Thanks for the introduction and for the invitation to give the talk.

1 Good evening, my name is Ron Yeh, and I am a PhD student at Stanford University. Today, I will describe my research on integrating regular Paper with Computing Devices to provide mobility and collaboration in a range of tasks.

2 My talk will be broken down into three main parts. First, I will talk about our efforts in augmenting paper notebooks [show notebook].

3 Second, I will describe how we have used paper to augment Wall and Table interactions.

4 Finally, I will talk about my work on a toolkit to enable designers and developers to build novel paper-centric interactions.

5 But before any of that, I will motivate the talk by describing some of the background research I did.

6 Since I am a computer scientist, I work with computers 8 hours a day. I have four monitors, three laptops, two desktop computers, and several handhelds [show oqo]. However, for many disciplines, like design, or field biology, the computer is not the central tool. It is merely a means to facilitate fluid work with other media, from large sheets for brainstorming, to individual physical leaf specimens.

7 When I began this work, I was working on the BioACT grant to build tools that integrate computer science and field biology. I started out by interviewing 23 field biologists at Stanford University, and the Jasper Ridge biological preserve. I observed them in the field, and even took a 20 week class to learn the field methods of researchers at Jasper Ridge. From this work, I quickly realized that paper notebooks are the central tool for field biology work. They use notebooks to organize all their observations, experimental data, and even (occasionally) samples!

8 I followed a class to the Los Tuxtlas Rainforest, where I observed and participated in experiments out in the field.

9 From my observations, we saw that the process of capturing field data is relatively easy. However, accessing the notebook bound data, transforming it to a form that was analyzable, and sharing the notebook data, was much more difficult.

10 This is because there is a clear difference in advantages between our paper tools and our computers. Paper notebooks are robust; the have infinite battery life, and zero boot time. If I throw a notebook to the ground, it will still work (this is not true with my tablet PC). However, as we know… computing devices are great at search, data storage, and sharing of information.

So how do we combine these advantages?

11 In my work, I use the Anoto digital pen. This is a commercial technology that utilizes an inking pen that can capture your handwriting. There is a camera near the ballpoint tip, and a clock that can timestamp your ink strokes. As you write on paper that is printed with a special dot pattern, this camera will track the dots (seen on the left), and record every single stroke that you write. The dots are unique such that the pen can detect different pages, and different locations (x & y) within a page.

12 This is a quick demo of what you can do. On the left, is a photograph of a sheet of notes a biologist has written. Once she plugged the pen into her USB port in her computer, she can upload her notes and they will be rendered as a vector image that you can see on the right.

13 With this technology, we decided we could help solve the access, transformation, and sharing problems that field biologists had, by augmenting their paper notebooks.

14 The project I will now describe is called ButterflyNet. ButterflyNet aids a field scientist's work by assisting with content that the captures using his paper notebook and his digital camera.

15 Here is one interaction that ButterflyNet provides. Suppose that a biologist takes some handwritten notes, at 4:43pm (1643). Then, he takes a photograph a minute later. ButterflyNet will automatically consider these two pieces of content linked, so that the Biologist can find his photos faster when he is back at the lab. ButterflyNet automatically associates content by time.

16 ButterflyNet also enables a biologist to associate photos explicitly, through a manual hotspot gesture. The person takes a photograph, or browses to an existing photograph, and then draws a box in his notebook. Back at the lab, this photograph will be displayed in the digital version of his notes.

17 This first video will describe how hotspot linking works. [Flávio]

18 This is what the ButterflyNet Browser looks like.

19 <Go to 20>

20 <Go to 21>

21 <Go Back to 20>
When a user navigates his notes, he is driving the interaction on the left. On the right, the photographs will appear automatically. The ButterflyNet system will search for photos that most closely match the notes on the left.

22 This is an example of the result of hotspot linking. See how the photograph is pasted into the notes.

23 This next video will describe the browser, and show how the biologist can export handwritten notes to a Multimedia spreadsheet. [Flávio]

24 As you saw briefly in the video, the Pen can talk to our computer in realtime, via Bluetooth wireless. With the Navigate by Pen interaction, the user can tap on a physical page, and ButterflyNet will flip to the right page in the digital notes, and find the related photographs for you.

25 ButterflyNet has been used by biologists at Jasper Ridge, by researchers at NASA, Anthropologists at Intel, and Design Students at Stanford University.

Do we have any questions at this point? If you have any, please ask Flavio and he can type them to me.

26 Next, I will describe how we have used paper to Augment Walls and Tables.

27 This is my vision for the office of the future. You will have a printer embedded into your wall. Every morning, when you walk in, a large sheet of interactive (e)Paper prints out, with all your calendar, todos, maps, sensor readings, etc. If you need to take this sheet with you for any reason, you can tear it off the wall, roll it up, and take it with you to a meeting.

28 We will prototype this future with a Wide Format Printer. The GIGAprints project couples the digital pen, with a large printer, and multiple computing devices, to provide the real-time output that paper cannot provide.

29 Here is an example GIGAprint. This is a wall of Photos that would be useful in a museum, or in a news agency, where you need to find interesting photos very very quickly. With this GIGAprint, the user walks up to the wall and taps on a photo of interest. The photo is then downloaded to his mobile phone and he can walk away with it.

30 Here is the Video of the Photo Wall.

31 We have also built a network monitoring tool that couples the high resolution of the printer with the realtime feedback of an overlaid projector. The print shows the data representing all the computers in our laboratory. 600,000 data points are represented in the graphs. With the print, the realtime projection, and a digital pen, a network administrator can find out when any of our computers have been infected with a worm or denial of service attack.

32 This video demonstrates the network monitoring tool.

33 We have also augmented large paper maps, where tapping on a location on the map, will initiate a query on a handheld device for all the geo-tagged photos that were taken “near” the location where you tapped.

34 We can detect multiple pens in realtime. You can see this in our video demonstration of the Twistr competitive game. Where a player has to find and tap on their photographs before their opponent does.

[Flávio Plays Video 4]

35 In the last section of my talk, I will describe how we can design and build these paper applications.

36 I am currently building a toolkit, that enables designers and developers to build these interactions. The R3 toolkit can be found at hci.stanford.edu/paper.

37 We have two main goals for this toolkit. First, I'd like to lower the threshold for building this class of applications. Building ButterflyNet and GIGAprints was each an extremely tough task, taking many person-hours. If we can reduce the time to development, then that would help immensely. Second, I'd like to provide a high enough ceiling that expert developers can use the toolkit as a platform to do interesting research and build useful tools.

38 So how does one build a Paper Application? First, you need to design and print the paper user interface. Second, you organize the application logic, and connect Sheet objects (which represents pieces of paper), to Pens, and other Devices. Finally, you write custom Graphical User Interfaces and Event Handlers for interesting areas on the page.

39 To design and print a paper interface, we will first need to create the look and feel in some tool that can export a PDF document. I tend to use Adobe Illustrator.

Then, through Java code or our design tool, you can add Active Regions which the pen can interact with. These regions will be layered with the dot pattern.

You export a new PDF from our software, with the pattern overlaid. Finally, you can print this paper interface to a laser printer or Wide-format inkjet printer that is 600 DPI or better.

40 Here, is an O'Reilly poster that I have augmented with active regions. This poster shows a timeline of different programming languages.

41 If we zoom in, we see that I have added a region on top of the Postscript timeline.

42 And zooming further, we see how the dots are arranged, in a grid where each dot is jittered up, down, left, or right.

43 After we have added active regions to our page, we will now need to add event handlers to each region. If you are familiar with developing GUI software, in Swing or Windows Forms, this will be familiar to you.

For example, here, I add an event handler that will detect clicks on a region. Every time a person taps his pen on my region, this event handler will be called, and your code will be run.

44 SO Here are the basic building blocks of any R3 application. There are Sheets that represent pieces of paper. Regions represent content, or any any interesting area on the sheet. Handlers are attached to regions, and they receive notification whenever a person writes in or taps on that region.

There are devices, such as this handheld PC, and Pens.

The application object wraps all of the above objects, and the toolkit object manages the application.

45 Now I will go through a quick example, to show you how I can make a paper application that will advance my powerpoint slides for me.

46 Here is most of the code that we will need to create this application. We will go through it quickly, line by line.

47 First, you create a new Application object.

48 You create one sheet of paper (8.5 x 11 inches in this case).

49 You create a new Active Region, that is like a button, or a menu.

50 You add an event handler to this active region, to advance the slides.

51 Add the region to the sheet, and add the sheet to the application.

52 Allow the application to detect a single pen.

53 And Finally, ask the toolkit to start your application.

54 What does Event Handler look like?

55 This is a handler that will detect the direction of a pen mark. Up, Down, Left, Right, the 8 standard compass directions.

56 Depending on the direction, we can switch and do different things. In this case, I can ask the system to type the Right Arrow Key, to advance my powerpoint. [show demo]

57 I can also do different things for different directions, such as go forward when I mark right, and go backward when I mark left.

58 Now, the process of building Paper Applications with the R3 toolkit was modeled after building graphical user interfaces. For example, in Swing, you create a JFrame, which is the Application Window. You can add a button, and then add listeners that will be invoked every time someone clicks that button.

59 In R3, you create a Sheet, which is the sheet of paper. You can add a Region, and event handlers, that will be notified every time someone interacts with that region! The difference is that we can handle multiple pens, so that the Pen ID will be sent in as extra information every time the event handler is notified.

60 In this next video, we will show you other applications that we have built using this toolkit. It includes a tool for augmenting Video conferencing, a tool for sharing Blog entries, a virtual Audio Tour Guide, and interactions around an augmented map. [Flávio]

61 Very quickly, I'll go over how I made the Audio Guide application.

62 The Map was created by tiling flyby photographs from the Windows Live Mapping application.

63 I imported this into Adobe Illustrator, where I added decorations for active regions that a person could tap on.

64 Then, I go to Java to create the application object.

65 I ask the toolkit to start the empty application, and then I can design the rest of the paper interface by using our interface designer.

66 I simply drag the PDF I exported from illustrator into our designer …

67 The interface designer is a plugin for Adobe Acrobat. You can drag out interactive regions, and save this information to an XML file, which can be loaded in your software. In your software you will then add the active regions by typing actual Java code.

68 In conclusion, today I have described to you our research efforts in integrating paper and digital tools. ButterflyNet demonstrated how we could improve interactions around a paper notebook. GIGAprints showed how we could use large sheets of paper to augment work around tables and walls. Finally, the R3 toolkit enables hci practitioners to design new paper-centric applications. These three projects are open source, and can be found on hci.stanford.edu/paper.

69 Thank you very much. If you have any questions, I can take them now.