

## Possible Opportunities at the Intersection with Empathic Systems

Peter Dinda\*                      Gokhan Memik  
Department of Electrical Engineering and Computer Science  
Northwestern University

The Empathic Systems Project<sup>1</sup> is a decade-long effort at Northwestern University that has focused on a new approach to core problems in computer systems and computer architecture. Hardware architectures and systems software are the foundation of watches, phones, laptop and desktop computers, servers, clouds, and the networks that stitch all these elements together. While the software and hardware are largely invisible to the user, the decisions that are continuously made by them (often every millisecond) have a direct impact on the performance and quality experienced by the user, as well as the energy efficiency and environmental impact of the user's actions. Almost all extant hardware and software assumes a *canonical* user in making these decisions. The Empathic Systems Project's approach instead is to make these decisions in a way that incorporates measurement of the satisfaction of the *individual* user.

Our approach has been quite successful in a variety of domains such as scheduling, processor, display and network power management, network quality of service management, remote display, and volunteer computing. We have generally found that for any given fixed operating point in these domains there is considerable variation across users in the measured satisfaction for the point. Because of this variation, tailoring the operating point to the individual user, both raising it (e.g. increasing performance) for more demanding users, and lowering it (e.g. decreasing performance) for more tolerant users, leads *both* to more satisfied individual users and lower resource consumption across users.

The research we conduct combines both core systems/architecture work (we build and test systems software and hardware) *and* human/computer interaction work (we conduct studies involving human subjects). Our work thus far has been entirely outside of the health care area, but we are interested in becoming involved. Our current efforts focus on phones, clouds, and biometrics. In the following we describe two possible opportunities we have gleaned from the experience described above, and from (limited) interaction with those in the health care area, particularly mobile health.

One opportunity we see is a new way of affecting behavior through smart phone control below the level of the app or page the user is interacting with. College students, and everyone else, are spending more and more time glued to their smart phones and other mobile devices. This is time that could be potentially be used to motivate physical activity or other healthy pursuits. Often the user of the phone desires to act more healthily on principle, at a high cognitive level, but then does not follow through. Of course, reminders are one way to force long-term goals into the short-term interaction with the device.

As an alternative, consider tying the performance of the user's phone to the user's progress toward the user's own high-level health goals. As the user progresses, he earns credits which are spent to make the phone and cloud perform better. Not only does this provide continuous feedback to the user, as opposed to the event feedback of notifications, it also avoids interrupting the user's workflow on the phone. It is also a negative feedback mechanism---as the user ignores his high-level health goals, the phone gradually becomes less usable, and thus less of a distraction from those goals.

In our work, we often have the need to run studies where we modulate the performance of different aspects of the phone or other computer in order to determine their effects on user satisfaction. Currently, we can modulate CPU performance (for native apps), JavaScript

---

<sup>1</sup> <http://empathicsystems.org>

performance (for web apps), and network connection performance (for native and web app interaction with the cloud). These technologies could be readily used to prototype and evaluate the concept described above.

The second opportunity we see relates to biometrics, in which there is clear mutual interest. In our work, we desire direct, continuous measurement of user satisfaction with the current performance of the system. One approach we have taken to this is to use biometrics; we leverage the raw measurements not to determine any health measure, but rather to compute a measure of user satisfaction with performance that can then be fed forward. An empathic system could then be envisioned as a control system that is modulating performance to maintain a stable satisfaction level.

We have evaluated this model with human subjects studies “in the lab” with promising results. We have frequently built our own hardware to support this work, and we are currently developing our own highly configurable variant of a health monitoring band / health watch in order to support a study to evaluate the ideas “in the wild”.

Considerably flexibility in sensor selection is necessary for our work and off-the-shelf health bands and platforms are insufficient, hence our hardware efforts. From our considerable distance (we could be very wrong), we have noticed that the mobile health (and perhaps also student health) community seems to also be somewhat limited by off-the-shelf products. This can sometimes take on surprising dimensions. For example, sensing that the user is sitting using sophisticated filtering of accelerometer data from a health band or phone seems like a complex and fraught software solution compared to a simple hardware solution of having a strain or pressure sensor on the user. As a side project, we advised an undergraduate group who did just that using our nascent platform.

It seems to us that an open, extensible, hardware platform for biometric sensing would be of benefit to the college health community. Such a platform would enable the selection of sensors at the level of electronic components instead of at the level of retail products. Additionally, a campus deployment could provide data both for health and for empathic systems. Finally, the right platform could involve the “maker” community, of which many college students are a part. Indeed, innovation in sensors, hardware, and software, as well as apps, could even become part of the output of a deployment of a platform at a college.

Peter Dinda is a professor in the Department of Electrical Engineering and Computer Science at Northwestern University. He holds a B.S. in electrical and computer engineering from the University of Wisconsin and a Ph.D. in computer science from Carnegie Mellon University. He works in experimental computer systems, particularly parallel and distributed systems, and has authored over 120 scientific papers, authored or is a major contributor to several large publicly available codebases, and holds five patents. His research currently involves virtualization and operating systems for distributed and parallel computing, programming languages for parallel computing, and empathic systems for bridging individual user satisfaction and systems-level decision-making. He is a Fellow of the IEEE. You can find out more about him at [pdinda.org](http://pdinda.org).

Gokhan Memik is a professor in the Department of Electrical Engineering and Computer Science Department at Northwestern University. He received the B.S. degree in Computer Engineering from Bogazici University and a PhD in Electrical Engineering from the University of California at Los Angeles. His research area is computer architecture. He is the author or co-author of 2 book chapters and over 100 refereed journal/conference publications, including several best paper nominations and awards. He also co-authored NetBench and MineBench, two widely used benchmarking suites for networking and data mining applications, respectively. You can find out more about him at [www.eecs.northwestern.edu/~memik](http://www.eecs.northwestern.edu/~memik).