Device and Physical Data Independence for Multimedia Presentations

RICHARD STAEHLI, JONATHAN WALPOLE, and DAVID MAIER

Department of Computer Science and Engineering, Oregon Graduate Institute of Science & Technology, Portland, Oregon ⟨{*staehli, walpole, maier*}*@cse.ogi.edu*⟩

Multimedia computing promises access to any type of visual or aural medium through digital networks. But can a given multimedia document be effectively accessed everywhere? The presentation of data must adapt to both the available communications bandwidth and the output device resolution. Current multimedia systems assume that applications require the highest possible quality and handle resource overloads through ad hoc methods, such as video frame dropping. To support a variety of applications with lower quality requirements, we need both new standards for scalable data encoding and new techniques for communicating application quality requirements. This paper describes a new approach for specifying *quality of service* (QoS) requirements based on functionality rather than on data encoding and device capabilities.

The potential of distributed multimedia computing can be achieved by offering *device-independent* and *physical-data-independent* service interfaces. Logical data independence is also desirable, but we omit discussion of it here in the interest of brevity. Device and physical-data independence are well known principles of database system design. In multimedia systems they have the following meaning:

—The same content can be presented on devices that have different resolution and bandwidth characteristics.

—The location and encoding of stored data should be transparent to the user.

Device independence is already supported by some content authoring standards. For example, the emerging ISO MHEG standard uses *virtual coordinates* for content layout [Meyer-Boudnik and Effelsberg 1995]. However, MHEG and most other authoring standards identify content with a particular encoding of data. The presentation quality for such content typically depends on the encoding, the presentation engine, and the available resources. If requests for multimedia services are to have the same meaning on any platform, they should specify QoS requirements that are device- and physical-data-independent.

QoS specification. We propose a three-step methodology for QoS specifications: defining an ideal presentation, choosing an error interpretation, and constructing a user mode.

An *ideal presentation* is the set of expected output values for every point in the presentation space and time. The ideal output values may vary continuously over the coordinate space of real numbers, unlike the actual output, which has finite resolution and discrete values. As a consequence, the specification of an ideal presentation is device-independent, like a PostScript document. Figure 1(a) illustrates the specification of an ideal presentation through a *content* descriptor that may be reused in many presentations and a *view* descriptor that specifies a particular mapping of content onto device and realtime coordinates. Figure

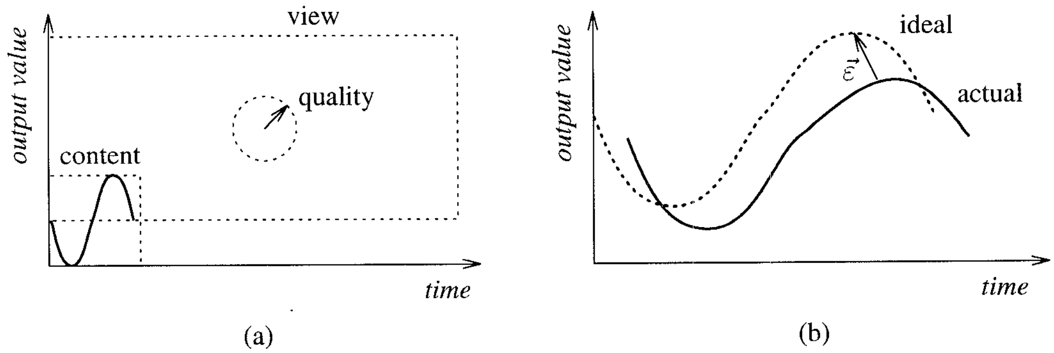

Figure 1. Content, view, and quality specify QoS.

1(b) shows the mapped content with a dotted line.

An *actual presentation* will deviate from the ideal because of device limitations and choice of presentation algorithms and scheduling policy. Device limitations such as screen color depth may require dithering or some other approximation of the ideal color values. Video resolution is limited by the pixel dimensions of an output window, and sample rates for both video and audio are limited by device bandwidth. The choice of compression, decompression, and rendering algorithms can introduce errors in the output values. The choice of scheduling policy affects the timing of those output values. However, just as the specification of an idea presentation is device-independent, the specification of allowable error should be independent of the mechanisms used for a presentation.

An *error interpretation* maps each point in an actual presentation to a point in the ideal presentation. Figure 1(b) shows an error interpretation $\vec{\varepsilon}$ for a single point in an audio presentation. The vector $(\varepsilon_t, \varepsilon_v)$ says that the value v at time t should have occurred at time $t + \varepsilon_t$ and should have had the value $v + \varepsilon_v$. An interpretation of error in a video presentation must also account for displacement errors in both x and y dimensions.

The definition allows many different error interpretations for a given pair of ideal and actual presentations. It is tempting to define a "correct" interpretation of error based on the intended correspondence of output events with content values for a particular implementation. But we want to constrain presentation outputs, not the implementation.

Finally, presentation quality requirements can be defined in terms of a *user model.* A user model estimates subjective presentation quality from an error interpretation. We have described a user model based on an error vector of *shift*, *rate*, *jitter*, and *resolution* components for each coordinate dimension and a *synchronization* error component for the timing error between outputs [Staehli et al. 1995]. These error components constitute a detailed error interpretation that can more closely model human perception. The normalized magnitude of the error vector is computed by weighting each error component according to user sensitivity. The user model "accepts" a presentation if the normalized error magnitude is within a specified limit everywhere for some error interpretation. This approach is conservative. Alternative user models might bound the average error or place other constraints on the distribution of error over an entire presentation.

QoS-Driven Presentations. Multimedia systems can provide better service if the QoS requirements of each client are known. A QoS specification can serve as a throttle to reduce resource use: requesting, for example, 24 frames/second

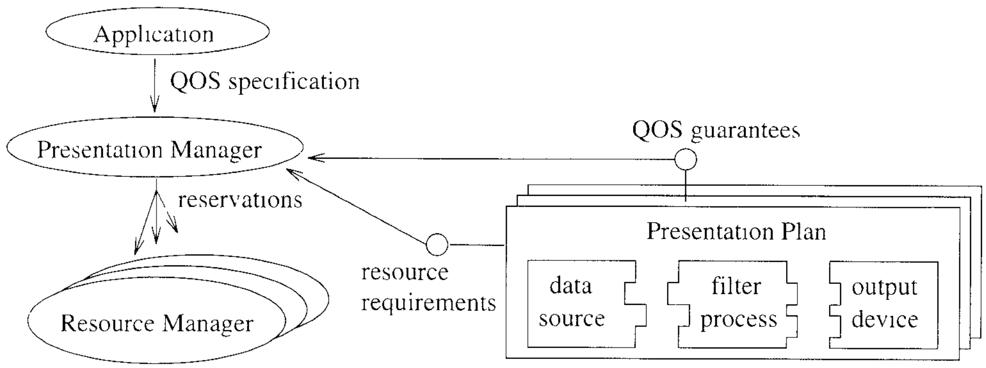

Figure 2. Presentation planning.

video when a data source could supply 60. Resource regulation is essential in a shared environment. A QoS specification can also indicate which use of resources provides the best quality for a particular presentation. For example, in a bandwidth-constrained environment, an action video might be best presented at 320 × 240 pixels and 15 frames/second, whereas a video of a chalkboard lecture uses the same bandwidth more effectively with 640 × 480 pixels and 4 frames/second.

Some systems guarantee performance, whereas others may only provide best-effort service. QoS specifications are needed to drive resource-management decisions in both cases. Best-effort resource management still involves making decisions about how to trade one kind of resource consumption for another. This set of decisions is a planning problem that can be guided by the QoS specification. Making guarantees (which can be hard or statistical) requires an end-to-end resource reservation approach and an admission test [Jones et al. 1995; Campbell et al. 1993; Nahrstedt and Smith 1995]. Figuring out which of many different resource allocation plans is best is an optimization decision that can be guided by the QoS specification.

Figure 2 illustrates a high-level architecture for an admission test. A presentation manager receives the QoS requirements for a presentation from an application. A presentation plan is *feasible* if it can guarantee the QoS requirements and if the presentation manager can reserve resources for the plan. The admission test can choose to execute the feasible presentation plan with the fewest resource requirements.

CONCLUSION

Multimedia systems are only beginning to realize the flexibility inherent in digital computing. More work is needed to understand QoS requirements for multimedia presentations and to exploit those requirements for optimal resource management. Device- and physical-data-independent QoS specifications allow applications to say *what* multimedia services are required without restricting *how* they are implemented.

REFERENCES

CAMPBELL, A., COULSON, G., GARCIA, F., HUTCHISON, D., AND LEOPOLD, H. 1993. Integrated quality of service for multimedia communications. In *Proceedings IEEE INFOCOMM '93* (San Francisco, April), IEEE.

JONES, M. B., LEACH, P. J., DRAVES, R. P., AND BARRERA, J. S., III. Support for user-centric modular real-time resource management in the Rialto operating system. In *Proceedings of NOSS-DAV 95* (Durham, NH, April), IEEE Communications Society, 55–66.

MEYER-BOUDNIK, T. AND EFFELSBERG, W. 1995. MHEG explained *IEEE Multimedia 2*, 1 (Spring), 26–38.

NAHRSTEDT, K. AND SMITH, J. M. 1995. The QoS broker. *IEEE Multimedia 2*, 1 (Spring), 53–67.

STAEHLI, R., WALPOLE, J., AND MAIER, D. 1995. Quality of service specifications for multimedia presentations. *Multimedia Syst. 3*, 5/6 (Nov.), 251, 263.