it
then click on the
paintbrush icon



Painting/drawing palette



Paint brush, fill bucket,
eye-dropper(color selector), eraser

Brush size



color selection

shapes



Simulated distributed computing

- ◆ Create a ‘computer’
- ◆ Make copies of it
- ◆ Have our ‘distributed computer system’ do something interesting



Create a 'computer' :



2) Keep it

3) Rename it: "p0"
(for "processor #0")



1) Draw a picture of
your 'computer'

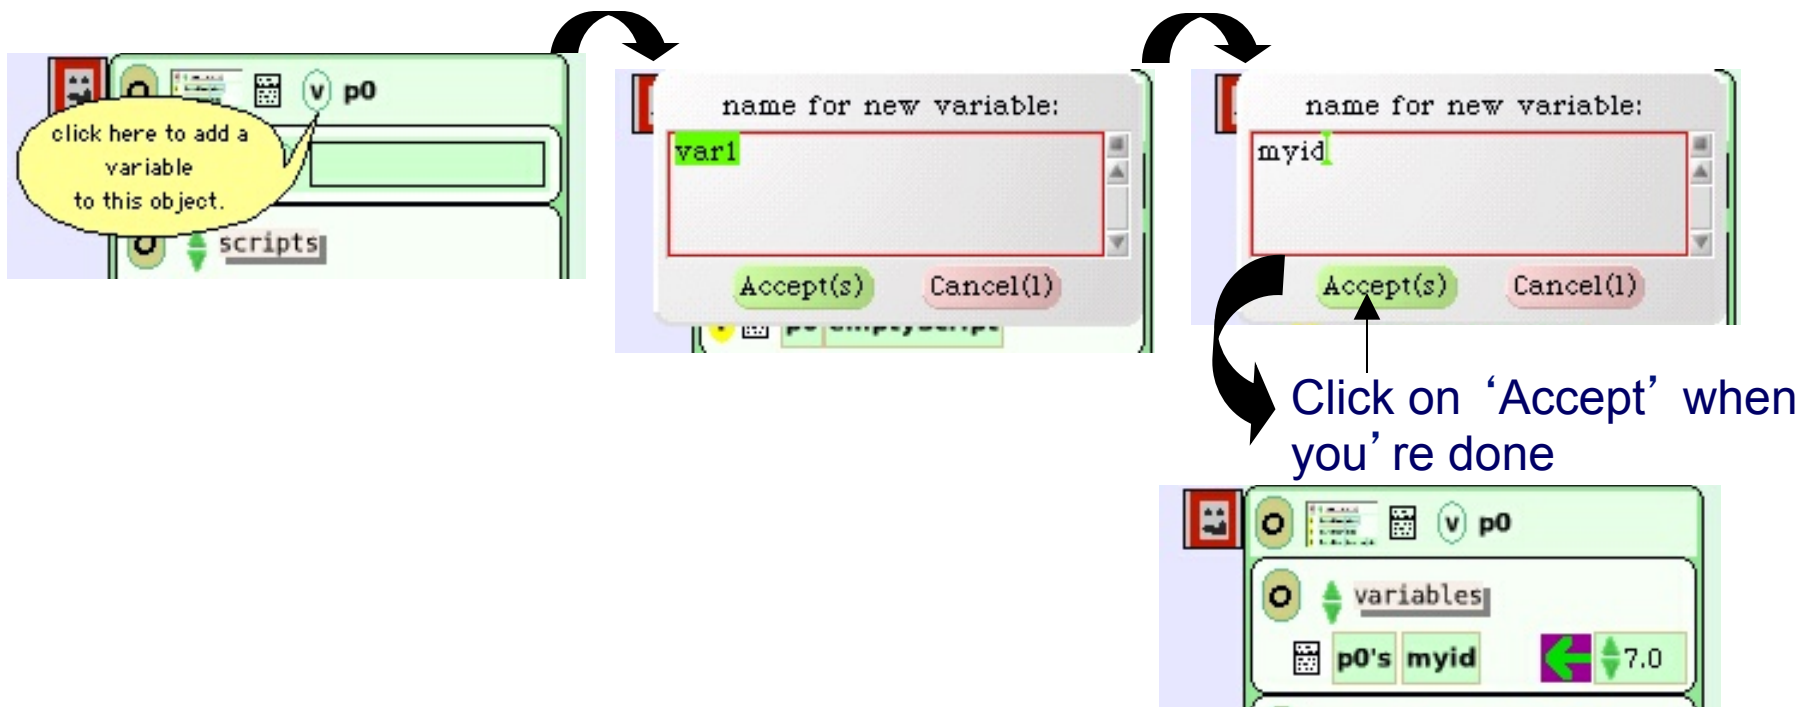
"mousing over" a painted object should reveal a set of "halos". Otherwise, on Windows, Alt-click the object. On Linux, Ctl-click. On Macs, Cmd-click.

4) Open its Viewer
(click the "eyeball" halo)



Object-oriented programming

An object's Viewer lets us 'program' the object. A Viewer contains categories of draggable 'tiles'.

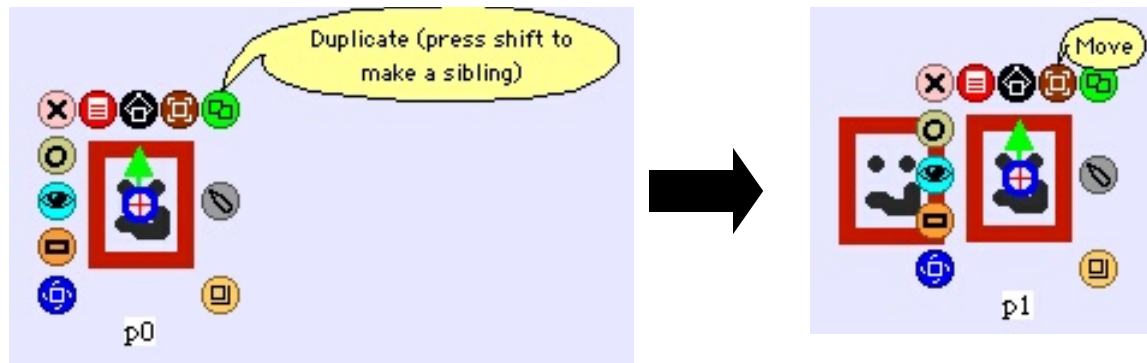


Create a variable 'myid' for the object 'p0'

(myid is a variable used in MPI programs that identifies each computer)



Clone p0 to make p1,p2,p3



Clone p0 by making a sibling (hold 'Shift' while pressing the Duplicate halo). Notice the name of the new sibling is 'p1'.

Make another sibling of p0 (NOT of p1!) → p2

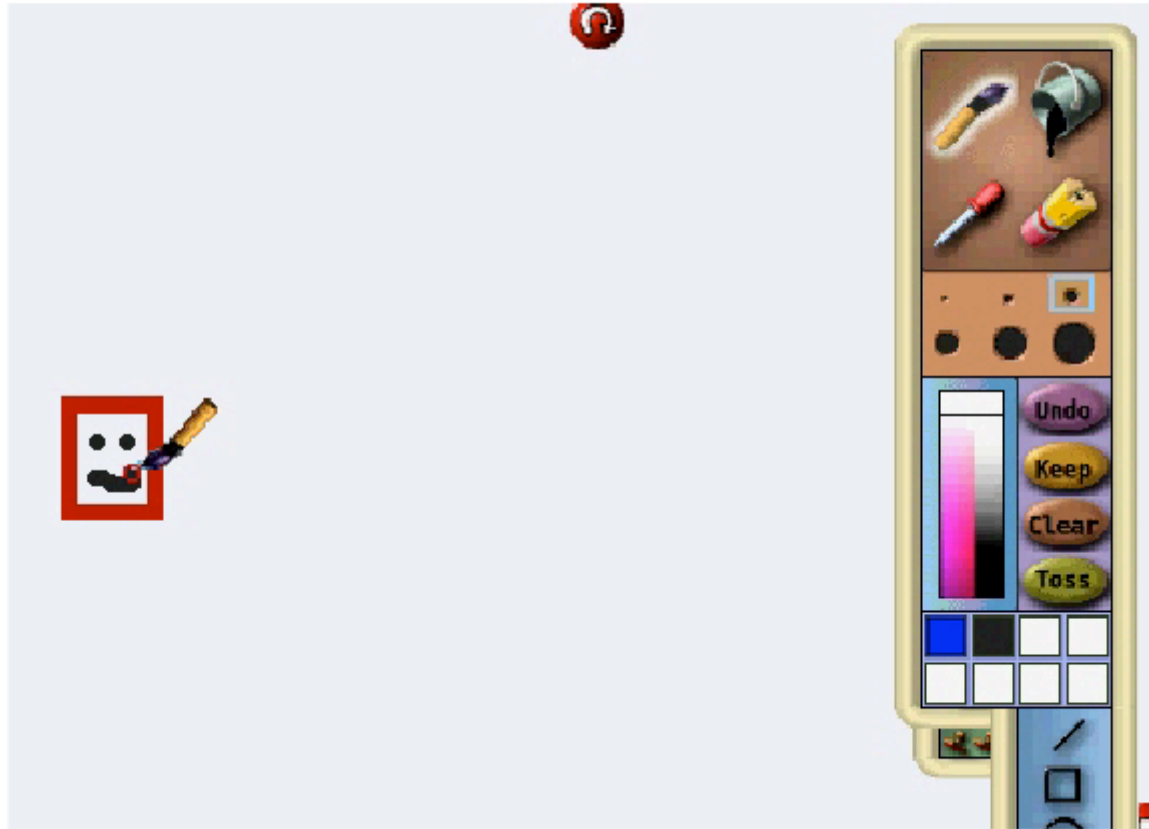
Make a third sibling of p0 → p3



If you need to move p0-p3, you can 'Shift' rubberband them to select the group.



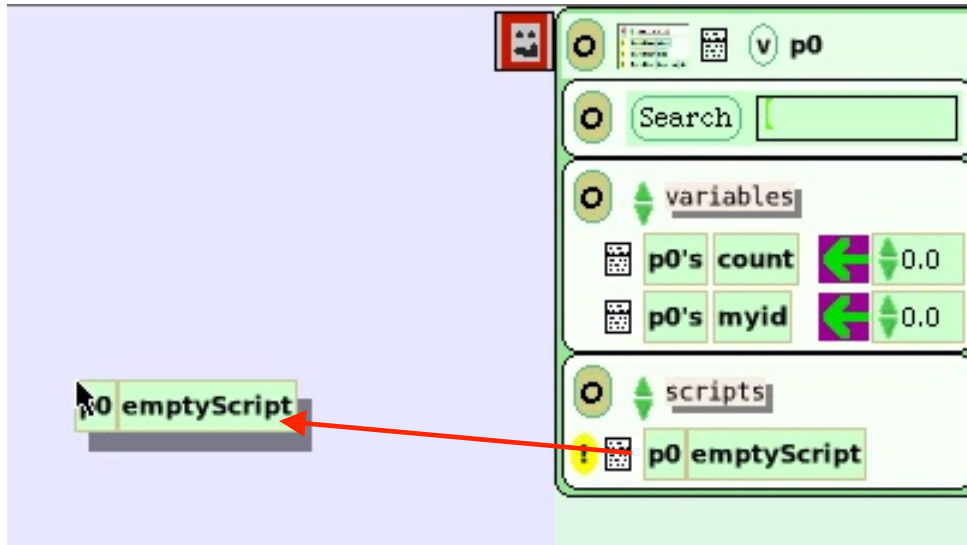
Flash movie: creating p0-p3



When making siblings of p0, be sure to press 'Shift' key while selecting the Duplicate halo.

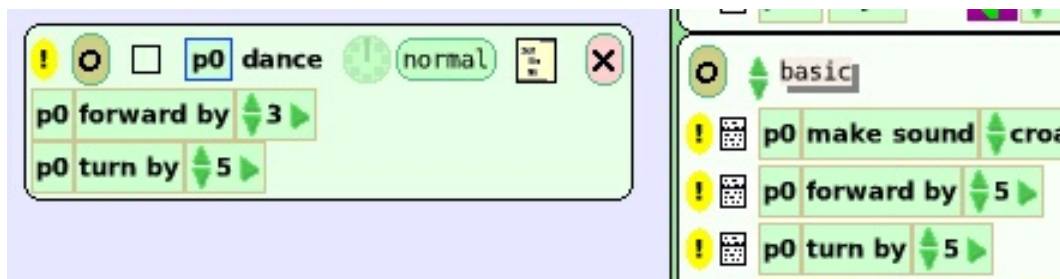


Silly example: 'dancing' computers

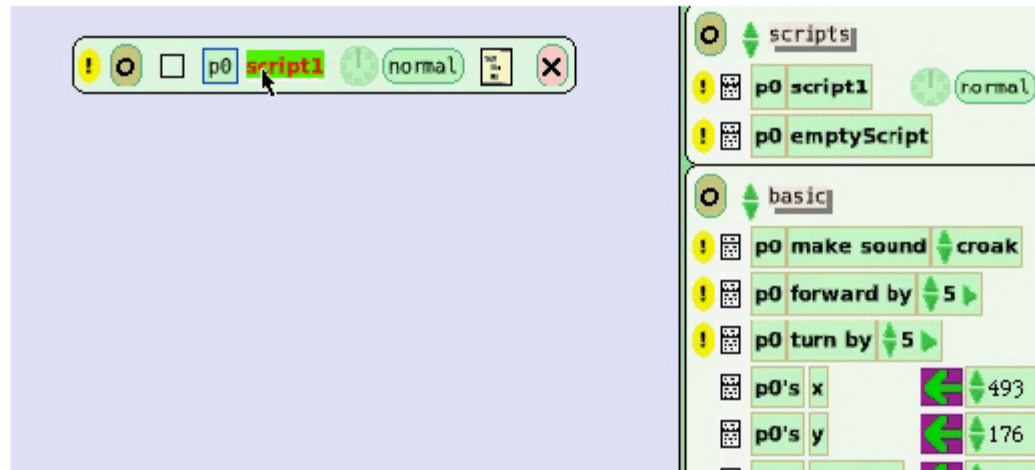


Create a 'dance' script for p0:

- move forward
- turn




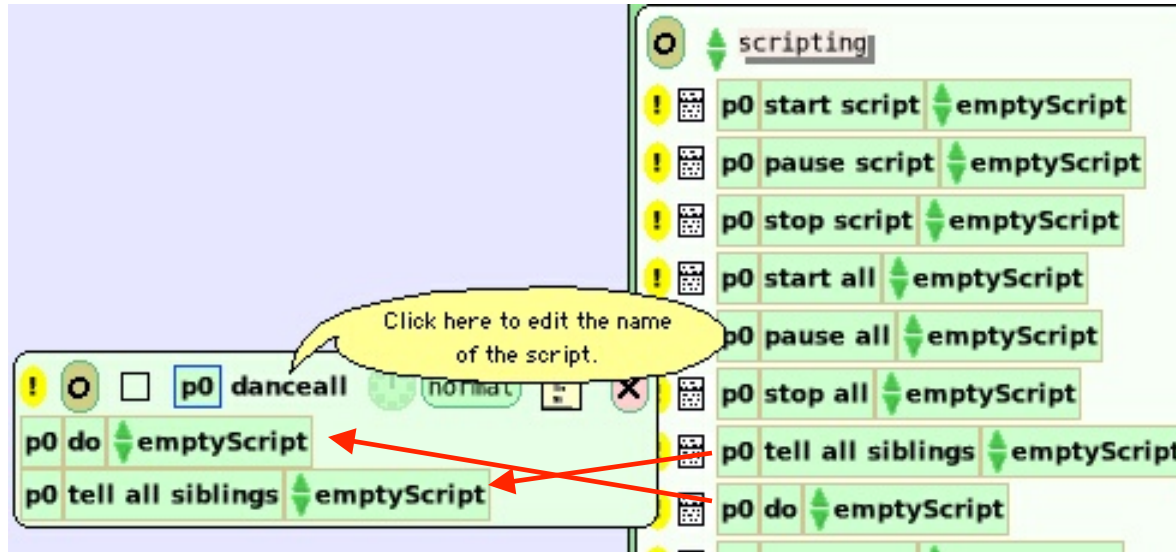
Flash movie: create a script



Create a script (a program) by dragging/dropping tiles (i.e. hold/release mouse button)



Make all computers dance



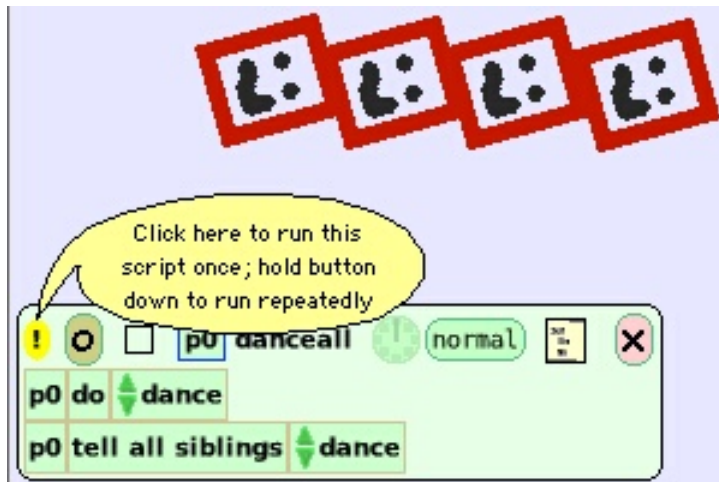
Create another script, name it 'danceall' and have it invoke the 'dance' script for p0 and all its siblings.



Sync'd line dancing

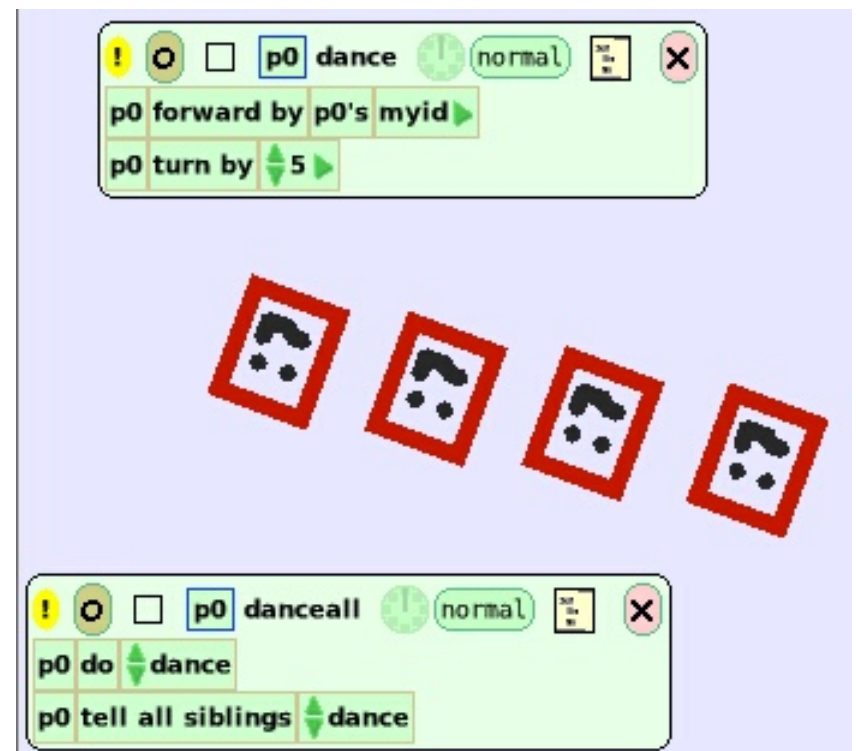
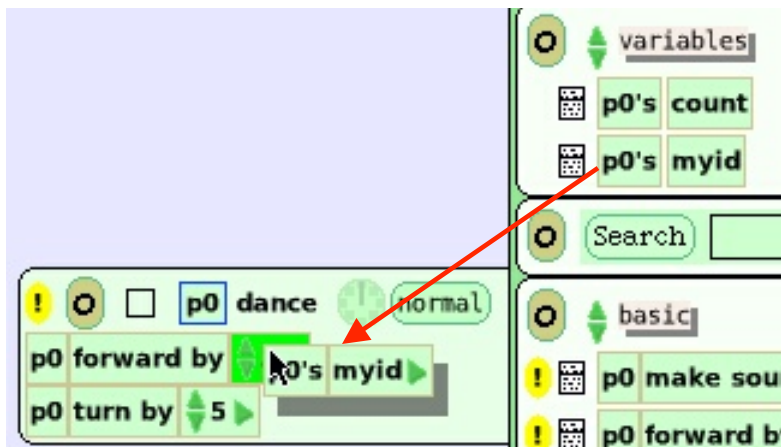


Running the 'danceall' script will cause all computers to dance the same circle dance.




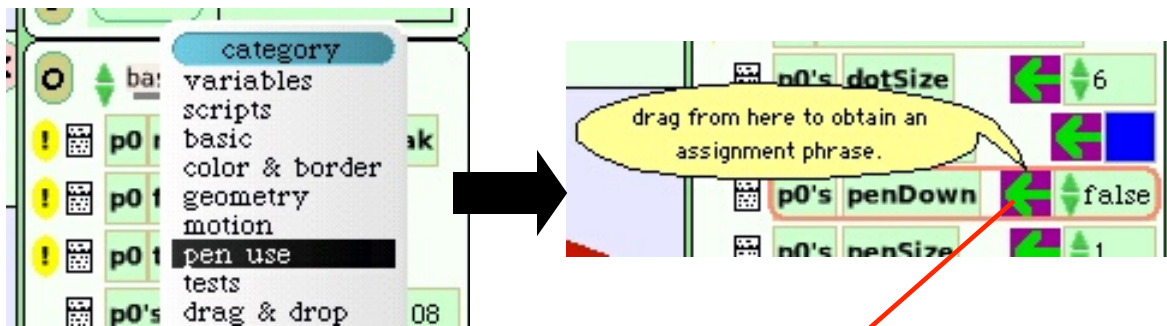
Asynchronous dancing

Replacing the “forward by” fixed value with the variable ‘myid’ will make each computer have a unique dance.



Draw my dance


In the 'pen use' category, drag the 'penDown' tile into our script to draw our paths.



Calculating the value of pi (π)

An introductory MPI program is to compute the value of pi using a Monte-Carlo approach.

Each processor:



- “throws darts” at a square board of unit radius
- if a dart lands inside the unit circle, increment a counter

Since the area of the square board is $=4$ (2×2) and the area inside the unit circle is $\pi r^2 = \pi$

Then we know that $\pi/4 \sim (\text{\#darts in circle}) / (\text{total \#})$ and we can approximate π



Each of 4 processors has a different colored dart – results after a few throws

! O p0 dartall normal [stop] [close]

pt0 do throwdart

pt0 tell all siblings throwdart

pt0's count = 3 Red

pt1's count = 4 Green

pt2's count = 2 Blue

pt3's count = 4 Black

! O pt0 throwdart normal [stop] [close]

pt0's x ← 150 - random 300

pt0's xnorm ← pt0's x / 150

pt0's y ← 150 - random 300

pt0's ynorm ← pt0's y / 150

pt0's z ← pt0's xnorm * pt0's xnorm

pt0's z increase by pt0's ynorm * pt0's ynorm

pt0 stamp

Test pt0's z < 1

Yes pt0's count increase by 1.0

No

The pixel size of our dart board is 300x300; we normalize it to be [-1,1]

Results after several throws

```
! [O] [ ] p0 dartall [! paused] [ ] [X]
pt0 do throwdart
pt0 tell all siblings throwdart
```

```
! [O] [ ] pt0 throwdart [! normal] [ ] [X]
pt0's x ← 150 - random 300
pt0's xnorm ← pt0's x / 150
pt0's y ← 150 - random 300
pt0's ynorm ← pt0's y / 150
pt0's z ← pt0's xnorm * pt0's xnorm
pt0's z increase by pt0's ynorm * pt0's ynorm
pt0 stamp
Test pt0's z < 1
Yes pt0's count increase by 1.0
No
```

pt0's count = 218

Red

pt1's count = 207


Green

pt2's count = 202

Blue

pt3's count = 208

Black



The calculation of pi is left as an exercise...
(get the Squeak project from <http://sda.iu.edu/K-12>)





Thanks!

We gratefully acknowledge the developers of Squeak and the squeakland.org community of users.

