Computer Graphics International 2011

The Workflow Scale

Why 5x Faster Might Not Be Enough

Eric Enderton · Daniel Wexler

Abstract This essay discusses qualitative versus quantitative accelerations of user tasks, in the context of computer animation production. A workflow regime is defined as a range of system response times in which the artist's relationship to the task is qualitatively similar. Radical new technology is much more likely to succeed when it brings an artist's workflow into a new regime, providing a discontinuous improvement in efficiency and final image quality. More modest technology revisions can supply smaller speed-ups, but these are merely consumed by Blinn's Law, i.e., the tendency for system response times to remain constant as technology improves, due to increased input complexity. We propose a list of workflow regimes and their ranges of response time.

Keywords interactivity · Blinn's Law · technology adoption · computer graphics · computer animation production

1 Three Examples

Research and development creates ever more efficient methods for artists to edit computer animated scenes. But there are two types of speed-ups: incremental and breakthrough. An incremental speed-up makes the user proportionally more efficient, allowing either faster edit cycles or more complex inputs. But a breakthrough – such as from overnight runs to hourly runs, or from batch to interactive – changes the workflow and the user's qualitative experience.

While these concepts could apply to a wide range of industries and tasks, let us consider a hypothetical example at a computer animation studio. Suppose you are an artist who is given the opportunity to use a new lighting tool that is five times faster than the old one. The new system is modern and intriguing, but it has the usual rough edges, missing features, integration issues, and learning curve. At what point will you tell your supervisor that you have adopted the new system, and no longer require the old one? We claim the answer depends on the absolute system response time (in this case, the render time).

Ten hours (old) versus two hours (new). Before, you worked all day on new parameters, estimating their effect with small tests. You submitted the job once, before going home, knowing that the results would be examined by your team the next morning. With the new system, you see the results while you still remember what you changed. You may feel more leeway to experiment, and perhaps find a new direction to show your supervisor after lunch. Unless the new system produces significantly inferior results, you will quickly abandon the old system.

Ten minutes versus two minutes. While it is bracing to get so many more turnarounds per hour, you still have to find something else to do during the two minute renders. By the time the output appears, you must actively compare it to the previous image, and double-check the parameter value you just changed. Your workflow has not changed. And if you get distracted by your interim task (such as email), your turnarounds per hour may go back down. If the new system is missing features you found effective, or causes other delays, then you might very reasonably stick to the old system. You are more productive there.

E. Enderton
NVIDIA Corp.
E-mail: eenderton@nvidia.com

Ten seconds versus two seconds. At ten seconds, you were careful not to mistype parameter values. At two seconds, you type and hit the render button without worrying about it. Before, you thought about the shot as a whole while waiting for renders. Now, you are still thinking about the same parameter you just changed when the new image appears. Before, you never turned on "auto render." Now, you find you sometimes hold the mouse down during the whole two seconds, and you're experimenting with auto render. This has changed your relationship to the image. If you like the new relationship, you are inclined to work through any problems with the new system – so that you never have to wait for ten seconds again.

In the "ten hour" and "ten second" scenarios, the new technology is likely to be adopted, because it has changed the artist's workflow. The speed-up it provides crosses a boundary; the artist's subjective experience and relationship to the task at hand has shifted significantly. The technology has succeeded. By contrast, the "ten minute" scenario supposes the same speed-up factor (5x), yet it remains within the same general workflow, and adoption is an uphill battle. The new system has failed.

2 Workflow Regimes

Not everyone's experience will exactly match these fictional anecdotes. Yet, in this era of rapid technological change, it seems everyone has encountered a qualitative change of their experience when taking up a new digital tool – whether that new tool is GUI color sliders, displacement sculpting with live ambient occlusion, or simple SMS text messaging. Accelerating a task can change the user's workflow in a way that reaches beyond mere time scaling of the activity. If this is mediated by human psychological abilities, such as the decay time of short-term textual or visual memory, then it should not surprise us if there is a universal time scale for workflow change.

The time axis of this scale would be the overall system response time: the time to render and display a new image, for example, or the time to recalculate and display a simulation, once the user has requested it. Let us define a *workflow regime* to be a range of system response times over which workflow does not change. Within a regime, the user's subjective experience is similar; their mental relationship to the task at hand does not qualitatively change. Here is our initial attempt at identifying and quantifying the workflow regimes for digital content creation tasks. Please refer to Figure 1, which graphs these regimes by response time. (Regime boundaries are of course inexact.)

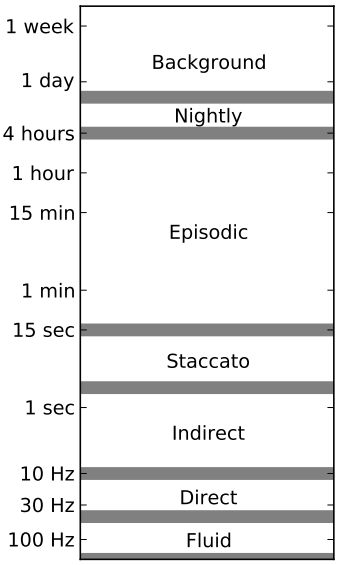

Fig. 1 The workflow scale: response time to user input (log scale) versus workflow regime. Accelerating a task has the greatest impact when it moves the task to a new workflow regime.

Direct. At frame rates of 20 Hz or better, we have an impression of direct physical connection to the image: the image changes as we drag the mouse.

Indirect. As the frame rate drops towards 1 Hz, the visceral connection breaks down into something more conscious. We feel the system is heavy or sluggish, and we may overshoot. But these are still physical metaphors; we feel physically connected, albeit indirectly. This connection is critical for focused and rapid task achievement. (Ironically it is often billed as "direct manipulation.")

Staccato. At response rates of several seconds, we experience a series of separate responses rather than any smooth scale. We are no longer driving the car, so much as sending it messages. While our attention may wander away from the particular parameter we are adjusting, it does not have time to wander from the task entirely.

Episodic. During a response time of a couple minutes, we have time to disengage from both the parameter and the larger task. We may tend to other windows on our desktop. As minutes stretch to hours, we get coffee, attend meetings, and so on. Using the system becomes more technical. But by returning to the task periodically over the course of the day, we can still explore options.

Nightly. Here, only one run per day can be completed. In the traditional rhythm of film production, artists spend the work day setting up jobs that run on the server farm overnight.[1] The team views the resulting footage in "dailies" each morning; experimentation is much less valuable than forward progress.

Background. As response time (and resource consumption) stretches to multiple days, we start referring to both the running job and the task we are trying to complete as a "background task." Eventually the decision to launch a job at all becomes too big for one person.

Naturally, the scale can be extended. The smaller time scale towards 60-120 Hz invokes a sense of fluidity or realism, and is important for video games and immersive simulations. Larger time scales occur in some research and industrial settings.

3 Blinn's Law

For Industrial Light & Magic's visual effects work on the 1993 film *Jurassic Park*, typical overnight render times were 4 hours per frame. When the same company worked on the 1997 movie sequel, using both faster computers and more advanced software, frame times were typically an identical 4 hours. The imagery in the sequel was of course more complex – more dinosaurs, richer lighting, etc. While the precise value of the "threshold of pain" for render time varies by institution, this effect has been widely observed. Jim Blinn famously remarked, "All frames take 45 minutes," and the effect is known as Blinn's Law.

This is why a series of incremental speed-ups do not accumulate to become a workflow breakthrough. Each small speed-up, such as those provided by Moore's Law, tends to be used for greater input complexity rather than greater interactivity. Blinn's Law keeps the system within one workflow regime. Eventually however, input complexity reaches a point of diminishing returns; when the task output is "good enough," the task can actually get faster. (Animation and rendering for films and video games have apparently not reached this point.)

Raising input complexity is, of course, a choice; Blinn's Law is not enforced. Particularly in cases where results can be easily layered or combined, an artist who makes faster iterations on simpler input may converge more quickly on the desired final output, even a complex output.[2] In principle, a studio may choose to reduce input complexity in order to enter a better workflow regime. In practice, this leap is rarely made. The structures and rhythms of the institution are tuned to the old workflow. A new workflow often begins as experimentation by a few rebels, and ends up as a new department.

4 Conclusions

A workflow regime change results in a discontinuous improvement in final image quality by facilitating a deeper interactive relationship between the artist and their art. Technologists should be aware that improvements that fail to cross a regime boundary run a risk of failing to be adopted.

We hope that this preliminary survey of workflow regimes and their impact can inspire insights into the costs and benefits of technology development. Although here we consider only film production, we conjecture that the same principles apply to digital content creation in general, particularly computer game development, and even to other computer-mediated tasks, such as data mining and scientific visualization, Finally, we call for more formal psychological research to examine the relationship between users' qualitative experience and productivity.

[1] Note that the nightly job is to run all the frames in the shot, on many computers. Individual frames may still run hourly, even on a single computer.

[2] Similarly, a simpler tool with cruder output but a tighter feedback loop can be more effective than a complex tool with longer feedback.