



# Mobile Graphics

**Kari Pulli** ■ *Nokia Research Center Palo Alto*

**Scott Klemmer** ■ *Stanford University*

**M**obile phones form a ubiquitous graphics platform. Today, according to the International Telecommunications Union, more than 3 billion people are mobile subscribers—more than double the number of PCs in the world. New devices have significantly richer graphics, processing, interfacing, and networking capabilities than their predecessors. Mobile devices, however, feature some fairly intrinsic properties that distinguish them from PCs. Battery and heat dissipation concerns limit the available power, and the small form factor constrains the physical size of input devices and graphical displays. In addition, the use context for mobile devices is different: they are available at all times, not just at your desk; and the interactions are often briefer and more episodic in their nature. Plus, new sensing technologies—most notably camera and GPS—enable new opportunities for interaction.

This special issue aims to highlight novel research on mobile graphics and interaction, with particular attention to emerging research that addresses the unique attributes of mobile devices.

Some graphics applications make the most sense in the mobile context. For example, we need maps when we're on the move. They help us navigate from one location to another, browse nearby places and objects of interest, or access location-based services. Antti Nurminen presents a mobile 3D city map system that runs on mobile phones. He addresses several practical problems: collecting data and texture maps; compressing the data to fit the devices' small RAM size; sending the map, geometry, and texture data from a server over the air to the device; rendering the views in real time on the device; and trying to maximize the battery time while doing all that. The system uses visibility cells to avoid drawing parts of the city not visible in urban canyons, and it integrates a real-time status of public transportation to the maps by drawing buses and street cars. One design tradeoff is between lower-resolution texture maps that fit the device memory better and render

faster versus the illegibility of important landmarks such as company logos on a building facade.

Harlan Hile and Gaetano Borriello address indoor navigation using camera phones. Many office spaces or institutions such as hospitals have similar-looking corridors and lack distinguishing landmarks that help users orient themselves. Navigation aids could be useful in these environments. While outdoor navigators can position themselves using a GPS signal, this functionality isn't available indoors. Instead, the authors use a Wi-Fi base station to roughly identify a camera phone's location. Their system uses the camera phone image to detect features, such as doorways, and match them to a building floor plan to accurately locate the camera viewpoint with respect to the floor plan. Then it overlays arrows on the camera image, directing users to their destination. In this system, the camera phone serves as a data-capture device (taking the image and sensing the closest Wi-Fi station). All the computation takes place on a server (image segmentation, matching to the floor plan, calculating the camera pose, annotating the original image), and the final annotated image appears on the camera phone display. The entire cycle in the current system takes approximately 10 seconds.

The previous system is an example of augmented reality (AR), where real-world imagery is augmented or annotated with digital information. While Hile and Borriello augmented the camera image with arrows aiding navigation, Wolfgang Broll and his colleagues use real-time AR for gaming that's embedded in real-world outdoor environments. They give examples of two major trends for future AR games: small, easy-to-set-up, user-modifiable AR games and AR games that are tightly interwoven to their physical environment. In the first example, players hunt for magic-potion ingredients using an ultramobile PC with a video see-through interface. The system is easy to modify, learn, and transport to new environments. The second system is tightly

integrated on location at the historical center of Cologne, Germany. As the users move around the game area, the system time-warps them to various historical or future eras, where they meet magical creatures that give them tasks to solve. The authors describe their graphical system, which scales across different devices and processing levels, as well as the authoring tool for creating game content.

This issue's Projects in VR department also features an AR system. Erich Bruns and his colleagues present an overview of PhoneGuide, their adaptive museum guidance system. It uses camera-equipped mobile phones for on-device object recognition. Using Bluetooth beacons, the system knows users' rough location, and the vision system provides a prioritized object list for a user to select from. The system records the selections for later use in further training the system. The museum visitors are finally provided with location and object-aware multimedia content.

Mobile devices are easier to carry than desktop or even laptop computers. However, many applications that a user might want to access on the road aren't available on handheld devices. Some applications offer remote access. However, most such applications have been developed for a large screen and peripherals such as a keyboard and mouse, and so they're often difficult to operate from handheld devices. Fabrizio Lamberti and Andrea Sanna describe a framework that automatically analyzes a desktop application's user interface; creates a summary of it using an extensible description language; and lets the user define a new, more compact graphical user interface that is easier to use when remotely accessing it from a PDA or a mobile phone.

Interactions with mobile devices is still a research challenge. These devices have fewer and smaller keys than desktop keyboards. Most devices provide a five-way joystick button (left-right, up-down, press), but that's not nearly as versatile as a mouse. New touch displays allow direct manipulation using a finger or stylus, but using them obscures the small display. Further, touch usually affords only 2D manipulation, whereas many tasks would benefit from 3D manipulation. Martin Hachet and his colleagues present a prototype for 3D elastic control on mobile devices. The elastic part allows easier, more intuitive rate control. The 3D part lets users combine panning and zooming or change the camera heading and pitch angles and the forward-backward motion without having to switch between different modes—for example, toggling between panning and zooming.

Graphics hardware has increased graphics performance on desktops faster than Moore's law has increased CPU performance, but this has come at the

cost of greatly increased silicon size and power consumption. The key driver for mobile devices isn't pure performance, but sufficient performance while minimizing power consumption. Ben Juurlink and his colleagues present GRAphics AcceLerator (GRAAL), a framework for benchmarking, designing, simulating, and estimating the power consumption and silicon area of embedded 3D graphics accelerators. They concentrate on tile-based rendering architectures, which allow most memory accesses for graphics to take place locally, on-chip, which consumes significantly less power than accesses to the off-chip frame and z-buffers. Although tile-based architectures incur the overhead of having to store the tile geometry before rendering, this article studies a few alternatives on how to manage the scene and graphics state to optimize overall system performance.

**T**his special issue concludes with a survey on the state of the art in mobile graphics research by Tolga Capin and his colleagues. The survey describes the limitations of mobile devices with respect to graphics and interaction. Researchers have developed many solutions to overcome these limitations, and the survey covers mobile graphics hardware development to address the power and computation issue. The authors address various ways of sharing storage and processing capacity as well as transmitting data between mobile devices and remote servers. They review developments in 3D displays and ways to visualize objects in a small display as well as sensors such as accelerometers and cameras that facilitate new interaction techniques. The survey concludes with a roadmap for future mobile graphics research.

We hope you enjoy both the survey and collection of specific projects we've assembled here. ■■

**Kari Pulli** is a research fellow at Nokia Research Center. He has been an active contributor to several mobile graphics standards and recently wrote a book about mobile 3D graphics. Pulli received a PhD in computer science from the University of Washington and an MBA from the University of Oulu. Contact him at [kari.pulli@nokia.com](mailto:kari.pulli@nokia.com).

**Scott Klemmer** is an assistant professor of computer science at Stanford University, where he codirects the Human-Computer Interaction Group. He received his BA from Brown University in art-semiotics and computer science; his MS and PhD from UC Berkeley in computer science. He is a recipient of the Microsoft Research New Faculty Fellowship and Sloan Fellowship. Contact him at [srk@cs.stanford.edu](mailto:srk@cs.stanford.edu).