COMPUTING edge

> Smart Health
> Wearables
> Software
> Quantum Computing

JULY 2019
www.computer.org

IEEE

IEEE COMPUTER SOCIETY

TCHPC IEEE Computer Society Technical Consortium on High Performance Computing

Call for Nominations: IEEE CS TCHPC Early Career Researchers Award for Excellence in High Performance Computing – 2019

The IEEE Computer Society TCHPC Early Career Researchers Award for Excellence in High Performance Computing recognizes up to 3 individuals who have made outstanding, influential, and potentially long-lasting contributions in the field of high-performance computing within 5 years of receiving their PhD degree as of January 01 of the year of the award. It is sponsored by the IEEE Computer Society **Technical Consortium on High Performance Computing (TCHPC)** and its member Technical Committees:

- Technical Committee on Parallel Process (TCPP)
- Technical Committee on Computer Communications (TCCC)
- Technical Committee on Distributed Processing (TCDP)
- Technical Committee on Cloud Computing (TCCLD)
- Task Force on Rebooting Computing (TFRC)
- Technical Committee on Computational Life Sciences (TCCLS)

Nominations: A candidate must be nominated by member(s) of the community. An individual may nominate at most one candidate for this award. The nomination application must be submitted via email to **tchpc-awards@computer.org** as a single PDF file and should contain the following details:

1. Name/email of person making the nomination (self-nominations are not eligible).
2. Name/email of candidate for whom the award is recommended.
3. A statement by the nominator (maximum of 500 words) as to why the nominee is highly deserving of the award. Note that since the award is for outstanding contributions, the statement and supporting letters should address what the contributions are and why they are both outstanding and significant. The nomination should also list the names and email of up to 3 persons who have provided letters supporting the nomination.
4. CV of the nominee.
5. Up to three letters of support from persons other than the nominator – these should be collected by the nominator and included in the nomination.

Important Dates:

- Nomination Deadline: August 15, 2019
- Results Notification: September 15, 2019

Award Selection Committee: The award selection committee consists of:

- **Lin Gan,** Tsinghua University and National Supercomputing Center in Wuxi, China
- **Vladimir Getov,** University of Westminster, UK (Chair)
- **Minyi Guo,** Shanghai Jiao Tong University, China
- **Ananth Kalyanaraman,** Washington State University, USA.
- **Christine Morin,** Inria, France
- **Irene Qualters,** Los Alamos National Laboratory, USA
- **Satoshi Sekiguchi,** National Institute of Industrial Science and Technology, Japan
- **Min Si,** Argonne National Laboratory, USA
- **Yogesh Simmhan,** Indian Institute of Science, Bangalore, India
- **Edgar Solomonik,** University of Illinois at Urbana-Champaign, USA

Note that members of the selection committee cannot be nominators or provide support letters.

Award & Presentation Note: Awardees will be presented a plaque and will be recognized by IEEE Computer Society and TCHPC websites, newsletters, and archives. The awards will be presented at the SC19 conference that will be held in Denver, CO, USA, 18–21 Nov. 2019. Details of the conference can be found at http://sc19.supercomputing.org/.

For more information, please email tchpc-awards@computer.org.

COMPUTING edge

IEEE COMPUTER SOCIETY computer.org • +1 714 821 8380

Printed with inks containing soy and/or vegetable oils

IEEE

IEEE COMPUTER SOCIETY

Circulation: ComputingEdge (ISSN 2469-7087) is published monthly by the IEEE Computer Society. IEEE Headquarters, Three Park Avenue, 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; voice +1 714 821 8380; fax +1 714 821 4010; IEEE Computer Society Headquarters, 2001 L Street NW, Suite 700, Washington, DC 20036.

Postmaster: Send address changes to ComputingEdge-IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Printed in USA.

JULY 2019 • VOLUME 5, NUMBER 7

COMPUTING edge

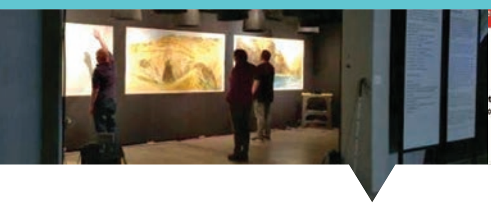

44
Quantum Computer Scale-up

Magazine Roundup

The IEEE Computer Society's lineup of 12 peer-reviewed technical magazines covers cutting-edge topics ranging from software design and computer graphics to Internet computing and security, from scientific applications and machine intelligence to visualization and microchip design. Here are highlights from recent issues.

Computer

Rack-Scale Capabilities: Fine-Grained Protection for Large-Scale Memories

Rack-scale systems with large, shared, disaggregated, and persistent memory need solid protection and authorization techniques. The authors of this article from the February 2019 issue of *Computer* propose a solution that uses a memory-side capability enforcement processor, which gates memory accesses through extended capabilities, enables fine-grained access control beyond a single address space, and minimally disrupts the programming model.

Computing in Science & Engineering

Developing a Computational Chemistry Framework for the Exascale Era

Within computational chemistry, the NWChem package has arguably been the de facto standard for running high-accuracy numerical simulations on the most powerful supercomputers. To better address the

Published by the IEEE Computer Society

challenges presented by emerging exascale architectures, NWChem is being rewritten. Design of the resulting package, NWChemEx, is driven by exascale computing, as well as the team's involvement with the Molecular Sciences Software Institute (MolSSI). MolSSI is an NSF initiative focused on establishing coding and data standards for computational chemistry. As a result, NWChemEx is built upon a general computational chemistry framework called the simulation development environment (SDE), which is designed with a focus on extensibility and interoperability. This article from the March/April 2019 issue of *Computing in Science & Engineering* describes the modular approach of the SDE and how it has been used to implement the self-consistent field algorithm within NWChemEx.

IEEE Annals of the History of Computing

The Role of Governments in the Spread of Novel Computing Devices in the Nineteenth and Early Twentieth Century United States

Nineteenth and early twentieth century American governments profoundly shaped diffusion of novel mathematical instruments. The federal government ran an office that judged what inventions were patentable and a legal system for those who defended or challenged patent rights. Governments at all levels employed inventors. Sometimes new laws required extensive calculations promoting invention and sale of computing instruments. Governments were customers for mathematical instruments ranging from harmonic analyzers to tabulating machines. Government buyers also offered testimonials to businesses. From the 1890s, antitrust legislation led to some federal government oversight of corporations. Historians usually rely on documents to tell this story. This article from the January–March 2019 issue of *IEEE Annals of the History of Computing* begins from objects—a few known only from patent descriptions, but most surviving in museum collections.

IEEE Computer Graphics and Applications

Visualization of Clouds and Atmospheric Air Flows

The IEEE Scientific Visualization Contest 2017 addressed the arising challenges in the visualization and analysis of atmospheric cloud-resolving simulations. In this article from the January/February 2019 issue of *IEEE Computer Graphics and Applications*, the authors utilize direct and indirect methods to represent atmospheric attributes such as cloud water content and air pressure, as well as employ Eulerian and Lagrangian techniques for air flow visualization.

IEEE Intelligent Systems

Employing Topical Relations in Semantic Analysis of Traffic Videos

Motion patterns in traffic video can be directly exploited to generate high-level descriptions of video content, which can be used for rule mining and abnormal event detection. The most recent and successful unsupervised methods for complex traffic scene analysis are based on topic models. In this article from the January/February 2019 issue of *IEEE Intelligent Systems*, a topic-related sparse topical coding framework is proposed for more effectively discovering motion patterns in traffic videos.

IEEE Internet Computing

A Service-Oriented Perspective on Blockchain Smart Contracts

Smart contracts turn blockchains into distributed computing platforms. This article from the January/February 2019 issue of *IEEE Internet Computing* studies whether smart contracts as implemented by a state-of-the-art blockchain technology may serve as a component technology for a computing paradigm like service-oriented computing in the blockchain, in order to foster reuse and increase cost-effectiveness.

IEEE Micro

Image Recognition Accelerator Design Using In-Memory Processing

This article from the January/February 2019 issue of *IEEE Micro* proposes a hardware accelerator design, called object recognition and classification hardware accelerator on resistive devices, which processes object recognition tasks inside emerging nonvolatile memory. The in-memory processing dramatically lowers the overhead of data movement, improving overall system efficiency. The proposed design accelerates key subtasks of image recognition, including text, face, pedestrian, and vehicle recognition. The evaluation shows significant improvements on performance and energy efficiency as compared to state-of-the-art processors and accelerators.

IEEE MultiMedia

Edge Caching and Computing in 5G for Mobile AR/VR and Tactile Internet

Both AR/VR and tactile internet applications require massive computational capability, high communication bandwidth, and ultra-low latency that cannot be provided with the current wireless mobile networks. By 2020, long-term evolution (LTE) networks will start to be replaced by fifth-generation (5G) networks. Edge caching and mobile edge computing are among the potential 5G technologies that bring content and computing resources close to the users, reducing latency and load on the backhaul. The aim of this article from the January–March 2019 issue of *IEEE MultiMedia* is to present current state-of-the-art research on edge caching and computing with a focus on AR/VR applications and tactile internet and to discuss applications, opportunities, and challenges in this emerging field.

IEEE Pervasive Computing

The Golden Age of Privacy?

The General Data Protection Regulation (GDPR), CoE 108þ, the California Consumer Privacy Act (CCPA); this has been a busy year for privacy legislators around the world. What will be the effects of all this lawmaking? Read more in the October–December 2018 issue of *IEEE Pervasive Computing*.

IEEE Security & Privacy

Public Auditing for Trusted Cloud Storage Services

Cloud storage can provide on-demand outsourcing of data services for organizations and individuals. However, because customers may not fully trust that cloud service providers meet their legal expectations for data security, techniques for auditing the cloud have attracted increasing attention. In this article from the January/February 2019 issue of *IEEE Security & Privacy*, the authors present an architecture of public data auditing, review existing methods or mechanisms for various auditing objectives, and discuss trends and possible future developments.

IEEE Software

Implementing Large-Scale Agile Frameworks: Challenges and Recommendations

On the basis of 13 agile transformation cases over 15 years, the authors of this article from the March/April 2019 issue of *IEEE Software* identify nine challenges associated with implementing Scaled Agile Framework, Scrum at Scale, Spotify, Large-Scale Scrum, Nexus, and other mixed or customized large-scale agile frameworks. These challenges should be considered by organizations aspiring to pursue a large-scale agile strategy. The authors also provide recommendations for practitioners and agile researchers.

IT Professional

Internet of Everything as a Platform for Extreme Automation

Internet of Everything (IoE) is expected to reinvent the business and the automation wheels altogether. From processes to business and manufacturing frameworks, everything is expected to change with the change in data available and the smart connectivity between people and machines for critical decision-making. IoE is bringing productivity and competitiveness to higher levels along with opening up many doors to new and exciting opportunities. IoE expands on the concept of the Internet of Things by connecting devices and people in one network. This connection goes beyond the basic M2M communications to enable a democratization of skill and how it is being delivered globally. An integral part of this is to be able to transmit touch in perceived real time, which is enabled by suitable robotics and haptics equipment at the edges, along with unprecedented communication network capabilities. Read more in the January/February 2019 issue of *IT Professional*.

Computing and Your Health

Computing is making its mark on healthcare. Technologies like electronic medical records, cloud computing, and smartphones have already transformed medical treatment and health data analytics. This issue of *ComputingEdge* describes innovative technologies that are helping improve health outcomes, whether through sophisticated virtual-reality systems or simple wearables.

"Enhancing Accessibility and Engagement for Those with Disabilities," from *IEEE Pervasive Computing*, details three projects aimed at improving quality of life for people with a specific impairment: a sound-awareness system for hearing loss, a proxemic audio interface for vision loss, and soft haptic toys for autism. "Using Virtual Reality to Increase Motivation in Poststroke Rehabilitation," from *IEEE Computer Graphics and Applications*, presents a system for helping stroke survivors regain limb movements through virtual-reality games that mimic real-life scenarios.

Wearables are another computing technology commonly used in healthcare applications. In *IEEE Intelligent Systems*' "How Will the Internet of Things Enable Augmented Personalized Health?," the authors argue that data from wearables and environmental sensors can be leveraged for better preventative healthcare. Meanwhile, *IEEE MultiMedia*'s "Health Media: From Multimedia Signals to Personal Health Insights" emphasizes both the benefits of low-cost wearable devices and the challenges in translating the data they collect into actionable information.

This *ComputingEdge* issue also discusses hardware's influence on software and vice versa. *IEEE Software*'s "Software Components" evaluates the case for mass-producing software parts—similar to how the electronics industry makes standardized resistors, capacitors, diodes, and transistors. *IEEE Micro*'s "Persistent Memory: Abstractions, Abstractions, and Abstractions" considers the effects of software design on nonvolatile main memory.

Finally, this *ComputingEdge* issue includes a *Computer* article about quantum computing. "Quantum Computer Scale-up" predicts a trajectory for qubit hardware that's based on the history of semiconductors and Moore's law.

Enhancing Accessibility and Engagement for Those with Disabilities

Kyle Rector
University of Iowa

Pervasive computing technology can enhance quality of life for those with disabilities by providing access to timely information and helping them to navigate their environment independently. Three research projects focusing on different impairments demonstrate the importance of including target users in the design and implementation of an accessibility system.

As pervasive computing devices increase in quantity and capability, they offer abundant new opportunities to help users, especially those with disabilities, in various environments. For example, the University of Maryland's Project Sidewalk (sidewalk.umiacs.umd.edu) is collecting information about accessibility problems in Washington, DC, and other cities from volunteers who virtually explore neighborhood streets and then using that data to "improve city planning, build accessibility-aware mapping tools, and train machine learning algorithms to automatically find accessibility issues."

In general, pervasive computing technology can enhance quality of life for those with disabilities by providing access to timely information and helping them to navigate their environment independently. In this article, I describe three research projects that focus on people with different impairments: deafness and hardness of hearing, blindness and low vision, and autism. In each case, including target users in the design and implementation of the accessibility system has been critical to its success.

LOCATION-INDEPENDENT SOUND AWARENESS FOR THE DEAF AND HARD OF HEARING

UbiEar[1] is a system that notifies those who are deaf or hard of hearing (DHH) about important sound events through their smartphone. To determine what types of sound events to detect and to guide their design, the researchers surveyed 60 DHH students from 10 to 26 years of age. The researchers found that the system should detect four social sound events (doorbell ringing, knocking on door, people crying, people coughing) and five early warning sound events (fire alarm, smoke alarm, kettle boiling whistle, microwave oven warning, and police siren). In addition, it needed to work in any location, report sound events with minimal delay and few errors, consume minimal battery power, and require very little user input.

UbiEar classifies sound events using a lightweight convolutional neural network (CNN). Training occurs in the cloud to preserve battery life, while sensing, sound detection, data processing, and sound recognition occur on the mobile device to minimize delay. The system optimizes each computation step on the device to reduce battery load. For example, it uses adaptive sensing to put the microphone to sleep when no sound of interest is detected.

The researchers developed a prototype user interface to learn what sounds students wanted UbiEar to detect and what types of notifications they wanted to receive. They then conducted a second user study with 86 DHH students and found that most of them wanted information on the same sound events from the initial survey conveyed through a mixture of notification mechanisms, including a flickering screen and flashlight notifications, with the ability to personalize these.

The researchers evaluated UbiEar use by the 86 students over 2 days with more than 10 million audio clips from their own collection (including 60,000 that were synthesized with different background noises), along with relevant audio clips from two other online datasets. The researchers found that the system achieves high accuracy (over 90 percent) for the nine acoustic events, and performs better than other baselines including a generic CNN, a deep neural network, a random forest classifier, and an AdaBoost model. They also found that over a 10-hour usage period, UbiEar only drains the phone battery by 16 percent. After the 2-day deployment, the students reported medium to high satisfaction with respect to notification delay and accuracy of the detected sound events.

UbiEar provides a great example of how work at the intersection of pervasive computing and accessibility can profoundly impact people with disabilities. The researchers had to consider both system engineering challenges, such as limited battery life, and the target users' needs throughout the design process to meet their goals. With this mobile technology, people who are DHH can become more aware of potentially important social and early warning sound events in almost any environment.

UbiEar provides a great example of how work at the intersection of pervasive computing and accessibility can profoundly impact people with disabilities.

PROXEMIC AUDIO INTERFACES FOR VISUALLY IMPAIRED ART ENGAGEMENT

I recently completed a project with Meredith Ringel Morris and Neel Joshi of Microsoft Research in collaboration with two artists.[2] As with UbiEar, our goal was to create a novel technology that delivers information to users with an impairment, in this case blindness and low vision, in an accessible format. However, we also aimed to enhance their engagement with the environment—in particular, museums.

People who are blind or have low vision have fewer opportunities to visit museums and have a fully accessible experience. They must go during scheduled accessible tours (which typically occur infrequently, such as once a month), walk through the museum with a sighted friend or guide, or learn a new audio guide technology. To address this problem, we wanted to create a more accessible way to explore artwork. Similar to the UbiEar team, we consulted with stakeholders to inform our design—in our case, people who are blind or have low vision, artists, and museum staff. Based on these interviews we decided that, to create an aesthetically satisfying experience, we needed a technology that could verbally describe a work of art as well as convey emotion and mood.

After exploring various interfaces that could support such an experience, we opted to create a proxemic audio interface. Proxemics is a theory introduced by anthropologist Edward T. Hall in the 1960s that people place themselves at a certain distance from others by relationship type. Visual proxemic interfaces have adopted this concept by increasing the level of detail or information about a target object as a person moves closer to it. However, previous research had not explored proxemic audio interfaces.

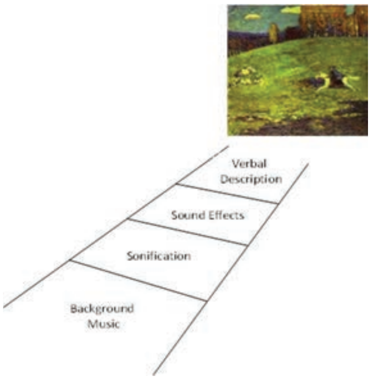

Figure 1. Eyes-Free Art was a prototype proxemic audio interface for museums that increased detail about a work of art (in this case, Wassily Kandinsky's *The Blue Rider*) in four zones as one moved closer to it: from background music to a sonification of the work's colors to sound effects related to objects in the piece to a verbal description of it.

We designed a prototype interface called Eyes-Free Art using a Microsoft Kinect that increased audio detail about a work of art in four zones as a person moved closer to it: from background music to a novel sonification of the work's colors to sound effects related to objects in the piece to a verbal description of it (Figure 1). Users could hear more details by moving their body and hands. We implemented this interface for five paintings with various colors and objects, and conducted a lab study with 13 blind and low-vision people. The subjects appreciated the interface's interactivity, but their feedback also led us to make several changes. These included ensuring that the transitions in the audio presentations between zones were smoother and, in contrast to visual proxemics interfaces, moving the detailed verbal description to the most distant zone to provide more context for the rest of the experience.

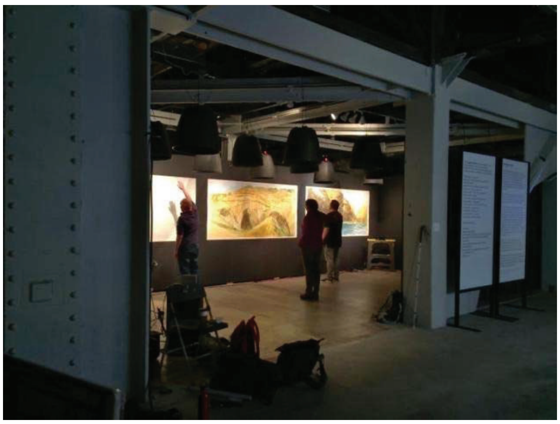

Figure 2. *The Oregon Project* was an interactive art installation inspired by Eyes-Free Art that used Kinect-based proxemic audio interfaces and various sound recordings to engage visually impaired as well as sighted patrons.

Following our lab study, we collaborated with artists Keith Salmon and Dan Thornton to create an accessible interactive art installation called *The Oregon Project* with multiple paintings that numerous people could explore simultaneously (Figure 2). Four Kinects tracked people in the room, and parabolic speakers delivered the audio presentations. Before entering the room, the visitors heard a detailed verbal description of *The Oregon Project* by the artists. As one moved closer to individual paintings, audio presentations seamlessly transitioned from environmental sounds evoking the Oregon plains to paint-stroke sounds, which visitors could control with their body and arms. *The Oregon Project* was deployed at a museum in the US and in the UK. Sighted as well as visually impaired patrons commented that, through the proxemic audio interface, they become more immersed in the artwork.

SOFT HAPTIC TOYS FOR CHILDREN WITH AUTISM SPECTRUM DISORDER

The third research project focused on creating a pervasive technology for children with autism spectrum disorder to enhance engagement with their environment. Jinsil Hwaryoung Seo, Pavithra Aravindan, and Annie Sungkajun of Texas A&M University implemented a six-month design process to create soft haptic toys for autistic children who suffer from over- or under-sensory development and stimulation.[3]

Children on the spectrum actively participated in design exploration sessions, which considered different materials, shapes, and interaction modes. They preferred the softest fabrics, such as microfleece and cotton, along with animal shapes. In addition, the researchers found that vibrations were more helpful in relaxing the children during times of stress than other types of feedback such as lights and sounds.

The soft haptic toys were constructed using the LilyPad Arduino, conductive thread, and other interactive components including pressure sensors and haptic motors. After the design exploration sessions, five children with high-functioning autism created their own toys. The researchers reported that the children were engaged in the building process and happy with the toys that they had created.

CONCLUSION

All three research projects demonstrate the importance of including target users while developing a pervasive computing technology to enhance accessibility and engagement for those with disabilities. Both users and designers benefit from this collaboration, leading to a more effective and impactful system.

REFERENCES

1. L. Sicong et al., "UbiEar: Bringing Location-independent Sound Awareness to the Hard-of-Hearing People with Smartphones," *Proc. ACM on Interactive, Mobile, Wearable, and Ubiquitous Technologies*, vol. 1, no. 2, 2017; doi.org/10.1145/3090082.
2. K. Rector et al., "Eyes-Free Art: Exploring Proxemic Audio Interfaces for Blind and Low Vision Art Engagement," *Proc. ACM on Interactive, Mobile, Wearable, and Ubiquitous Technologies*, vol. 1, no. 3, 2017; doi.org/10.1145/3130958.
3. J.H. Seo, P. Aravindan, and A. Sungkajun, "Toward Creative Engagement of Soft Haptic Toys with Children with Autism Spectrum Disorder," *Proc. 2017 ACM SIGCHI Conf. Creativity and Cognition* (C&C), 2017, pp. 75–79.

ABOUT THE AUTHOR

Kyle Rector is an assistant professor of computer science and a member of the Public Digital Arts Cluster at the University of Iowa. Contact her at kyle-rector@uiowa.edu.

This article originally appeared in
IEEE Pervasive Computing, *vol. 17, no. 1, 2018.*

Department

Using Virtual Reality to Increase Motivation in Poststroke Rehabilitation

VR Therapeutic Mini-Games Help in Poststroke Recovery

Paulo Dias
DETI/IEETA Universidade de Aveiro

Ricardo Silva
DETI/IEETA Universidade de Aveiro

Paula Amorim
Centro de Medicina de Reabilitação da Região Centro—Rovisco Pais

Jorge Laíns
Centro de Medicina de Reabilitação da Região Centro—Rovisco Pais

Eulália Roque
Centro de Medicina de Reabilitação da Região Centro—Rovisco Pais

Inês Serôdio
Centro de Medicina de Reabilitação da Região Centro—Rovisco Pais

Fátima Pereira
Centro de Medicina de Reabilitação da Região Centro—Rovisco Pais

Beatriz Sousa Santos
DETI/IEETA Universidade de Aveiro

Editor: Mike Potel, potel@wildcrest.com

***Abstract*—Virtual reality (VR) applications meet fundamental principles of rehabilitation: intensity, task oriented training, biofeedback, environments rich in stimuli, and motivation, all pivotal factors for the success of rehabilitation programs. This paper describes the development process of a set of VR minigames developed to increase the motivation of stroke patients while performing repetitive upper limb movements.**

TECHNOLOGY MAY SIGNIFICANTLY improve the lives of people suffering from incapacity or deficiency affecting millions worldwide. Virtual reality (VR) is already used to help patients endure pain and disease treatment[1–4] as well as recover from stroke,[5,6] among other applications in medicine. VR has a significant potential for rehabilitation[7,8] as it allows the creation of virtual environments (VEs) providing multiple stimuli and fostering the improvement of motor and cognitive capacities while motivating and engaging the patients. Moreover, VR applications may meet the four basic principles of rehabilitation: intensity, task oriented training, biofeedback,

Digital Object Identifier 10.1109/MCG.2018.2875630
Date of current version 6 March 2019.

and motivation, all pivotal factors for the success of rehabilitation programs.[6,9,10,11]

The following benefits of using VR in rehabilitation have been reported in the literature:[7] better performance, improvement of the affected limb and cognitive functions, neuroplasticity stimulation, and greater autonomy in the daily life activities, while increasing the patients' motivation and collaboration during the rehabilitation program. In particular, some authors have "found evidence that the use of VR and interactive video gaming may be beneficial in improving upper limb function and ADL (Activities of Daily Living) function when used as an adjunct to usual care (to increase overall therapy time) or when compared with the same dose of conventional therapy."[7] This makes VR an exciting tool in the future of therapy, "not only because it was proven to be effective among sick and healthy subjects, but also because it had very little side-effect and was much safer than other aggressive or offensive therapies."[8]

Recently, affordable sensors developed by the gaming industry have been explored for rehabilitation.[6,12] This synergy between benefits and affordable technology makes VR applications a natural approach for stroke rehabilitation, one of the main causes of incapacity worldwide. Aware of this potential, and concerned with the lack of motivation of stroke patients while performing repetitive upper limb movements in acute, subacute, and chronic phases, a group of professionals at "Centro de Medicina de Reabilitação da Região Centro—Rovisco Pais," a National Rehabilitation Center in Portugal, contacted the Universidade de Aveiro to develop VR therapeutic serious games aimed at increasing motivation by providing everyday life context to the movements. Several VR applications were developed using a Leap Motion sensor (www.leapmotion.com) to track upper limb movements. These applications help patients perform relevant shoulder, arm, and hand movements, while immersing them in an informal game-like VE. This paper describes the development of the applications and the main results of a study involving a group of 12 patients of the rehabilitation center.

VR APPLICATIONS

With the goal of maximizing the usefulness and efficacy of the applications and taking into consideration the specific nature of their users and context of use, the initial phase of the process involved a series of visits to the rehabilitation center and meetings, first with a group of interested physiatrists, and later also with physical and occupational therapists. These meetings helped establish a common ground of mutual understanding of what patients need and what the technology can provide, thenceforth facilitating the communication between the teams. The first outcome of these meetings was the awareness that the ideal VR platform should encompass not only a set of "minigames" to motivate patients during the essential, but tedious sessions of upper limb rehabilitation (the initial goal), but also the possibility of personalization of the games as well as remote monitoring of the patients' progress, allowing a better follow-up of the patients' evolution beyond the rehabilitation center. This is a very important feature allowing patients to actively participate in their program at home. As a result of this initial phase, a set of decisions concerning the design and implementation of the applications were made, and the physiatrists and therapists stayed involved during the process, regularly giving feedback and helping establish intermediate goals.

> The VR platform should encompass not only a set of minigames to motivate patients, but also the possibility of personalization and remote monitoring of patient progress.

The Leap Motion controller was selected as the sensor to monitor both coarse gestures (shoulder or elbow movement, detected due to change of hand position) and fine movements (finger pinches) since it detects the position, orientation, and current state of the hand. The games were developed in Unity3D (unity3d.com). This platform allows the creation of VEs as well as game logic and facilitates the virtual world creation interface as well as native integration with an Oculus Rift DK2 Head Mounted Display (www.oculus.com/rift) and the official Leap Motion SDK package.

The system includes a backend server controlling access to the database and the front-end three-dimensional applications used by patients, as well as a configuration web page. This allows for storage and management of game

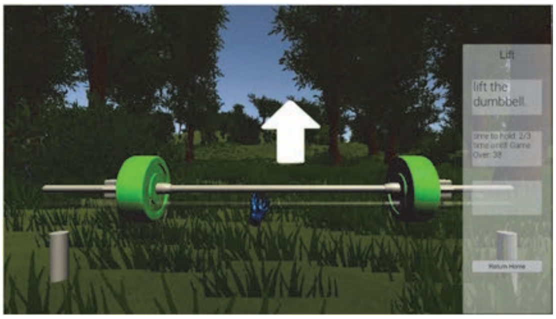

Figure 1. "Lift" game: user lifts barbell to a specified height a target number of times.

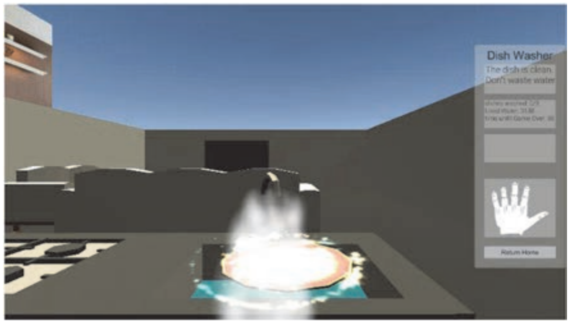

Figure 3. "Dish Washer" game: user opens and closes the hand to wash the dish a number of times.

configuration data (game instance, number of iterations, maximum completion time, difficulty level, and other aspects of the game), and game results (task completion, time elapsed, and specific values concerning the patient's movements as the longest distance reached).

Our first goal was to define which gestures were relevant for the exercises to be performed by patients during the games. The "Enjalbert Test" was selected as the basis for the applications to be developed since it was already used to evaluate patients' progress at the rehabilitation center.[13] The test, a five-level scale, is used to access the current state of the upper limb movement recovery for a poststroke victim and includes different movements, ranging from 0 (no upper limb movement) to 5 (fine pincer movements with all fingers).

- Lifting and holding the hand in place (shoulder).
- Bringing the hand to the mouth (shoulder and elbow).
- Opening and closing the hand (hand).

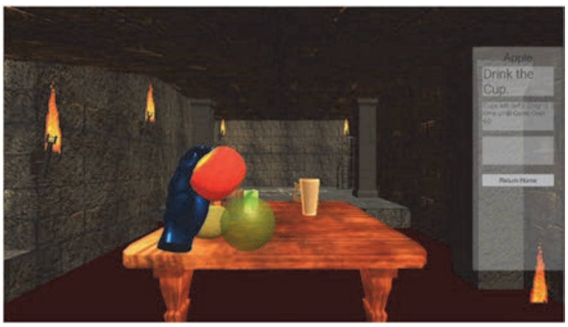

Figure 2. "Apple Eater" game: user takes an apple to the mouth a number of times.

- Executing fine pinch movements with the index and middle fingers (hand).
- Executing fine pincer movements with the ring and pinky fingers (hand).

An important requirement was that the games should evoke real life situations and be aimed at helping patients recover capacities for an independent life. Thus, it was decided to develop five minigames, focused on movements involved in progressing through the Enjalbert scale. The games developed to exercise the first three gestures passed a first round of tests with patients (in the same order as the list above).

The games should evoke real life situations and be aimed at helping patients recover capacities for an independent life.

- Lift: The patient should lift a barbell above a specified line (see Figure 1) and hold it for a predefined time before bringing it back down. This action should be repeated for a predefined number of times.
- Apple eater: The patient should reach one of the two apples (see Figure 2) on a table and bring it to the mouth.
- Dish washer: The patient should wash the dishes, opening and closing their hand to turn ON and OFF the sink's faucet (see Figure 3). The patient must keep the hand open until the dish is entirely clean.

Two more games were developed to exercise "finger pinch" movements that required users to pick objects from a box and drop them on a table

Figure 4. VR system used at the rehabilitation center. 1) Computer. 2) Monitor. 3) Oculus Rift head mounted display. 4) Leap Motion controller. 5) Speaker.

using different pinch gestures. However, due to the unreliability of the Leap Motion controller for very fine gestures, doctors concluded these games were not responsive enough to be tested with patients.

Data such as the duration of each movement, number of repetitions, height of the barbell line, number of apples on each side of the table, number of dishes, or what is considered an open hand are configured through a backend web page. A calibration application was also developed to configure the games according to the patient's condition, essential for allowing the patients to accomplish the task. With this application the limits for values such as ''maximum height when lifting arm'' or ''maximum hand opening'' can be set for each patient and updated according to the patients progress along their rehabilitation program.

As part of the development process, several rounds of preliminary tests were performed at the rehabilitation center with the help of doctors, therapists, and volunteer patients who played the games. This formative evaluation phase had a twofold purpose: Identify and correct possible limitations of the applications and assess whether the patients liked and were motivated by the minigames. Some modifications were made, mostly regarding the distance between the virtual hand resting position and the interaction objects, since in an initial phase applications were only tested by users with full control of their upper limb and these issues were not noticed. Another relevant improvement was the addition of a score and a ''success'' sound effect at the completion of the task, in a way to provide positive feedback and encouragement, and allow for competition among patients, features that were considered important to increase motivation. On the other hand, when patients did not attain the goal, discouraging sounds or negative messages were not given so as to avoid patient frustration.

Beyond testing the minigames, these preliminary testing sessions were also meant to instruct the therapists on how to use the system, especially the configuration settings, as they would be the main users.

USER STUDY

A VR system was installed at the rehabilitation center to enable its patients to use the developed applications. The VR setup is composed by the following elements, as shown in Figure 4.

- A desktop computer to run the applications and local backend server (marked "1" in the figure).
- A 4k definition monitor to display the VE, when running the applications in a nonimmersive setting ("2").
- An Oculus Rift DK2 HMD (head mounted display) to display the VE, when running the applications in a fully immersive setting ("3").
- A Leap Motion controller to track the position and orientation of the patient's hands, so they can be represented and used in the VE ("4").
- A speaker positioned in front of the patient to provide audio feedback ("5").

To evaluate the developed minigames, a pilot study was conducted after a formal authorization by the rehabilitation center ethics committee and a careful selection of the patients that should participate. The aim of this study was to establish which selection standards should be applied regarding which patients could use the applications and benefit from them, as well as to obtain data regarding the patients' satisfaction with the games.

The main questions to be answered by our study were:

1) At what level of recovery could the patients start using the minigames?

2) Which exclusion criteria should be used?
3) Which particular stroke sequelae cause unusual results in a patient's capability and enjoyment when playing?
4) Is this type of treatment well accepted by the patients?
5) Is there a preference regarding the level of immersion (nonimmersive versus full immersion)?

A group of 12 patients (six female) aged between 39 and 71 in several phases of recovery and suffering from different stroke sequelae were selected to test the applications, using both the immersive and nonimmersive versions of the games. The patients used the applications while seated and then answered a questionnaire regarding their satisfaction with the minigames, always accompanied by a developer and a therapist.

First, the patient was instructed about the test and the equipment he/she would be using. The patient then played the minigames twice, using the computer screen as the display and the Oculus Rift DK2 HMD. To prevent bias, half the patients used the nonimmersive version of the applications first, while the other half started by using the HMD.

After the games were concluded (successfully or not) the patient answered orally part of the questionnaire (concerned with familiarity with technology; nonimmersive versus fully immersive and general questions). The remaining sections of the questionnaire (doctor/therapist credentials; patient information; occupation therapy) were answered by the doctors and therapists.

RESULTS AND DISCUSSION

Although most patients were not familiar with computer games (9 out of 12 had never played videogames) or VR (10 out of 12), the minigames were well accepted both in the rate of success and in satisfaction. Only two patients were not able to successfully complete all three minigames and only one was not able to complete any of them. All the patients who were able to play came away satisfied, claiming to have enjoyed the experience and expressing interest in including VR as part of the rehabilitation therapy. "Lift" was the preferred game, followed by "Dish Washer."

Although no significant differences in performance or acceptance between the fully immersive and nonimmersive games were noted, when asked, most patients claimed to prefer the fully immersive version of the system. Two patients preferred the nonimmersive version of the minigames. One of the patients had never used a computer before and found full immersion to be too invasive. The other patient suffered from proprioceptive sensitivity and, as explained by the doctors, stroke victims with this particular sequela feel the need to look at their hand in order to execute the movements. Thus, not being able to see the real hand when using the HMD may have caused the patient to feel less attracted to the fully immersive version, despite being able to successfully play in both versions.

No patients expressed feeling any type of cybersickness during or after playing. This was expected as none of the minigames involves any virtual full body movements, or rapidly changing images, which are important causes for this kind of side effect.

When asked in which setting, individual or social, they would prefer to play the games, the results were approximately the same for both. It was also noted that one patient participating in the study had previously played the minigames, during the preliminary tests, with greater success. Although being able to complete all three games both times, during the study the patient was suffering from depression, which was considered to be a plausible cause for the decrease in performance.

No hardware related issues specific to Oculus Rift were observed; however, there are visual and cognitive stroke sequelae (besides depression) that may hamper the usage of immersive VR in patients' therapy. Examples of such conditions are hemineglet and assomatognosia, involving deficits in recognizing the hemispace and hemibody contralateral to the injured brain hemisphere

Some limitations were found regarding the use of the Leap Motion sensor as a tracker. Two specific issues were considered relevant: the position of the sensor on top of the table proved too hard to reach by patients in early phases of recovery. This obstacle was overcome by placing a board on the patient's lap and positioning the sensor on it. In the "Apple Eater" game, because the patient's mouth position in the virtual environment was static, unless the patient kept his/her back straight throughout the full exercise, this position would no longer correspond to the actual mouth area of the patient. This issue was amplified by the fact that the patients would lean forward to reach the objects, and was alleviated by reminding the patients to keep their back straight during the procedure.

CONCLUSION

Overall, the potential use of the minigames in occupational therapy in poststroke rehabilitation was very well received by patients, doctors, and therapists, with its major benefit being the increase in a patient's motivation for recovery through the use of fun and relaxed environments, which successfully distract the patient from the dull clinical setting at an affordable cost.

The collaboration continues with the development of more applications, both aimed at upper limb movement recovery and rehabilitation for other stroke sequelae and further tests to introduce this approach in the routine therapy of the rehabilitation center at least in some phases of their recovery. The next phase will be the evaluation of the efficacy of this approach as an additional therapeutic instrument in the rehabilitation of poststroke patients in acute, subacute, and chronic phases, through a longitudinal study involving a larger number of patients with a wider variety of conditions, both at the rehabilitation center and at home. For instance, if gender or age correlations were noticeable, this study would provide guidelines on how to use VR with different patients.

Augmented reality based physical therapy games using smartphones might also be a promising direction as they lower the barrier to greater home-based use and technological literacy of the population is increasing. Compared to conventional approaches, AR alternatives allow adapting the exercises to the patients' interests and habits potentially increasing their motivation. Nevertheless, immersive VR-based games may be ultimately more engaging.

ACKNOWLEDGMENTS

The authors would like to thank the students and patients who participated in the tests and user study. This work was supported in part by FCT, under its Project UID/CEC/00127/2013.

REFERENCES

1. R. M. Baños *et al.*, "A positive psychological intervention using virtual reality for patients with advanced cancer in a hospital setting: A pilot study to assess feasibility," *Supportive Care Cancer*, vol. 21, no. 1, pp. 263–270, 2013.
2. C. V. Maani *et al.*, "Virtual reality pain control during burn wound debridement of combat-related burn injuries using robot-like arm mounted VR goggles," *J. Trauma, Injury, Infection, Critical Care*, vol. 71, supplement S125–S130, 2011.
3. D. R. Patterson, M. P. Jensen, S. A. Wiechman, and S. R. Sharar, "Virtual reality hypnosis for pain associated with recovery from physical Trauma," *Int. J. Clin. Exp. Hypnosis* vol. 58, no. 3, pp. 288–300, 2010.
4. S. M. Schneider and L. E. Hood, "Virtual reality: A distraction intervention for chemotherapy," *Oncol. Nursing Forum*, vol. 34, no. 1, pp. 39–46, 2007.
5. S. Cho *et al.*, "Development of virtual reality proprioceptive rehabilitation system for stroke patients," *Comput. Methods Programs Biomed.*, vol. 113, no. 1, pp. 258–265, 2014.
6. M. Covarrubias, M. Bordegoni, M. Rosini, E. Guanziroli, U. Cugini, and F. Molteni, "VR system for rehabilitation based on hand gestural and olfactory interaction," in *Proc. 21st ACM Symp. Virtual Reality Softw. Technol.*, 2015, pp. 117–120.
7. K. E. Laver, B. Lange, S. George, J. E. Deutsch, G. Saposnik, and M. Crotty, "Virtual reality for stroke rehabilitation," *Cochrane Database Systematic Rev.*, 2017. Accessed: Feb. 26, 2018, [Online]. Available: http://doi.wiley.com/10.1002/14651858.CD008349.pub4
8. Z. Liu, S. Wangluo, and H. Dong, "Advances and tendencies: A review of recent studies on virtual reality for pain management," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Cham Switzerland: Springer, 2016, pp. 512–520.
9. M. Covarrubias, A. Mansutti, M. Bordegoni, and U. Cugini, *Interacting Game and Haptic System Based on Point-Based Approach for Assisting Patients After Stroke*. Cham Switzerland: Springer, 2014, pp. 289–296.

10. P. Kiper, L. Piron, A. Turolla, J. Stożek, and P. Tonin, "The effectiveness of reinforced feedback in virtual environment in the first 12 months after stroke," *Neurologia i neurochirurgia polska*, vol. 45, no. 5, pp. 436–44, 2011.
11. L. Piron, *et al.*, "Exercises for paretic upper limb after stroke: A combined virtual-reality and telemedicine approach," *J. Rehabil. Med.*, vol. 41, no. 12, pp. 1016–102, 2009.
12. A. D. Gama, P. Fallavollita, V. Teichrieb, and N. Navab, "Motor rehabilitation using kinect: A systematic review," *Games Health J.*, vol. 4, no. 2, pp. 123–135, 2015.
13. M. Enjalbert, J. Pelissier, and D. Blind, "Classification fonctionnelle de la préhension chez l'hemiplégic adulte," *Hémiplégie vasculaire de l'adulte et medicine de rééducation*, J. Pélissier Ed., Masson, Paris, 1998, pp. 212–223.

Paulo Dias is an Assistant Professor with the Department of Electronics Telecommunications and Informatics, University of Aveiro and a Researcher with the Institute of Electronics and Informatics Engineering of Aveiro. His research interests include visual computing and robotics. Contact him at paulo.dias@ua.pt.

Ricardo Silva received the MSc degree in computer engineering from the University of Aveiro, Portugal, and he developed this work in the scope of his dissertation. Contact him at ricardojsilva@ua.pt.

Paula Amorim is a Physical and Rehabilitation Physician with the Portugal Centre Region Rehabilitation Medicine Centre Rovisco Pais. Her current research interests include telereabilitation, virtual and augmented reality. Contact her at pamorim@roviscopais.min-saude.pt.

Jorge Lains is a Physical and Rehabilitation Physician with the Portugal Centre Region Rehabilitation Medicine Centre Rovisco Pais. Contact him at jorgelains@roviscopais.min-saude.pt.

Eulália Roque is an occupational therapist at the Portugal Centre Region Rehabilitation Medicine Centre Rovisco Pais. Contact them at toc@roviscopais.min-saude.pt.

Inês Serôdio is an occupational therapist with the Portugal Centre Region Rehabilitation Medicine Centre Rovisco Pais. Contact her at toc@roviscopais.min-saude.pt.

Fátima Pereira is an occupational therapist with the Portugal Centre Region Rehabilitation Medicine Centre Rovisco Pais. Contact her at toc@roviscopais.min-saude.pt.

Beatriz Sousa Santos is an Associate Professor with the Department of Electronics Telecommunications and Informatics, the University of Aveiro and a Researcher with the Institute of Electronics and Informatics Engineering of Aveiro, Portugal. Her research interests include virtual and augmented reality. Contact her at bss@ua.pt.

Contact department editor Mike Potel at potel@wildcrest.com.

This article originally appeared in IEEE Computer Graphics and Applications, *vol. 39, no. 1, 2019.*

How Will the Internet of Things Enable Augmented Personalized Health?

Amit Sheth
Kno.e.sis-Wright State University

Utkarshani Jaimini
Kno.e.sis-Wright State University

Hong Yung Yip
Kno.e.sis-Wright State University

The Internet of Things refers to network-enabled technologies, including mobile and wearable devices, which are capable of sensing and actuation as well as interaction and communication with other similar devices over the Internet. The IoT is profoundly redefining the way we create, consume, and share information. Ordinary citizens increasingly use these technologies to track their sleep, food intake, activity, vital signs, and other physiological statuses. This activity is complemented by IoT systems that continuously collect and process environment-related data that has a bearing on human health. This synergy has created an opportunity for a new generation of healthcare solutions.

The paradigm shift from reactive medicine to proactive and preventive medicine is primarily motivated by economic imperatives such as the rising cost of healthcare, as well as continued improvements on quality of life and longevity. According to the Centers for Medicare and Medicaid Services (CMS), in 2016 the cost of healthcare in the US reached $3.6 trillion and is expected to increase to $5.5 trillion by 2025 (www.advisory.com/daily-briefing/2017/02/16/spending-growth). On the other hand, the global smart healthcare industry is expected to reach $169.30 billion by 2020.

It's also projected that by 2019, 87 percent of US healthcare organizations will have adopted Internet of Things (IoT) technology (www.i-scoop.eu/internet-of-things-guide/internet-things-healthcare), of which 73 percent will be used to reduce cost, and 64 percent will be for patient monitoring.

IoT data itself isn't adequate to understand an individual's health and associated aspects of well-being and fitness; it's usually necessary to look at that individual's clinical record and behavioral information, as well as social and environmental information affecting that individual. Interpreting how well a patient is doing also requires looking at his adherence to respective health objectives, application of relevant clinical knowledge and desired outcomes, such as the patient's preference for quality of life versus longevity and expert knowledge.

Augmented Personalized Healthcare (APH) is a vision (http://wiki.knoesis.org/index.php/Augmented_Personalized_Health:_How_Smart_Data_with_IoTs_and_AI_is_about_to_Change_Healthcare)[1] for exploiting the extensive variety of relevant data and medical knowledge using artificial intelligence (AI) techniques to extend and enhance human health and well-being. It anticipates the use of physical, cyber, and social data obtained from wearables and IoT devices; clinical information including electronic medical records (EMRs); mobile applications supporting targeted interactions and engagement with the patients; and web-based information including web services (such as those providing health-relevant data on allergens and air quality), social media (such as posts by patients with similar concerns and conditions), and extensive online knowledge bases of clinical practice and medicine. Data can be collected at the personal, public, and population levels, and be combined with knowledge that affects human health. Augmentation refers to aggregating this data and converting into actionable information that can improve health-related outcomes through better and more timely decisions. This embodiment of APH is an entirely new approach to human healthcare in comparison with the current episodic system of periodic care primarily centered around healthcare establishments (such as clinics, hospitals, and labs).

APH involves continuous monitoring, engagement, and health management in which, instead of treating a patient for a disease, the focus shifts to involving the patient in preventing disease, predicting possible adverse outcomes and intervening to mitigate or eliminate them through proactive measures, and trying to keep citizens healthy and fit with continuous lifestyle changes. Rather than only focusing on the management of chronic conditions, APH proposes a holistic approach for improving the overall quality of life.

Patient-generated health data (PGHD) is the heart of APH. It's primarily generated by IoT devices and captures the digital footprint representative of patients' health over time with finer details that are distinct from the data generated in clinical settings through EMRs and personal health records (PHRs). The two main IoT categories for patient health monitoring are wearable sensors and environmental sensors. Wearable sensors are portable sensors that patients wear most, if not all, of the time. These close-vicinity wearable sensors monitor patients' physiological markers, such as heart rate, breathing rate, and blood pressure. They're designed to integrate into patients' daily routines to enable passive and continuous sensing and monitoring for timely interventions. Environmental sensors, on the other hand, are sensors that collect environmental data relevant to patients. These sensors are normally not portable, but can sometimes provide critical information for health management. For example, weather data, such as humidity, pollen index, and air quality, are important for managing asthma. However, this data provides a coarse population-level measure, and wouldn't account for the differences of individual patients. To mitigate this, different sensors can be utilized as a complementarity. For example, Foobot (https://foobot.io) monitors indoor air quality and reflects a closer overview of a patient's environment. Hence, these sensors enable personalization and allow both physicians and patients to monitor asthma at a finer level.

STAGES OF TECHNOLOGY-ENABLED HEALTH AUGMENTATION

In this section, we review various stages of augmented health management strategies using APH technology.

Self-Monitoring

Currently, doctors see patients infrequently or as needed for new conditions; or, for people with chronic conditions and disease, they are at well-defined time intervals (monthly, quarterly, etc.), depending on the established medical protocols and severity of the medical condition. A doctor's understanding of a patient's condition often comes primarily from the patient's self-description (self-reporting) in addition to the observations gathered by the clinician during the visit. This has limitations, of course, as sometimes not all significant events or issues are recalled at the time of the visit. In addition, the exact timing, location, and reasons for the triggering event might not be available. With continuous monitoring using IoT-enabled sensor devices, wearables, and a periodically administered contextually relevant questionnaire, we can better capture relevant aspects of a patient's surroundings, diet, activities, and other factors related to health. All of these aspects, when analyzed, can help to determine possible and precise contributing factors of patients' conditions or level of well-being. PGHD plays an important role in supplementing existing avenues for collecting clinical data and filling in information gaps on a routine basis, thus generating a more comprehensive picture of long-term patient health (www.healthit.gov/policy-researchers-implementers/patient-generated-health-data).

Self-Appraisal

Self-appraisal describes the patient's ability to evaluate the relevance of a variety of data and observations within the context of his or her general health objectives or specific health concerns. Wearable devices are used to keep track of patients' day-to-day activities. However, there's a big gap between simply having the access to relevant data generated from self-monitoring and being able to analyze and interpret the data in a useful way. Patients are interested in understanding if they are keeping up with progress toward their health goals. Consider, for example, using a Fitbit[2] to measure the number of steps taken each day and the quality and duration of sleep. Is this data helping the patient fulfill their desired objective or do they need to do something more? What is the distinction between expending 1,700 calories and 2,200 calories per day vis-à-vis the objective to shed 5 pounds in next 3 months to improve the management of diabetes? What about the existence and impact of any abnormal behavior on body activity? For example, for patients taking asthma meds, if their Fitbit shows a heart rate of 100+ while asleep, is that a serious enough condition to require clinical consultation?

Self-Management

Self-management refers to the patient's decisions and behaviors that impact management of their chronic conditions. Generally, the impact that a patient may seek is getting back in line with the prescribed medical care plan or agreed upon health objective. Patients empowered with IoT-generated PGHD have a better sense of their health condition and can make informed decisions about care as opposed to episodic clinical visits in which patients aren't aware of their state until diagnosed. An APH technology that intends to support self-management is expected to identify actionable information, such as increasing weight-bearing exercise or reducing consumption of energy-dense food. An APH technology can aid patients by providing alerts about potential triggers (such as high pollen counts) or feedback on adherence (such as unexpected weight gain or not meeting activity targets), which can be used to keep patients on course. This improves the effective use of IoT for both data collection and relevant data/analysis/alert access by the patient. It can then provide alternatives to the patient to take steps to better adhere to the physician-specified care plan to reduce adverse impact due to deviation from the plan or improve the outcome of the objective (for example, an APH technology used to promote self-management for patients who are obese can use IoT capabilities, such as activity monitoring and fluid consumption to also measure increases in activity level and targeted water intake if weight gain continues after use of oral steroid has ended).

Intervention

The next step up in health management is clinical intervention, which includes a change in the care plan prescribed by the clinician. An APH technology can use the data it gathers to help clinicians provide PGHD, environmental, and other data as well as corresponding analysis and interpretation to help evaluate and adjust a patient's clinical plan. The timely analysis of IoT data can yield insights for early intervention before a patient's situation deteriorates. In the case of kHealth Asthma, which developed an APH technology for managing asthma in children, the observed deterioration of asthma symptoms through PGHD can suggest change in medication or its dosage, develop trigger-avoidance plans, and so on. The IoT data collected from individuals can help clinicians develop a personalized, patient-centric recommendations for use in the healthcare system with implicit feedback and support adherence to physician-prescribed protocols.

Disease Progression Tracking and Prediction

Going beyond immediate and short-term management of health concerns, it would be highly rewarding on the individual and public health levels if the longitudinal collection of personalized health data including PGHD and environmental data could facilitate tracking how a disease is progressing, predicting significant changes in health status, and identifying and taking remedial actions. For example, for a patient who is pre-diabetic and has an A1C score higher than 6, it would be highly valuable to be able to track the score's changes and issue an alert when it has reached diabetic status (A1C $\geq$ 6.5), thus predicting the high probability that a patient becomes diabetic and requires insulin treatment. For an asthmatic patient who is overweight and is on long-term steroid medication—which may lead to a number of adverse situations including higher energy intake—it would be important to track associated weight gain and compute the probability of worsening of asthma severity. A more straightforward strategy would be for the clinician to periodically review the patient's data and make an educated judgment on the disease progression. A more advanced strategy would be to use the personalized health data and analyze it vis-à-vis published clinical studies and longitudinal data collected with relevant cohort population.

The objective is to devise more proactive interventions and incorporate more nonmedical solutions, such as lifestyle changes that are often very effective but take a longer time to show benefit. Individual and public health will greatly improve based on what we learn from such strategies, which will become evidence-based enhancements of widely accepted clinical pathways and protocols.

APH SHOWCASE AND APPLICATION SCENARIO

The knowledge-enabled healthcare (kHealth) initiative at Kno.e.sis is an example of an APH framework to enhance decision making and improve health, fitness, and well-being (http://bit.ly/kAsthma). The early prototyping and testing involved kHealth–ADHF, a mobile app and a sensor kit designed to reduce readmission in patients with Active Decompensated Heart Failure (ADHF). kHealth–ADHF involved continuous monitoring using targeted questions driven by application-specific (cardiovascular) knowledge as well as sensors to record blood pressure, heart rate, and body weight. These measurements provided observational data via Bluetooth to the mobile app, which also asked the patient pertinent questions and analyzed answers and collected data and then generated alerts.

A follow-on application, kHealth–asthma, designed to better control asthma in children, extends physical data collection using cyber and social data. Figure 1 shows an instance of the kHealth–Asthma application collecting multimodal data to monitor pediatric asthma, which is a multifactorial and multifaceted disease. We are running a trial with 200 patients with asthma, collecting possibly the broadest modality of data, with an average 124 readings collected per day (2 tablet readings per day, 24 Fitbit readings per day, 2 Peak flow readings, 96 Foobot readings per day), for a duration of 1, 3, and 6 months. This vast amount of multimodal and multisensory data poses a big data challenge (due to the data variety and associated challenges with integration) in comparison to other mHealth studies like Google Verily, IBM, and Swiss startup

Docdok.health project, as well as the Stanford wearable study that deals with fewer modalities and a smaller sample size. To address the aforementioned problem, kHealth supports contextual (condition-specific) annotation, integration, and interpretation of sensor data using Semantic Sensor Network (SSN) ontology. Furthermore, kHealth supports contextualized actionable feature selection in PGHD to generate Smart Data using SSN and domain-specific knowledge sources. Utilization of Smart Data provides timely medical intervention and remediation measures. In a broad sense, a knowledge graph is a knowledge base that provides semantic annotation using its characteristic functionalities like fact extraction, named entity recognition, relationship identification, locale-specific information, event extraction, and intent identification to enrich information.

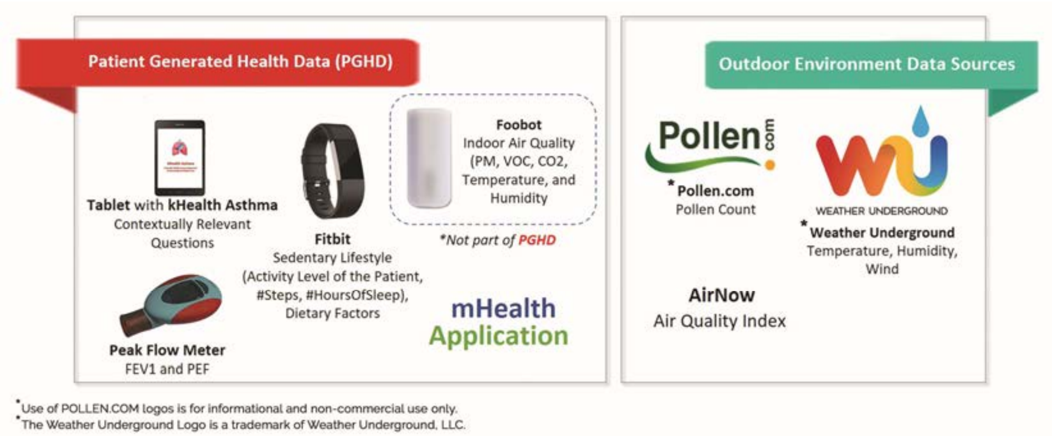

Figure 1. kHealth, an mHealth application that gathers patient-generated health data (PGHD) through contextually relevant questions (tablet), sensors (Foobot, peak flow meter), wearable (Fitbit), and from external data sources for contextual (such as location-specific) environmental data.

We describe the kHealth APH approach with an example. Sara is a 10-year-old girl. With the help of our kHealth kit (Figure 1), she is able to monitor her daily activities, helping her and her clinician to intervene and update her treatment plan accordingly. Self-monitoring refers to the data collection using the mobile devices and sensors. Sara completes her mHealth application questionnaire twice a day, collects her daily activity level and sleep pattern using Fitbit[2] (http://bit.ly/1VdkW3I) and places an indoor air quality monitor to measure her indoor environment at home. Self-monitoring in many cases, won't be helpful if collected data aren't acted upon. In general, we don't want a technology to make a clinical decision and change the care plan, but to enable adherence to a care plan specified by the patient's physician. Self-appraisal refers to the process of self-monitoring and self-reconciliation of the observed data. Self-management helps a patient to make better judgment or action within the scope of the care plan. For instance, a patient might observe that every time he goes out and has allergic symptoms, the APH systems has found a strong correlation with high pollen. The intervention involves looping in the clinician for monitoring the severity level of patient and alter the care plan. Activities involved in the intervention may include the addition of new medication, changing the course of intake of medication or suggesting preventive medication considering the historical observations of the patient by involving the clinical and support services in care and health monitoring process. Assessment of post-intervention processes is crucial for reclassification of patient's disease. For instance, the collected evidences can help the physician to reclassify the asthma from mild persistent to moderate persistent and adjust the care plan with modification in the medication.

CHALLENGES IN CONVERTING BIG DATA INTO SMART DATA

A variety of studies involving PGHD and other health-relevant data using a broad variety of IoT are ongoing for developing personalized digital care solutions for a variety of health related objectives, as shown in Figure 2. These systems need to deal with a host of data related challenges such as accessing, storing, querying, and managing large volumes of highly dynamic data, and systems related challenges such as interoperability and integration, security, privacy, trust, scalability, and reliability.[3] We characterized the ongoing efforts along two critical dimensions: abstractions (making the data meaningful and interpretable with respect to an individual's health) and actionable information (supporting decision making and actions informed from the data). These involve addressing challenges in data analysis including semantic data modeling, annotation, knowledge representation (for example, modeling for constrained environments, complexity issues, and time/location dependency of data), and so on. While statistical analysis of the data collected helps one identify correlations, it's widely observed that a correlation doesn't necessarily imply a causation. Working with domain experts (that is, clinicians for the health applications) for understanding correlations between observations from different modalities is the key in associating meaning to the variations in observations that can then support derivation of causations. Another challenge is that the clinicians, health practitioners, and patients cannot keep up with an enormous amount of data being generated. Patients can't interpret the data in the context of health conditions and objectives and clinicians don't have time to look at it. There's an urgent need to convert the raw data into Smart Data. By making sense out of big data (http://j.mp/SmData), Smart Data provides value from harnessing the challenges posed by volume, velocity, variety, and veracity of big data, in turn providing actionable information and improving the decision-making process. Smart Data is focused on the actionable value achieved by human involvement in data creation, processing, and consumption phases for improving the human experience (http://bit.ly/HumanExperience). We propose the following evidence-based semantic perception approaches: (a) contextualization, (b) abstraction, and (c) personalization.

Contextualization

Contextualization refers to data interpretation in terms of knowledge (context). PGHD consists of demographic and medical information from EMRs and time-series data collected from various environmental sensors, physiological sensors, and public web resources. Contextualization supports ranking a patient's diagnosis and patient similarity based on demographics and PGHD. It deals with low-level fine-grain data covering various facets by determining the data type and value, and then situates it in relation to other domain concepts, thus developing a meaningful interpretation of results. A large body of existing research on ontologies and semantic web techniques and technologies can be leveraged for this purpose.[4] However, relying solely on description logics or formal knowledge representation alone is often not sufficient to understand the complex nature of many health conditions. Probabilistic graph models from representing knowledge graphs, combined with machine learning and NLP on relevant data is an alternative some recent approaches have used.

Abstraction

Abstraction is a computational technique that maps and associates raw data to action-related information, taking into account personal details but ignoring inessential differences to provide an integrated view of proper remediation measures. For example, high activity translates to different workout durations based on age, weight, current health, weather, and sport; or a low risk of heart problems depends on demographic and genetic information, as well as diet. In some cases, abstraction can be embedded on the device: for example, question-answering systems in mobile health applications as a way of indirectly supervising personalization of healthcare. However, one of the challenges is the need to formalize normalcy and detect an anomaly. Anomaly detection is nontrivial because the notion of normalcy itself is intrinsically dynamic, based on spatiotemporal and personal context. It also requires personalization and the ability to uncover various correlations among multimodal data streams and discovering medically relevant abstract interpretations and the factors that influence them. The challenge itself can be overcome if sufficient patient data can be obtained through large-scale clinical studies, followed by identification of correlations, and then analyzed and explained by those with domain knowledge and expertise to derive causations.

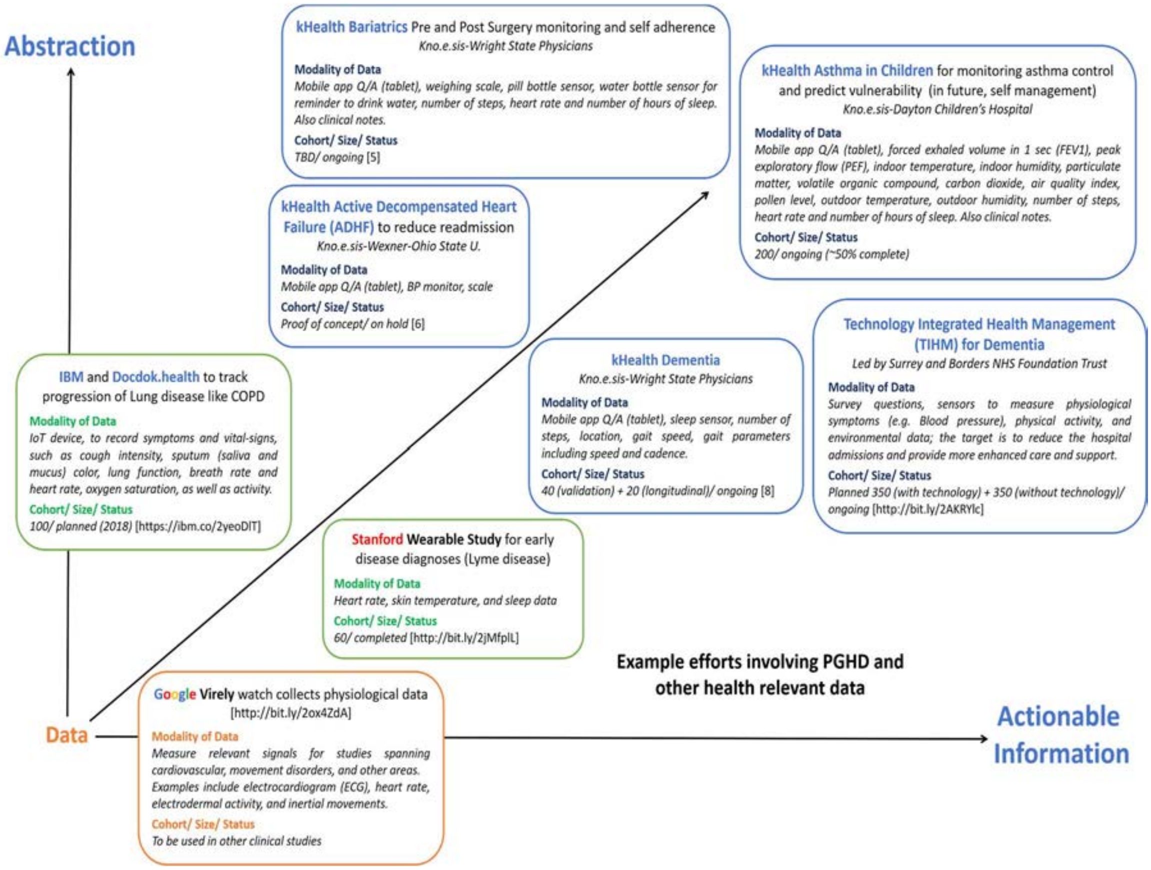

Figure 2. Example efforts involving PGHD and other health-relevant data. Converting PGHD data into actionable information. [References for kHealth Bariatrics,[5] kHealth ADHF,[6] kHealth Dementia,[8] TIHM[7] provide further details.]

Personalization

Personalization in healthcare refers to the determination of a treatment plan based on severity of disease, the prevalence of triggers, and vulnerabilities vis-à-vis the use of past and current health data. For example, a low-dose SABA (Short-Acting Beta Agonists) might help someone keep asthma symptoms in check during the fall season, but it might not work for another patient who needs a higher dosage due to more severe asthma and a greater prevalence and intensity of triggers during the spring season. IoT data provides an opportunity for personalization of future course of action and treatment plans by taking into account the contextual factors such as patient's health history, physical characteristics, environmental factors, activity, and lifestyle.

With contextualization, abstraction, and personalization in place, the next problem is how to synthesize a personalized vulnerability score for a patient's given medical condition or disease with respect to relevant health management objectives to better establish a control level, and to quantify and express the effectiveness of remedial measures in a manner readily accessible to both patient and clinician.

CONCLUSION

In terms of IoT and health, most current efforts are focused on data collection and improving understanding what the data implies at a basic level. Collected data needs to be analyzed and validated with EMRs that capture patient-care objectives, plans, and information self-reported by patients. The key aspect is generating actionable information that will be acceptable, easy to use and integrate into clinical pathways used by other systems, and can be given at the right time with the right modality to the end beneficiaries (clinicians and patients) with appropriate information governance procedures and privacy/security measures. The privacy issues in healthcare data can be dealt with by homomorphic encryption schemes, differential privacy, and data perturbation. Depending on the privacy level needed, additional cryptographic services can be introduced into the framework.

Transitioning from cohort-based treatment to more personalized treatment, basic statistical computing with causality and machine learning algorithms won't suffice. There's a need to combine and integrate machine learning and data analytics with reasoning engines and knowledge bases, thus propelling us into the realm of augmented personalized healthcare management and well-being applications and services.

ACKNOWLEDGMENTS

We thank Payam Barnaghi for his review and helpful suggestions. This work was partially supported by the National Institutes of Health (NIH) under grant R01 HD087132-01. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

REFERENCES

1. A. Sheth et al., "Augmented Personalized Health: How Smart Data with IoTs and AI Is about to Change Healthcare," *Proc. IEEE 3rd Int'l Forum Research and Technologies for Society and Industry* (RTSI), 2017, pp. 1–6.
2. J. Bian et al., "Exploring the Association Between Self-Reported Asthma Impact and Fitbit-Derived Sleep Quality and Physical Activity Measures in Adolescents," *JMIR mHealth and uHealth*, vol. 5, no. 7, 2017.
3. S. Sharma, K. Chen, and A. Sheth, "Towards Practical Privacy-Preserving Analytics for IoT and Cloud Based Healthcare Systems," *To appear in IEEE Internet Computing*, 2018.
4. A. Sheth and K. Thirunarayan, "Semantics Empowered Web 3.0: Managing Enterprise, Social, Sensor, and Cloud-based Data and Services for advanced applications," *Synthesis Lectures on Data Management*, vol. 46, no. 6, 2012, pp. 1–175.
5. S. Sohail et al., "A Multisensory Approach to Monitor Bariatric Patients' Postsurgical Behavior and Lessen Weight Recidivism," Obesity Week 2017. 34th ASMBS Annual Meeting, 29 Oct - 2 Nov 2017.
6. A. Sheth, P. Anantharam, and K. Thirunarayan, "kHealth: Proactive Personalized Actionable Information for Better Healthcare," *40th Int'l. Conf. Very Large Databases Workshop Personal Data Analytics in the Internet of Things* (PDA@IOT), 2014.
7. S. Enshaeifar et al., "Internet of Things for Dementia Care," *IEEE Internet Computing*, vol. 22, no. 1, 2018, pp. 8–17.
8. T. Banerjee et al., "Validating a Commercial Device for Continuous Activity Measurement in the Older Adult Population for Dementia Management," *Smart Health*, 2017; doi.org/https://doi.org/10.1016/j.smhl.2017.11.001.

ABOUT THE AUTHORS

Amit Sheth is the LexisNexis Ohio Eminent Scholar and the executive director of Kno.e.sis, the Ohio Center of Excellence in Knowledge-enabled Computing and BioHealth Innovations at Wright State University. His research interests include physical-cyber-social big data and computing for human experience, supported by knowledge-empowered artificial intelligence (AI) techniques. Sheth received his PhD from Ohio State University. He is an IEEE and an AAAI Fellow. Contact him at http://knoesis.org/amit.

Utkarshani Jaimini is a PhD candidate in computer science at Kno.e.sis Center at Wright State University. Her research interests include applied machine learning to solve sustainability and healthcare problems, predictive analytics, information extraction, augmented personalized healthcare, and the semantic web. Jaimini received her bachelor's degree in computer science engineering from The LNMIIT University. Contact her at utkarshani@knoesis.org.

Hong Yung (Joey) Yip is a PhD student at Kno.e.sis Center at Wright State University. His research interests include conversational artificial intelligence and semantic web and their applications to augmented personalized health. Yip received his MS in bioinformatics from University of Malaya. Contact him at joey@knoesis.org.

This article originally appeared in
IEEE Intelligent Systems, *vol. 33, no. 1, 2018.*

Health Media: From Multimedia Signals to Personal Health Insights

Susanne Boll
University of Oldenburg

Jochen Meyer
OFFIS Institute for Information Technology

Noel E. O'Connor
Dublin City University

Ever since the emergence of digitization, we've used the term *multimedia* to represent a combination of different kinds of media, such as text, images, audio, and videos. As new sensing technologies emerge and become omnipresent in our daily lives, the definition, role, and significance of multimedia is changing. Multimedia now represents the means for communicating, cooperating, and monitoring numerous aspects of daily life at various levels of granularity across multiple applications, ranging from personal to societal. Multimedia has become an integral part of the tools and systems that are providing solutions to today's societal challenges—including challenges related to health and healthcare.

In this article, we explore the relevance and contribution of new signals in this broader interpretation of multimedia for personal health. We present how core multimedia research is becoming an important enabler for applications with the potential for significant societal impact. We review the new multimedia signals in health, which we term *health media*, the research challenges they pose, and the trends in multimedia they are driving. We illustrate this with recent applications and systems from our own and related work. The new signals stem from a multitude of new smart health devices that are increasingly available to everyone and commonly accepted in our daily lives. However, the signals are only the first step in a journey towards a healthier lifestyle. Sensing, or even sensing followed by data analysis, is only part of the story; only by closing the "sense-analyze-decide-actuate" feedback loop will we encourage healthier lifestyles that will improve users' personal health.

We present the pathway from multimedia signals to a new generation of future personal digital health technology. First, we briefly review the new technological developments that enable continuous tracking and tracing of personal health information. These new signals relate to our personal health but do not yet deliver information or even recommendations to users. Next, we show how data from low-cost everyday devices can contribute to relevant health-related information. Finally, we look at the role of the now-empowered user facing the opportunity and challenge of exploiting and acting on health media. Only when the "sense-analyze-decide-actuate" loop is complete can the new multimedia signals objectively contribute to personal health. We discuss the roles of usage, usability, visualization, and health interventions in the success of this new area of digital health.

 2469-7087/19/$33.00

SENSING: THE NEW MULTIMEDIA SIGNALS IN HEALTH

Smart Health Devices

Over the last decade, a great variety of wearable and ambient sensors have been developed and are constantly being improved and incorporated in people's daily lives. These sensors measure physiological parameters in a non-intrusive, accurate manner and produce new types of health media. Key drivers are the ongoing miniaturization of sensors and actuators and advances in battery technology. The result is that we can now routinely sense a plethora of health-related data such as steps taken, miles cycled or run, heart rate, sleep quality, and muscular effort.[1] As an illustration, Figure 1 shows a snapshot of the smart health devices owned by a group of personal-health researchers during a 2015 Dagstuhl seminar on Health Behavior Change.[2]

Figure 1. A multitude of everyday digital health devices.

Devices and measurement that were once used only in medical-related contexts are now appearing in the consumer market. Early examples of these devices were sports watches targeted at the amateur athlete for tracking heart rate, pace, and speed during workouts, which enabled analysis of individual performance and performance improvements over time. Consequently, the physical appearance of the device was targeted at individuals already active in sports and interested in tracking and exercising. In a next wave of wearable devices, we observed activity trackers that use accelerometer data to quantify physical activity, mainly by counting steps. With the emergence of activity trackers for self-monitoring, such as the very prominent Fitbit, the market now addresses not just the sports and data enthusiast, but also the average person possessing even a low level of technology affinity with a moderate interest in health. The devices became smaller, much more attractive in form and shape, and, in recent years, even a fashionable item[3] fusing aesthetics with smart health functionality. In our own work, we showed that fashionable appearance contributes to the device's regular use, which potentially affects the data coming from the device.4 We discuss this in more detail in the section on the role of use and usability of smart health devices in the "sense-analyze-decide-actuate" loop.

Accelerometer-based step tracking was only the beginning of digital health devices in a large mass market. Now, smart devices are available for numerous relevant health parameters beyond walking and running, such as resistance training, sleep monitoring, mindfulness practice, posture monitoring, weight management, body composition, breathing techniques, and cardiac health status. Millions of these devices are now available in the consumer market in many forms and shapes. We can conclude that smart consumer devices collecting new multimedia signals related to our personal health are now (potentially) available to everyone.[1] Furthermore, this is just the beginning of the emergence of health media. For example, we are already starting to see the emergence of biomedical sensors using organic and printed electronics integration pointing to futuristic tattoo-like sensors for vital signals.[5,6]

Health-Related Multimedia Signals

So what are the signals that constitute this new paradigm of health media? The signals are of a wide variety of data types, coming in specific data formats and value ranges, sample sizes, and temporal resolutions. Heart rates might be monitored and sampled continuously throughout the day, but weight might be measured only once a day. Meanwhile, nutrition might be monitored in a self-assessment notebook several times a day. Beyond the new sensors, the digital traces we leave behind in our social networks, instant messengers, calendars, and smartphones can also be highly valuable for inferring information about our health.[7] For example, lack of interaction with devices is a necessary prerequisite for sleeping,[8] our location and its change over time indicates potential physical activity, our posts on social media might reveal our eating behavior[9] or contain useful information about our current mental state, and deviations from normal behaviors might indicate a change in life that is worth considering.

Very often, there is no one single device for collecting certain health parameters. A good example of this is sleep monitoring. Early sleep monitors were dedicated body-worn devices that used accelerometers to identify the sleeper's movements and infer phases of sleep and wake, something that is supported by most activity trackers today. Sleep monitoring now is more frequently done with sensor strips or mats placed under the sheet or mattress. These sensors identify the sleeper's movement and heart rate and allow for a reasonable estimation of sleep duration and quality. Sleep monitoring also reveals a very important aspect of measuring health-related parameters with low-cost consumer devices: Some health behaviors are complex issues that are difficult to assess and cannot easily be grasped using simple one-dimensional measurements. However, data cleaning, fusion, and merging can act as the bridge from multimedia signals to health insights.

Health-Logging

A huge advantage of this kind of low-cost sensing is that it enables the collection of data over long periods of time, covering not just weeks or months, but years or decades. These sensors can thus provide unprecedented insights into a person's daily life and its changes over time. Extensive collection of personal data for some unknown future use is a characteristic of *lifelogging*. As a pioneer in the field, Gordon Bell coined the term in Microsoft's MyLifeBits project,[10] which was inspired by Vannevar Bush's 1945 vision of a "Memex" machine for storing and retrieving all personal information and memories. This approach aimed mainly at collecting multimedia data, such as photos, for documenting a person's daily life. The related term *Quantified Self* was coined by journalists Gary Wolf and Kevin Kelly in 2007.[11] With more and more sensing devices out there, the Quantified Self movement became popular among data and technology enthusiasts to track various aspects of their lives using technology, pen-and-paper, and other tools.

The motto "self-knowledge through numbers" points out two important aspects of Quantified Self. First, it is about just one person observing oneself. The ambition is not to provide generic findings by observing a large group of people with relatively few parameters, but rather to gain insights about just one person by deeply analyzing that person's own data traces. This resembles traditional "case studies" in medical literature, although here a person is examining herself or himself rather than another person. Professor Deborah Estrin described this as the "n=me" approach[12] and argued that this is a reasonable way of gaining knowledge, complementing but not replacing the traditional evidence-based approach with large user groups.

Health and wellbeing emerged early as one important field within this movement, with applications including food diaries, activity recognition, and supporting memory in dementia patients.[13] We have reached the point where *health-logging*—the long-term logging of health-related features—is now possible and goes beyond previous experiments that often only lasted three to six months.[14] These new continuous signals from an individual offer tremendous opportunities for future personal healthcare. Understanding such multimedia data in health contexts is one of the core challenges that the multimedia community needs to investigate further.

ANALYSIS: FROM MULTIMEDIA SIGNALS TO HEALTH INSIGHTS

On the way from multimedia signals to health insights, several challenges have to be solved. The new multimedia signals are interesting in a medical context. Yet there is a long way to go before we can effectively clean, fuse, and interpret the (small) data collections and finally produce health-relevant information.

The Discrepancy between Signals and Health Relevance

One challenge is that the quality of the data produced by low-cost consumer devices is not as high as that of gold-standard clinical assessments. With consumer devices, we get a completely new and different perspective. The data that is collected can be very different from that of an approved medical device or assessment. An activity tracker primarily delivers step counts per minute, unlike a research-grade accelerometer that can deliver high-resolution 3D accelerometry data; it is neither meant nor able to recognize fine-grained activities. There isn't even a clear distinction between an actual step and a mere leg movement; hence, two activity trackers can reasonably deliver different step counts for the same daily activity. A consumer sleep tracker measuring only movement and heart rate using unobtrusive sensors can merely provide a small subset of the data that a high-end polysomnography can produce in a sleep lab with its huge but expensive and highly intrusive sensors. However, most people will never assess their sleep in a sleep lab, whereas everybody can monitor their sleep using a consumer device over many years.

The broadness of data [from low-cost consumer devices] might compensate for the lack of complexity and mitigate the potentially lower or unknown precision.

Smart health device data will therefore be of a different class than that of medical devices. It will often be less precise, not just for technical reasons, but also because of other reasons such as faulty use. However, it is much more available; in the doctor's office, blood pressure is measured every few months or years, whereas at home, the measurement can easily be taken daily. Furthermore, for many applications, a large set of low-precision data is often better than an empty (or very small) set of high-precision data.[1]

While the data quality might be limited, the new signals enable the availability of longitudinal data and thus provide unprecedented insights into a person's daily life. This broadness of data might compensate for the lack of complexity and mitigate the potentially lower or unknown precision. It also opens the door for new analyses and new applications in which identifying recurring patterns, normal states, deviations, and changes, as well as providing informal support in non-standard situations, might be more important than the highest precision in measuring a certain vital parameter at a specific point in time.

In the end, we find that there is no one sensor or device collecting all health-relevant data, but rather a very heterogeneous and large world of many data collections that are related to our health. None of these signals alone necessarily provides an individual with a picture of his or her health. A singular look at a daily step count of 10,000 might not be helpful for estimating how heart-friendly a person's lifestyle is. Research is now addressing systems for more complex health questions that require monitoring multiple behaviors and need more complex data analysis to identify health states and outcomes.

From Data to Health-Relevant Insights

Health is highly complex and multifaceted, and it cannot be understood by addressing just one aspect such as physical activity. Personal health applications therefore require gathering diverse data covering multiple aspects of health. As it seems unlikely that there will soon be one device that can measure everything and that everyone will use that one device, the integration of health-related data from multiple sources and providers is a key requirement for future health systems. While devices such as activity trackers were initially criticized for their low precision compared to high-end accelerometers,[15] they are now understood as a new class of tools with distinct characteristics and acceptable precision[16] that is useful in epidemiological studies[17] and for understanding patient exercise adherence in rehabilitation scenarios.[18] Ensuring a high expressiveness of data in relation to health and health behavior is therefore necessary.

Fusing health-related data is not just a practical challenge of accessing multiple systems with different APIs, access rights, and data formats. It is also a challenge to make sense of highly heterogeneous data and make it usable for new health-related applications. The key here is to understand the raw data, and then analyze and enrich it to turn it into meaningful health-related information.

An approach that we have successfully used in various projects foresees multiple steps of data processing.[19] Beyond the more obvious ones such as abstraction from physical to logical devices and aggregation of data over time, we particularly suggest one layer that maps data from one or multiple (logical) devices into what we call "primary health features" (see Figure 2). Primary health features are the smallest meaningful unit of a user's health data and the atomic building blocks for a user's health goals. Mapping from logical devices to primary health features converts a device's measurements into observations about a user's health, such as physical activity pursued in bouts of at least 10 minutes, one instance of sleep with onset and duration, and the weight and body fat at a given moment of time, thus raising the data to a higher semantic level.

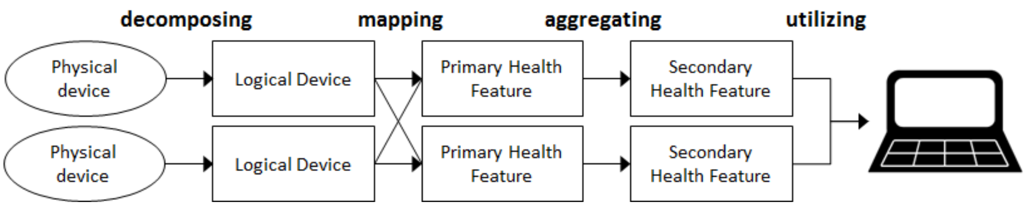

Figure 2. From multimedia signals to health insights.

The mapping from logical devices to primary health features can be canonical, such as when a scale directly delivers the values for weight and body fat as the primary health feature. However, the mapping also opens the door for advanced analyses that fuse data from multiple devices, integrate context and personalized knowledge, and deploy complex processing to, for example, identify the daily instance of "20 minutes cycling to work" based on time of day, GPS track, step counts, and past observations. Existing work, including our own, points to the rapid advances and impressive results of such techniques.[20,21]

Numerous other analyses can be envisioned here that, based on our digital multimedia data collections, infer information that is relevant for our health. This is where one of the core research topics of the multimedia community comes into play: making sense of rich but semi- or unstructured data. Our lifelogging photos can help us identify our eating behavior,[22] our tweets and Facebook posts can reveal our mental state,[23] and our interaction with digital devices can reveal our sleep circadian rhythm.[24]

DECISION AND ACTUATION: FROM HEALTH OPPORTUNITIES TO HEALTH ACTION

One of the hidden marketing promises of these consumer devices and apps is that merely tracking your health will make you healthier. However, many applications, particularly many medically oriented ones, primarily aim at measuring and presenting data, addressing the "sense" and "analyze" steps in the aforementioned feedback loop. The focus of applications for personal health data collections should rather be on the "decide" and "actuate" steps, providing actionable information and supporting the user in identifying changes, making decisions, and understanding his or her health.

Abandonment and Long-Term Use

The user is not just the ultimate consumer of the personal health information, but also the creator of the multimedia signals and the raw data. This dual role has a severe impact on any system making use of the data. Abandonment and lack of compliance are the most pressing challenges in short-term settings. In a study by Lazar et al.,[25] 80 percent of devices were abandoned after two months. And, in a study by Shih et al.,[26] 75 percent of 26 undergraduate students with activity trackers stopped using the tracker within four weeks. In real life, a market survey from 2014 found that one-third of all activity trackers are abandoned within six months.[27]

While there is no clear definition of when short-term use ends and long-term use starts, the current agreement seems to be that durations well beyond six months can indeed be considered long term.[28] We observed subtle but noticeable changes in self-tracking in this long-term setting,[14] shifting from behavior change and learning towards logging and confirmation. However, we also found that virtually everybody takes breaks from self-tracking that can last anywhere from days to years.[29] Luckily, users follow certain patterns in taking breaks, allowing researchers to predict the breaks (to some extent) and deal with the resulting gaps in the data.

Usability and Aesthetics of Smart Health Devices

Lack of usability is one of the reasons for abandonment of use. Devices that are supposed to be used every day and over a very long time clearly have particular usability requirements.[30] But not all devices can be worn all the time. For example, a user might experience a situation where regulations prohibit wearing wristbands for hygienic reasons. A lack of an appropriate wearing position on the body might prevent a user from wearing an ambient device.

However, for devices that are meant to be worn the whole day, aesthetics and physical and social comfort are particularly important aspects of usability. When investigating the role of aesthetics in the success of a wearable health device, we designed and built WaterJewel (see Figure 3), a bracelet with discreetly integrated light spots that reflect the user's actual fluid-intake behavior through abstract light signals. In a participatory design process, we created two decorative designs (masculine and feminine). In a field experiment, we explored the use of the WaterJewel prototypes in daily life and compared WaterJewel to a prevalent mobile fluid-intake reminder application. We found that users thought WaterJewel was a decorative, discreet, and practical piece of smart digital jewelry and that they expected a high level of customizability. More important, participants drank more in total and more regularly using the WaterJewel bracelet compared to a normal smartphone application.

Figure 3. WaterJewel: a bracelet to promote better drinking behavior.

We also found that usability is a matter of the data that the device collects. Consolvo et al. identified accuracy of data, adequacy of data coverage, appropriate units of measurement, and the need to deal with manual data entry in addition to automatic sensing as key points for collecting nutrition and activity data for health and wellness.[31] Understanding the data from a system's point of view, as well as making sense of it and enabling the user to understand it, are therefore necessary for successful use and application of consumer health devices.

Visualization

Multimedia in personal health and personal healthcare comes with an enormous demand for visualization and interpreting multimodal health data. Visualizations of patient data are used for not only diagnosis and therapy, but also management of personal health and chronic disease, because they can provide a basis for reflection and decision-making. Visualizations can enable coaching and cooperation between different stakeholders such as patients, medical doctors, and caretakers.

UbiFit Garden[32] was one of the earliest visualizations of personal health data. It uses a glance-able display on a mobile phone to represent one dimension of personal health, namely physical activity. In our own work, we addressed visualization of heterogeneous personal health data on a mobile phone in a neutral, non-persuasive way. We found the need for two types of visualization: "my health, now" and "my health, in the past." Visualizing "my health, now" aims to give objective feedback about current health status for decisions about actions in the near future or comparison of the perceived health status with the objective measurement. Considering Shneiderman's visualization mantra, "overview first […] then details-on-demand," we further found that an overview should cover the last 72 hours and that the best timeframe for details-on-demand seems to be the last seven days. Data older than one week does not influence "my health, now" anymore and should rather be seen as "my health, in the past" with a different visualization. We developed multiple visualizations: glance-able displays for the overview such as one with a tree metaphor, a simple timeline style for the details-on-demand, and a more traditional graphs-and-tables presentation for past data (see Figure 4).

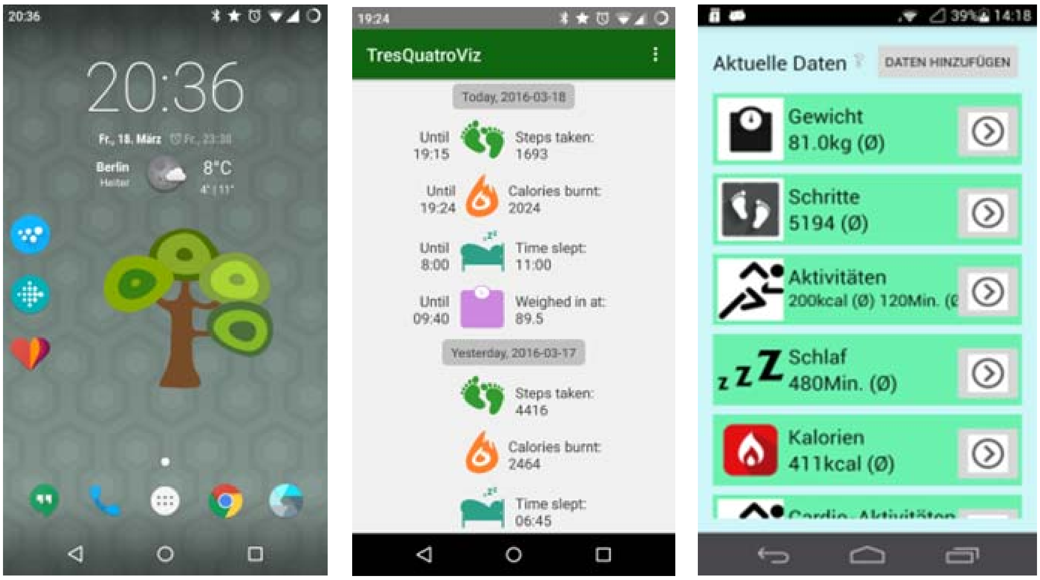

Figure 4. Visualizations of heterogeneous personal health data.

CHALLENGES

This new research field offers exciting opportunities but comes with challenges, as well. Beyond the research questions themselves, these new signals demand new approaches for research, especially because personal health is a long-term, possibly life-long, effort. Beyond the classical challenges of signal processing—sensor data fusion and filtering out noise—we also see numerous challenges related to the relevance and applicability of health media when it comes to making a real impact on our personal health.

Long-term studies: Conducting studies is a key tool for research; however, there are huge differences in understanding the setup of studies in computer science and medicine. Long-term studies covering years and decades would, in principle, be necessary to rigorously evaluate long-term tracking applications. However, this is currently implausible, not just because the effort is far too much for most computer science projects, but also because it is impossible to design such studies to account for the fast-changing technological world.

Open data collection: Collecting long-term tracking data is a big effort. Joining forces to collect and share the valuable data among different researchers for multipurpose analyses would be an obvious approach. However, sharing such data is hampered by issues such as lack of data interchange formats, difficulties in ensuring anonymity (such as with location data), and restrictive data protection requirements. More needs to be done to show the potential of health data to the public and institutions involved in collecting data to facilitate sharing of such data for research purposes.

Sharing [data] is hampered by issues such as lack of data interchange formats, difficulties in ensuring anonymity (such as with location data), and restrictive data protection requirements. More needs to be done to show the potential of health data to the public and institutions involved in collecting data to facilitate sharing of such data for research purposes.

Domain expertise: Technology development needs to proceed hand in hand with the relevant domain expertise, whether that expertise is related to clinical, human-performance, or psychological factors, to name but a few. This requires an ongoing dialogue between these fields with multidisciplinary research teams working together on specific projects. This is timely, given that these fields are becoming increasingly receptive to new technologies and the benefits they can bring. Technology research and development should be informed by a deep understanding of the relevant domain, leveraging state-of-the-art systematic reviews of what has proven effective (and ineffective) in the past.[33]

Usability and privacy: If it is not used, it is useless. The human factor plays an important role in this new area of multimedia signals for health. We can see that the planned use by a manufacturer or provider might not be the actual use by the individual. The tracking devices, for example, are not necessarily used regularly. Making the data valuable and relevant will play an important role in fostering interest and compliance in using the system, which will lay the foundation for personal health insights. Gaining health insights with smart devices also comes with the side effect that someone (else) is potentially getting access to our data. To encourage and motivate research in this area, greater transparency and understanding is needed regarding which stakeholders have access to our data and personal health insights and what that means for our privacy rights.

REFERENCES

1. J. Meyer and S. Boll, “Digital Health Devices for Everyone!,” *IEEE Pervasive Computing*, vol. 13, no. 2, 2014, pp. 10–13; www.computer.org/csdl/mags/pc/2014/02/mpc2014020010-abs.html.
2. “Life-long health behavior change technologies,” *Dagstuhl Seminar 15262*, 2015; www.dagstuhl.de/de/programm/kalender/semhp/?semnr=15262.

3. M. Maslakovic, "Merging fashion with fitness, ten activity trackers for women," *Gadgets & Wearables*, Gadgets & Wearables, 2017; www.gadgetsandwearables.com/2017/10/10/best-activity-trackers-for-women/.
4. J. Fortmann et al., "WaterJewel: design and evaluation of a bracelet to promote a better drinking behaviour," *Proceedings of the 13th International Conference on Mobile and Ubiquitous Multimedia* (MUM '14), 2014, pp. 58–67; https://dl.acm.org/citation.cfm?id=2677976.
5. S.K. Ameri et al., "Graphene Electronic Tattoo Sensors," *ACS Nano*, vol. 11, no. 8, 2017, pp. 7634–7641; https://pubs.acs.org/doi/abs/10.1021/acsnano.7b02182.
6. T. Ahmadraji et al., "Biomedical Diagnostics Enabled by Integrated Organic and Printed Electronics," *Analytical Chemistry*, vol. 89, no. 14, 2017, pp. 7447–7454; https://pubs.acs.org/doi/abs/10.1021/acs.analchem.7b01012.
7. M. De Choudhury, "Opportunities of social media in health and well-being," *XRDS: Crossroads, The ACM Magazine for Students - Health Informatics*, vol. 21, no. 2, 2014, pp. 23–27; https://dl.acm.org/citation.cfm?doid=2701297.2676570.
8. Z. Chen et al., "Unobtrusive sleep monitoring using smartphones," *7th International Conference on Pervasive Computing Technologies for Healthcare* (PervasiveHealth), 2013; https://ieeexplore.ieee.org/document/6563918/.
9. C.F. Chung, "When Personal Tracking Becomes Social: Examining the Use of Instagram for Healthy Eating," *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (CHI '17), 2017, pp. 1674–1687; https://dl.acm.org/citation.cfm?id=3025747.
10. J. Gemmell et al., "MyLifeBits: fulfilling the Memex vision," *Proceedings of the tenth ACM international conference on Multimedia* (MULTIMEDIA '02), 2002, pp. 235–238; https://dl.acm.org/citation.cfm?doid=641007.641053.
11. G. Wolf, "What is The Quantified Self?," *Quantified Self blog*, Quantified Self, 2011; www.quantifiedself.com/2011/03/what-is-the-quantified-self/.
12. D. Estrin, "Small data, where n = me," *Communications of the ACM*, vol. 57, no. 4, 2014, pp. 32–34; https://dl.acm.org/citation.cfm?doid=2580723.2580944.
13. Y. Yang et al., "Design and field evaluation of REMPAD: A recommender system supporting group reminiscence therapy," *Evolving Ambient Intelligence*, Springer, 2013; https://link.springer.com/chapter/10.1007/978-3-319-04406-4_3.
14. J. Meyer et al., "Making Sense in the Long Run: Long-Term Health Monitoring in Real Lives," *2017 IEEE International Conference on Healthcare Informatics* (ICHI), 2017; https://ieeexplore.ieee.org/document/8031158/.
15. P. Schneider et al., "Accuracy and Reliability of 10 Pedometers for Measuring Steps over a 400-m Walk," *Medicine & Science in Sports & Exercise*, vol. 35, no. 10, 2003, pp. 1779–1784.
16. T. Ferguson et al., "The validity of consumer-level, activity monitors in healthy adults worn in free-living conditions: a cross-sectional study," *International Journal of Behavioral Nutrition and Physical Activity*, 2015; https://ijbnpa.biomedcentral.com/articles/10.1186/s12966-015-0201-9.
17. B.A. Brühmann, M.E. Schmidt, and K. Steindorf, "Assessment of physical activity in epidemiological studies: Are questionnaires obsolete in the era of accelerometry?," *GMS Medical Informatics, Biometry and Epidemiology*, vol. 10, no. 1, 2014; https://egms.de/static/en/journals/mibe/2014-10/mibe000155.shtml.
18. O. Duff et al., "MedFit: The development of a mobile-application to enhance participant self-management of their cardiovascular disease," *Health Enhancing Physical Activity Europe Conference*, 2016; https://doras.dcu.ie/21273/.
19. J. Meyer and S. Boll, "Smart health systems for personal health action plans," *IEEE 16th International Conference on e-Health Networking, Applications and Services* (Healthcom), 2014; https://ieeexplore.ieee.org/document/7001877/.
20. A. Ahmadi et al., "Human gait monitoring using body-worn inertial sensors and kinematic modelling," *IEEE SENSORS*, 2015; https://ieeexplore.ieee.org/document/7370310/.
21. A. Ahmadi et al., "Automatic detection, extraction and analysis of unrestrained gait using a wearable sensor system," *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (EMBC), 2015; https://ieeexplore.ieee.org/document/7318786/.
22. G. O'Loughlin et al., "Using a Wearable Camera to Increase the Accuracy of Dietary Analysis," *American Journal of Preventive Medicine*, vol. 44, no. 3, 2013, pp. 297–301; https://ajpmonline.org/article/S0749-3797(12)00863-X/fulltext.

23. M. De Choudhury et al., “Characterizing and predicting postpartum depression from shared facebook data,” *17th ACM conference on computer supported cooperative work & social computing* (CSCW '14), 2014, pp. 626–638; https://dl.acm.org/citation.cfm?doid=2531602.2531675.
24. E.L. Murnane et al., “Mobile manifestations of alertness: connecting biological rhythms with patterns of smartphone app use,” *18th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '16), 2016, pp. 465–477; https://dl.acm.org/citation.cfm?id=2935383.
25. A. Lazar et al., “Why we use and abandon smart devices,” *ACM International Joint Conference on Pervasive and Ubiquitous Computing* (UbiComp '15), 2015, pp. 635–646; https://dl.acm.org/citation.cfm?doid=2750858.2804288.
26. P.C. Shih et al., “Use and Adoption Challenges of Wearable Activity Trackers,” *iConference*, 2015; https://hdl.handle.net/2142/73649.
27. “Inside Wearables: How the science of human behavior change offers the secret to long-term engagement,” *Endeavour Partners blog*, 2014; https://blog.endeavour.partners/inside-wearable-how-the-science-of-human-behavior-change-offers-the-secret-to-long-term-engagement-a15b3c7d4cf3.
28. E. Karapanos et al., “Wellbeing in the Making: Peoples’ Experiences with Wearable Activity Trackers,” *Psychology of Well-Being*, vol. 6, no. 4, 2016; https://ncbi.nlm.nih.gov/pmc/articles/PMC4908170/.
29. J. Meyer et al., “Identification and Classification of Usage Patterns in Long-Term Activity Tracking,” *Conference on Human Factors in Computing Systems* (CHI '17), 2017, pp. 667–678; https://dl.acm.org/citation.cfm?id=3025690.
30. J. Meyer et al., “Making Lifelogging Usable: Design Guidelines for Activity Trackers,” *International Conference on Multimedia Modeling*, 2015, pp. 323–334; https://link.springer.com/chapter/10.1007%2F978-3-319-14442-9_39.
31. S. Consolvo et al., “Designing for Healthy Lifestyles: Design Considerations for Mobile Technologies to Encourage Consumer Health and Wellness,” *Foundations and Trends in Human-Computer Interaction*, 2014; https://dl.acm.org/citation.cfm?id=2693094.
32. S. Consolvo et al., “Activity sensing in the wild: a field trial of ubifit garden,” *SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), 2008, pp. 1797–1806; https://dl.acm.org/citation.cfm?doid=1357054.1357335.
33. O. Duff et al., “Behavior Change Techniques in Physical Activity eHealth Interventions for People With Cardiovascular Disease: Systematic Review,” *Journal of Medical Internet Research*, vol. 19, no. 8, 2017; https://jmir.org/2017/8/e281/.

ABOUT THE AUTHORS

Susanne Boll is a professor of media informatics and multimedia systems at the University of Oldenburg, Germany. Contact her at susanne.boll@uni-oldenburg.de.

Jochen Meyer is director of the R&D Division Health at OFFIS Institute for Information Technology, Oldenburg, Germany. Contact him at meyer@offis.de.

Noel E. O’Connor is a professor in the School of Electronic Engineering and the director of the Insight Centre for Data Analytics at Dublin City University. Contact him at noel.oconnor@dcu.ie.

This article originally appeared in IEEE MultiMedia, *vol. 25, no. 1, 2018.*

Editor: **Gerard J. Holzmann**
Nimble Research
gholzmann@acm.org

Software Components

Gerard J. Holzmann

AT THE 1968 Conference on Software Engineering, mathematician and software engineer Doug McIlroy, alarmed by the sorry state of software development, made a strong pitch for the industrial production of software components.[1] Software systems, like bridges, houses, and cars, are built from parts. McIlroy noted that it didn't make much sense for every organization and developer to keep having to reinvent what's basically a common set of core components for software design. McIlroy envisioned an industry that could provide programmers a selection of mass-produced software parts, differing in accuracy, performance, and cost, to fit a broad range of possible applications.

An inspiration for McIlroy's presentation was the way in which the electronics industry had evolved. Electronics were commonly designed as sets of circuit boards populated with standardized components. There were, and still are, catalogs of resistors, capacitors, diodes, and transistors, with each item documented and marked with the intended range of use. For instance, resistors are marked with a standard color code that indicates their nominal value and the percentage by which their actual value can differ from that nominal value. The user can then make the tradeoff between paying somewhat more for greater precision or less when the highest level of accuracy isn't required. Why couldn't the same thing work in software?

Bricks and Bolts

Resistors, capacitors, and transistors can all be mass produced, like nuts and bolts and bricks, because they're typically used in large quantities. Every device built does, of course, have to use its own copy of all the components it uses, although all those copies are expected to have been produced to the same standards. Just about all circuit board components are standard. Rarely will a circuit designer have to develop an entirely new type of component that's not available in any catalog.

The situation is different in software. The developer needs only one encoding of each standard function, even if the software application that contains it is sold by the millions. The way to achieve accuracy in component design is also different. For instance, to get highly accurate resistors, we can measure the resistance of millions of items and select the ones that deviate by no more than the desired amount to achieve any level of accuracy. If we want to achieve greater accuracy for a software function, the best method is probably not to have large numbers of developers design a version and then select the best one, although in practice it does sometimes seem to work that way. It can be time consuming to rigorously prove the correctness of a software component, but once that work is done, the result should hold for every copy that's used later.

Another difference is that the large majority of the code written for a new application is usually unique to that application. In one of the larger programs I've maintained for a couple of decades, less than 15 percent of the functions originate in standard libraries. The remaining 85 percent are special-purpose functions unique to the application, starting with the main function itself. I suspect that the same is true for most software applications.

Libraries

The mass production in McIlroy's proposal referred to the creation of a larger-than-usual range of specialized features, target execution platforms, and environments that could be created for each basic routine. A good case can be made for this, but there doesn't appear to be a market for this type of software component industry. Pretty much all languages come freely with extensive libraries that encode everything from the most standard to the most exotic types of applications. Today, almost no one would consider writing his or her own library of trigonometric functions or regular-expression pattern matching.

Can we trust the reliability of all those libraries? Is the standard implementation of a sine routine in

Published by the IEEE Computer Society 2469-7087/19/$33.00

Java identical to the one that's available for C or Go? Who exactly takes responsibility for the accuracy of these functions, and how can you get bugs repaired quickly? After all, if you're not a paying customer, you no longer get to decide what's a bug and what's a feature, so you might just have to live with whatever the anonymous provider of the library decides. Unless, of course, you do decide to build everything from scratch.

Modules and Subsystems

Something else that might have changed since 1968 is the notion of what a software component actually is. To a carmaker, a component is the entire entertainment subsystem, the navigation module, the airbag unit, or the engine control unit. Today, no carmaker designs, builds, and programs its own versions of these modules: it purchases them from contractors that specialize in building them and that are responsible for fixing them if they're inadequate. Each of these modules typically comes with substantial amounts of embedded software to bring the hardware to life and give it its desired functionality. If there's a problem, your favorite repair place will not go hunting for any bugs, locate that burned-out transistor, or even try to upload new software into one of these modules. The repair shop will simply replace the entire module. Problem solved. In all these cases, a component is a subsystem, software included.

In the last 50 years, an industry has also been created for the development of specialized software tools. So, here we might need to consider specialized tools such as Photoshop or TurboTax to be identifiable software "components." Who would try to develop his or her own image-editing software, text editor, logic-model checker, or static-source-code-analysis tool today? Well, okay, that would be people like me, but we know we're a very small minority.

Software Tools

The notion of a software tool for solving problems was exemplified best in the design of Unix. If you're used to developing code on a Unix or Unix-like system, you'll have come to rely on standard tools such as `make`, `grep`, `sed`, `awk`, `sort`, `diff`, and `tr`. Each of these specialized tools can be used as a component part for solving larger problems. Each tool aims to solve one specific problem as efficiently as possible. And each tool is designed to support a standard I/O format, so that the output from one tool can be fed into any of the others.

We needn't be surprised that the design of Unix, spearheaded by Ken Thompson, originated in the group at Bell Labs that was created by Doug McIlroy. In this group, which I was fortunate to be part of, Doug McIlroy was a gentle force of inspiration behind a lot of the research that came out in subsequent years. As is well known, McIlroy also contributed the key concept of a Unix pipe as a simple notation to connect the standard output of one tool to the standard input of another. The pipe was the glue that was needed to let us build larger software systems from software components.

The Perpetual Crisis

The rapidly increasing size and complexity of software applications was discussed at length at the 1968 Conference on Software Engineering, famously leading to the first recorded use of the term "software crisis."[1] As noted in the discussion transcripts,

> *Particularly alarming is the seemingly unavoidable fallibility of large software, since a malfunction in an advanced hardware-software system can be a matter of life and death.*[1]

You probably would have believed me if I said that quote was from last year instead of 50 years ago.

An example of a large software system that was discussed at the 1968 conference was the OS for the IBM System/360. A chart (see Figure 1) illustrated the trend. The chart showed the size of IBM OS/360 increasing from around 1 million to around 7 million lines of assembly code, or the equivalent of about 1.4 million lines of C code today. The chart also illustrated the general trend of software size increasing by about a factor of 100 from 1958 to

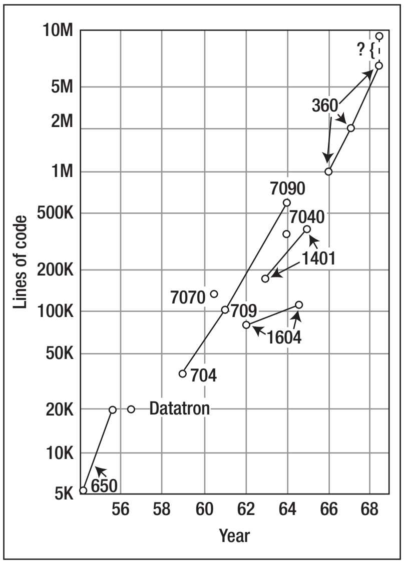

FIGURE 1. The growth in software size, 1956 to 1968.[1] The growth continues, although the pace has slowed.

ABOUT THE AUTHOR

GERARD J. HOLZMANN works on developing stronger methods for the design and analysis of safety-critical software as a consultant and researcher at Nimble Research. Contact him at gholzmann@acm.org.

1968. Luckily, that pace of growth didn't continue, or we would be seeing software applications with tens of billions of lines of code.

But clearly, the growth trend hasn't stopped, and we still feel we're nearing the point at which we'll lose all intellectual control over the software systems we're routinely creating. Of course, a well-designed large system isn't a single homogenous blob of intertwined code. It consists of many parts with, hopefully, well-designed interfaces and a limited number of functional dependencies so that all these components can be designed and checked independently. Anyhow, we can always hope.

Speed

One thing that has changed dramatically since 1968 is the speed and size of the machines we can use to execute our code. Even a humble Raspberry Pi C executes about 2,000 times faster than the fastest IBM System/360 model, the Model 75, did in 1965. For that matter, that Raspberry Pi is also more than 10 times faster than a Cray-1 supercomputer from 1975.[2–4]

Looking on the bright side again, these phenomenal gains in speed not only make our code run faster but also make it possible to analyze our code more thoroughly than ever before. In early compilers, for instance, many checks for consistency and correctness that could in principle have been included were left out to avoid slowing down compilation too much. Even so, in the '60s and '70s you could often literally take a coffee or lunch break before a larger application would finish compiling. Today's compilers can execute fast enough that they can routinely perform far more sophisticated types of checks of our code, without anyone noticing a difference in performance. And we now also have a range of dedicated static-source-code-analysis tools for even more rigorous types of checks. These tools are good enough that they've pretty much become standard in industrial software development.

Indeed, many application domains have a well-defined process for building larger software systems from reusable modules with well-defined interfaces. For instance, if you build spacecraft, the mission of the system generally differs from one vehicle to the next, but all the missions need certain specific types of functionality. Standard software modules needed on every interplanetary mission include the code for navigation (getting from Florida to Mars without too many course corrections), telemetry (getting mission data back to the ground controllers), resource arbitration (don't start driving your rover when it's busy drilling into a rock), and data handling (compression, packetization, and so on).

In the long run, it pays to develop generic software components for all these functions as robustly as possible, and to put them in the arsenal of software components on which mission designers can rely. With time and experience, these modules can be expected to get better and better, and handle more and more cases.

It should go without saying that this approach to the development of software components is well worth using, even if your job is not to fly the occasional mission to Mars.

References

1. *Software Engineering*, P. Naur and B. Randell, eds., Scientific Affairs Division, NATO, 1968; homepages.cs.ncl.ac.uk/brian.randell/NATO/nato1968.PDF.
2. L. Poulsen, "IBM 360/370/3090/390 Model Numbers"; www.beagle-ears.com/lars/engineer/comphist/model360.htm.
3. J. Armstrong, tweet, 17 Jan. 2018; twitter.com/joeerl/status/953711344783691777.
4. R. Longbottom, "Roy Longbottom's Raspberry Pi, Pi 2 and Pi 3 Benchmarks"; www.roylongbottom.org.uk/Raspberry percent20Pi percent20Benchmarks.htm.

This article originally appeared in IEEE Software, *vol. 35, no. 3, 2018.*

Persistent Memory: Abstractions, Abstractions, and Abstractions

Yan Solihin
University of Central Florida

Nonvolatile memory technologies have matured to the extent that commercial DIMM memory products are available today, for example Intel Optane DC Persistent Memory. Compared to DRAM, they provide higher density, better scaling potential, lower cost per bit, and are nonvolatile. Just like DRAM, they provide random access, byte addressability, and can be accessed through load/store instructions. However, writes are slow, expensive, and are limited in endurance. Nevertheless, on the whole, nonvolatile main memory (NVMM) is here and will replace or augment DRAM as main memory fabric.

Utilizing NVMM for capacity is obvious, but the biggest potential for NVMM is providing fabric for hosting data persistently. Currently, when program wants to keep data persistently across process lifetime, it must interface with the file system. When a process wants to work on data, it must create data structures in its address space, read data from a file to populate them. When the process terminates, its address space is reclaimed; data must be written to the file prior to that. The transition between memory and storage is expensive and unnecessary when the main memory is nonvolatile.

There are many challenges in utilizing NVMM. The fact is that memory and storage have had separate interfaces for decades and have been optimized in isolation for different goals. Hence, NVMM is a disruptive technology that requires revisiting this dichotomy for many years to come. The thesis of this paper is that how well systems and programs use NVMM depends on the design of abstractions. Abstraction affects the programming complexity, efficiency of persistency, and scalability.

In programming languages, there are in general three layers of abstraction: assembly language (AL), high-level language (HLL), and domain-specific language (DSL). These abstractions evolved over decades and provide distinct efficiently views of the machine. I argue that utilizing NVMM requires similar layers of abstractions as well.

Current research work in persistency programming mainly focuses on the persistency model. A persistency model defines *durability ordering*, i.e., in what order stores are made durable (reaching the nonvolatile domain) relative to the order in which they appear in the program, and *atomic durability*, in which multiple stores must be made durable together or none at all. An example is *epoch persistency*, which defines epochs separated by persist barriers. Within an epoch, all stores are made durable in any order. However, stores of the one epoch must be durable prior to any store in the next epoch is made durable. Instruction set architecture support for persistency includes instructions to force a cache line to reach the NVMM (e.g., CLFLUSH, CLFLUSHOPT, and CLWB), while a persist barrier is provided through a serializing instruction (e.g., SFENCE). Atomic durability is provided through durable

Digital Object Identifier 10.1109/MM.2018.2885589
Date of current version 21 February 2019.

transactions provided either in hardware or software.

Current persistency support overheads are too high. Even with speculation past a persist barrier,[1] substantial slowdown still results. Logging is another source of overheads, if performed in software. Logging in hardware restricts the count and size of durable transactions, hence a hybrid scheme promises the best of both worlds.[2]

> The thesis of this paper is that how well systems and programs use NVMM depends on the design of abstractions.

Current persistency model abstraction is also too low level, akin to AL abstraction, where programmers must be aware of the memory hierarchy and using appropriate assembly instructions. HLL persistency model is needed to simplify programming, for example language support that allows programmers to define which data is persistent, and which regions must be durable atomically. Abstraction simplicity leads to shorter software development time, and higher quality and reliability of the resulting code.

HLL model must also interact with the system. Hence, the system must provide data abstraction in order to host data persistently in the main memory. Persistency requires data to be encapsulated as a persistent memory object (PMO) that can be located across process lifetime, crashes and reboots, similar to a file. A PMO also needs system-managed permission, sharing, and efficient meta data. Recent research has looked into NVMM file system (e.g., PMFS, BPFS, NOVA, and NOVA-Fortis), or memory mapped files (e.g., Mnemosyne). Both approaches have significant drawbacks: the former incurs much software overhead, while the latter needs to reconcile two systems (virtual memory and file system) that are not necessarily compatible. I believe a PMO abstraction needs to adopt only necessary features from a file system and keep simple metadata.

Beyond HLL, DSL abstraction should also be developed. For example, key-value stores translate quite well to NVMM. Being higher level than HLL, DSL persistency abstraction needs not necessarily lead to high performance overheads, because certain semantics applicable for a particular domain may present an opportunity to relax the persistency model. For example, consider checkpointing in HPC systems. Instead of taking a snapshot of memory to create a checkpoint in the storage system, one may simply persist data structures regularly. If a failure occurs, one may recover by reading the data structures, and recomputing all noncompleted and partially completed regions of code.[3] We may also perform *lazy persistency*, where code regions are protected by error detection codes that detect incomplete persistency.[4] Upon failures, if some regions are detected to not have fully persisted, they are re-executed. Lazy persistency removes the need to perform logging, flush cache lines, and use persist barriers. The result is much faster execution and negligible write amplification, but at the cost of slower and more complex recovery.

Coming up with the right abstractions is very important for NVMM to harness its best potentials. In this paper, I touched upon a few aspects of abstractions, but many more aspects require future research.

REFERENCES

1. S. Shin, J. Tuck, and Y. Solihin, "Hiding the long latency of persist barriers using speculative execution," in *Proc. Int. Symp. Comput. Archit.*, 2017, pp. 175–186.
2. S. Shin, S. K. Tirukkovalluri, J. Tuck, and Y. Solihin, "Proteus: A flexible and fast software supported hardware logging approach for NVM," in *Proc. Int. Symp. Microarchit.*, 2017, pp. 178–190.
3. H. Elnawawy, M. Alshboul, J. Tuck, and Y. Solihin, "Efficient checkpointing of loop-based codes for non-volatile main memory," in *Proc. Int. Conf. Parallel Archit. Compilation Techn.*, 2017, pp. 318–329.
4. M. Alshboul, J. Tuck, and Y. Solihin, "Lazy persistency: A high-performing and write-efficient software persistency technique," in *Proc. Int. Symp. Comput. Archit.*, 2018, pp. 439–451.

Yan Solihin is the Director of Cybersecurity & Privacy Cluster and Charles N. Millican Professor of Computer Science at University of Central Florida, Orlando, USA. Contact him at yan.solihin@ucf.edu.

This article originally appeared in IEEE Micro, *vol. 39, no. 1, 2019.*

Quantum Computer Scale-up

Erik P. DeBenedictis, Sandia National Laboratories
Travis S. Humble, Oak Ridge National Laboratory
Paolo A. Gargini, IEEE International Roadmap for Devices and Systems

Quantum computers are creating buzz from Silicon Valley to the halls of Congress. But the hardware has only been prototyped at small scale. Could applying the engineering processes used to scale up semiconductors also accelerate quantum computer development?

A couple of decades after the invention of the transistor, the engineering community developed the integrated circuit to create computers out of large numbers of transistors. The initial invention launched an industry whose reinvested profits scaled up component count by a factor of a million—a growth principle now known as Moore's law[1]—and the computers built with them changed the world.

A few decades later, physicists and mathematicians came up with a new computing device called the qubit, which has the tantalizing, albeit unproven, potential to similarly change the world. Researchers have created small-scale prototype quantum computers and are now working with IEEE with the intent of making full-sized systems.

To better understand the current approach to quantum computer scale-up, let's first look at how transistors morphed into integrated circuits, and how integrated circuits then morphed into the colossus of classical computation today. In this column, we show a way forward that combines new quantum computer benchmarks with the processes used for semiconductor scaling.

SCALING UP MANUFACTURING

In some cases, scalability is easy to see. Figure 1a is a planar integrated semiconductor circuit, and Figure 1b is an integrated qubit circuit. They look pretty much the same, so we might assume that both will scale according to Moore's law. While this could be true, there's more to the story.

Figure 1c shows an earlier "flying-wire" integrated semiconductor circuit, with the transistors integrated as they are today but hooked up with gold wires above the surface. Manufacturing a million-component flying-wire integrated circuit would require placing millions of tiny gold wires during manufacture.

 Published by the IEEE Computer Society

EDITOR ERIK P. DEBENEDICTIS
Sandia National Laboratories;
epdeben@sandia.gov

The gold chandelier-like configuration in Figure 1d is the support structure of a quantum computer research demonstration, with the figure's qubit chip mounted at the bottom. The chandelier connects the cryogenically cooled quantum computer components to various stages of control electronics and then to the room-temperature electronics at the top. Scaling up the quantum computer from 7 qubits to 1 million qubits would require a proportional increase in the chandelier's complexity. Imagine scaling up the flying-wire chip or chandelier by 100,000×. They would become huge masses of mechanical structures that not only would be difficult to manufacture but would also contain numerous parts that could break during operation. It's apparent from just looking at them that they won't scale. However, the planar and qubit chips *do* have a tidy structure that looks like it will scale, and, in fact, these types of structures can be manufactured at scale. However, the quantum computer demonstration combines the chandelier and qubit chip, so the combined structure won't scale. No quantum computer design scales. Therefore, we need a better design.

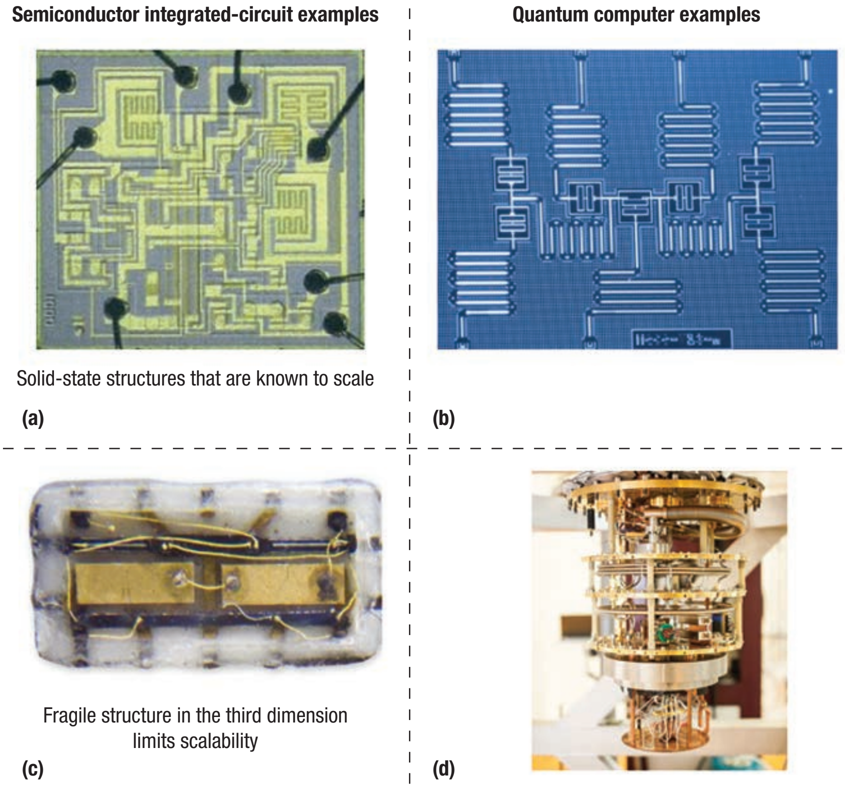

Figure 1. Scalable (a) semiconductor circuit (planar integrated circuit circa 1971 [555 timer]; https://en.wikipedia.org/wiki/555_timer_IC) and (b) qubit integrated circuit (IBM 7 qubit chip, at the bottom of the chandelier; https://spectrum.ieee.org/tech-talk/computing/hardware/tiny-quantum-computer-simulates-big-molecules). However, the qubit integrated circuit is just part of a quantum computer. (c) Flying-wire integrated circuit (Credit: Mark Richards. Courtesy of the Computer History Museum) and (d) quantum computer chandelier (dilution refrigerator cold stage; https://spectrum.ieee.org/computing/hardware/europe-will-spend-1-billion-to-turn-quantum-physics-into-quantum-technology), neither of which is scalable in manufacture because of the fragile structure in the third dimension. Therefore, the quantum computer comprising (b) and (d) is not scalable because of the chandelier.

SCALING UP FUNCTION

Planar integrated circuits such as those in Figure 1a are manufactured similarly to Henry Ford's assembly line for making large volumes of automobiles; however, another type of scalability arrived on the scene, becoming the key to computers' colossal impact on the world. We initially assumed that the size of each transistor would not change. This meant that chips with more complex functions would require more transistors—which would use more area, leading to fewer chips per wafer and more defects per chip. This led to cost–complexity curves such as those in Figure 2a, where the optimal point is at the bottom of each curve.

However, Gordon Moore's 1965 article[1] reported results from a sequence of several generations of integrated circuits that created more complex functions from one generation to the next, yet with the functions occupying the same chip area because the transistors were simultaneously getting smaller. Under this condition, the cost–complexity curves reproduced in Figure 2a moved toward both lower cost per component and more components per chip. On the basis of these assumptions, Moore predicted that the number of transistors that could be placed on each single die on a wafer would increase exponentially to 65,000 by 1975, as shown in Figure 2b. This exponential improvement is now called Moore's law.

If quantum computers were to scale up exponentially, it would be due to a

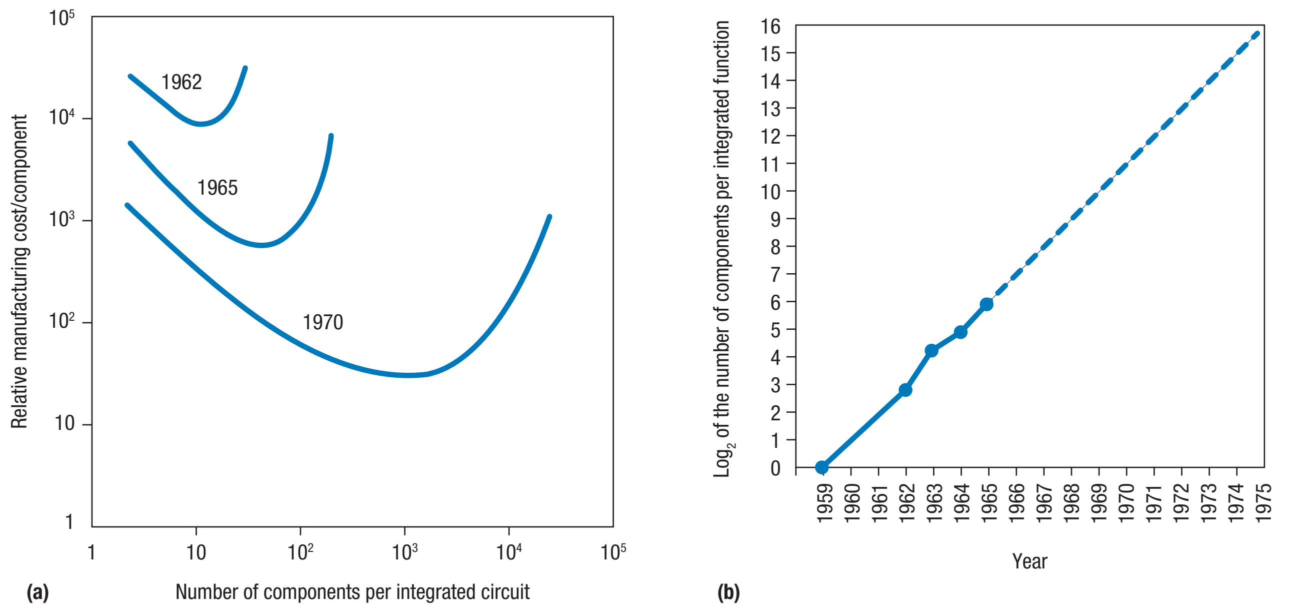

Figure 2. Data from (a) Moore's article[1] that supports (b) extrapolating optimum chip size to future years. The graph in (a) has been largely forgotten in public discourse, with the graph in (b) coming to represent Moore's law. However, the data will need to be reassessed for quantum computers before they can be considered scalable—if they actually are destined to be scalable.

different effect. Quantum computer components must be very precise to support quantum entanglement, but quantum entanglement can give exponential quantum speedup in some cases, or performance on specific problems that rises exponentially as the number of qubits increases. Quantum computers could scale exponentially if quantum chips could hold just one more qubit from one generation to the next—or if the qubits could be more precise from one generation to the next.

EFFECT OF MATERIALS

Bipolar transistors dominated the initial production of integrated circuits because they were faster and smaller than metal-oxide semiconductor (MOS) transistors, but bipolar transistors had an intrinsic scaling problem. Despite many attempts, it wasn't possible to reduce the number of defects in bipolar chips to MOS levels. This led to a market advantage for MOS, and, by 1975, MOS products had conquered the chip market.

Qubits have similar materials issues. The equivalent to the competition between bipolar and MOS transistors is Josephson junction (transmon) qubits, quantum dot qubits, and other types. All qubits are subject to *decoherence*—a loss of information when the engineered qubits interact with impurities in the material. The impurities become extra, parasitic qubits and produce errors when they couple too strongly with the engineered qubits.

The quest to reduce decoherence caused by materials impurities parallels the quest to reduce defects in bipolar and MOS transistors. Nobody currently knows how much decoherence can be reduced in any type of qubit, so nobody knows which type of qubit will be the most scalable in the long run.

REPEATING LESSONS FROM SEMICONDUCTOR SCALABILITY

What would it take to apply the process used to develop scalable semiconductors to quantum computers? The process includes manufacturing, function, and materials considerations, as described above. However, there are other issues, as Moore's short article detailed:[1]

- cost,
- reliability,
- material (silicon),
- yield,
- complexity,
- die size,
- interconnection space,
- heat,
- speed,
- power per unit area,
- design automation, and
- linear & RF.

The difference in design objectives between semiconductors and qubits isn't mysterious. David DiVincenzo documented the objectives for gate-type quantum computers,[2] and his objectives are well accepted by the quantum computing community as the following "DiVincenzo criteria":

1. A scalable physical system with well-characterized qubits.
2. The ability to initialize the state of the qubits to simple fiducial state, such as to |000...>.
3. A "universal" set of quantum gates.

This article originally appeared in Computer, *vol. 51, no. 10, 2018.*

TABLE 1. Tentative benchmark hierarchy.

Benchmark level	Gate-type quantum computer	Quantum annealer	Classical computer for reference
System	Yet to be developed	Yet to be developed	Clock period
Gate	Quantum volume, including quantum speed-up	Metric to be developed	Gate delay; speed-up not relevant
Device	Decoherence time	Decoherence time	Gain

4. Long relevant decoherence times, much longer than the gate-operation time.
5. A qubit-specific measurement capability.

The first DiVincenzo criterion uses the phrase "scalable physical system," which is essentially a three-word description of Moore's criterion and the enormous infrastructure developed by the semiconductor industry for managing computer system scale-up. We suggest that combining Moore's list with the DiVincenzo criteria would generate appropriate design objectives for scalable quantum computers.

QUANTUM COMPUTER BENCHMARKS

Efforts toward quantum computer scale-up started informally a few years ago, but were the topic of a more formal meeting 30–31 August 2018 among IEEE, representatives of the nascent quantum computer industry, and the US government. The smooth roadmap curves that plot metrics such as memory density and clock rate aren't appropriate for the earlier stage of quantum computers. So instead, the meeting participants developed the tentative matrix of quantum computer benchmarks shown in Table 1, which differentiates between the potential uses of quantum computing and the levels of engineering required to realize them. When this table is eventually filled in with specific metrics—some yet to be developed—and accepted by a consensus of the relevant parties, the resulting benchmarks will enable new insights into the progress of scaling up quantum computers.

We don't know whether quantum computers are destined to satisfy the applications-level expectations being bandied about from the media to the halls of Congress. There's a big gap between current small-scale prototypes and the millions of qubits required by proposed applications. The best hope lies in R&D to improve existing qubit designs until one or more approaches become scalable. Some of the world's largest companies are competing with one another and with well-funded start-ups to address the issues. However, the development of large-scale quantum computers might end up being too difficult for a single company. Our recent meeting recognizes that the IEEE community can work together to raise the probability of success.

ACKNOWLEDGMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the US Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. The views expressed in the article do not necessarily represent the views of the US Department of Energy or the United States Government. This manuscript has been authored by a contractor of the US Government under Contract No. DE-AC05-00OR22725. Accordingly, the US Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for US Government purposes. This document in part summarizes a meeting held 30–31 August 2018 and comprising the following individuals: Paul Bunyk, Steve Bush, Josh Combes, Tom Conte, Marcus da Silva, Erik DeBenedictis, David Ferguson, Paolo Gargini, Jean-Luc Gaudiot, Lee Gomes, Pat Gumann, Scott Holmes, Travis Humble, Alan Kadin, Scott Kozoil, Bruce Kraemer, Norbert Linke, Alex McCaskey, Cathy McGeoch, Peter Mueller, Oleg Mukanov, Paul Nation, Satyavolu Papa Rao, Borja Peropadre, Raphael Pooser, Moinuddin Qureshi, Mark Ritter, Andrew Sornborger, John Spargo, Himanshu Thapliyal, Bill Tonti, Elie Track, and Carl Williams.

REFERENCES

1. G.E. Moore, "Cramming More Components onto Integrated Circuits," *Electronics*, 1965, pp. 114–117.
2. D. DiVincenzo, "The Physical Implementation of Quantum Computation," *Fortschritte der Physik–Progress of Physics*, vol. 48.9, no. 11, 2000, pp. 771–783.

ERIK P. DEBENEDICTIS is a technical staff member at Sandia National Laboratories Center for Computing Research. He's a Senior Member of IEEE and a member of ACM and APS. Contact him at epdeben@sandia.gov.

TRAVIS S. HUMBLE is a distinguished scientist at Oak Ridge National Laboratory. He's a Senior Member of IEEE. Contact him at humblets@ornl.gov.

PAOLO A. GARGINI is chairman of the IEEE International Roadmap for Devices and Systems. Contact him at paologargini1@gmail.com.

FDA U.S. FOOD & DRUG ADMINISTRATION
CENTER FOR BIOLOGICS EVALUATION & RESEARCH

STAFF FELLOW: COMPUTER SCIENTIST

The FDA's Center for Biologics Evaluation and Research (CBER), Office of Biostatistics and Epidemiology (OBE) is recruiting to fill a Staff Fellow, Computer Scientist position in their **High Performance Integrated Virtual Environment (HIVE) Team**. The HIVE provides bioinformatics support services, algorithm development, and visualizations to assist researchers and regulators in answering complex scientific questions. For more information, please visit https://github.com/FDA/fda-hive.

The selected candidate will collaborate with other FDA scientists, including physicians, laboratory scientists, life sciences professionals, biostatisticians and epidemiologists to develop and provide high-performance bioinformatics tools, natural language processing solutions, informative data visualizations, and core data analysis services.

This position will be filled through the Service Fellowship program. The appointment is for up to four years with opportunities for renewal. **Applications will be accepted from U.S. citizens or Lawful Permanent Residents (green card holders) only**. No previous Federal experience is required. Appointment does not confer any entitlement to a position in the competitive service, and there is no entitlement to Merit Systems Protection Board (MSPB) appeals rights.

QUALIFICATIONS: Official transcripts required; Ph.D in Computer Science, Bioinformatics, Computational biology, Mathematics, Computer Modeling or related fields are encouraged to apply. **NOTE: Applicants must also demonstrate that they meet the basic requirements for the Computer Scientist, RG-1550 series to be considered.**

REQUIRED SPECIALIZED EXPERIENCE: Developing and maintaining high-throughput bioinformatics infrastructures for cluster and cloud computing platforms, high performance computing, artificial intelligence, NGS clinical genomics, and bioinformatics. Also, prior experience with software engineering and specific training in bioinformatics is desired.

SALARY: Salary commensurate with education/experience.

LOCATION: Silver Spring, MD

HOW TO APPLY: Submit electronic resume or curriculum vitae and supporting documentation to Brian Hively at brian.hively@fda.hhs.gov. Applications will be accepted through **August 13, 2019**. Please reference Job Code: **OBE-19-04-IEE**.

NOTE: This position may be subject to FDA's strict prohibited financial interest regulation and may require the incumbent to divest of certain financial interests. Applicants are strongly advised to seek additional information on this requirement from the FDA hiring official before accepting a position.

DEPARTMENT OF HEALTH AND HUMAN SERVICES IS AN EQUAL OPPORTUNITY EMPLOYER SMOKE FREE ENVIRONMENT

---

der Bundeswehr
Universität München

The research institute for Cyber Defence (CODE) at the Bundeswehr University Munich has been substantially expanded to become the research institute for "Cyber Defence and Smart Data" of the Bundeswehr and the Federal Government institutions. CODE was established in 2013 with the objective of bringing together experts from different faculties and scientific disciplines as well as expertise from industry and government agencies to conduct research in cyber and information space. CODE pursues a comprehensive, integrated, and interdisciplinary approach to implementing technical innovations and concepts for the protection of data, software, and ICT infrastructures in accordance with legal and commercial framework conditions. It has already established important strategic partnerships in this area. The objective of the expansion is to unite the Bundeswehr's and the Federal Government's research initiatives in the area of Cyber Defence and Smart Data and to establish the CODE research institute as the primary point of contact in the cyber and information domains of the Bundeswehr and the Federal Government.

Research and teaching in the area of cyber security is already being carried out as part of the bachelor's and master's programs in the Department of Computer Science. A new independent master's program in Cyber Security was launched on January 1st, 2018.

The Department of Computer Science at the Bundeswehr University Munich is seeking a professor for the following specialist area of its Cyber Defence und Smart Data research institute:

W3 University Professorship for Artificial Intelligence in IT Security

Artificial intelligence methods are used, among other things, to automate fact-based decision-making processes and to simulate and advance the behavior of individual actuators or entire swarms. In the context of IT security, they have the long-term potential not only to automate time-consuming manual analysis tasks of large amounts of data reliably and with low error rates by means of cognitive security, but also to contribute to the protection of networked systems. Therefore, they should be used for the detection of, for example, anomalies or attacks and for the design and implementation of preventive security mechanisms.

Due to the wide range of possible applications of Artificial Intelligence, the W3 professorship has an interface function within the Department of Computer Science and the CODE research institute and serves as a bridge to the engineering, humanities, and social sciences. An excellent, internationally oriented professor is sought who is particularly proven in several of the following areas in research and teaching with a clear reference to IT security:

- Autonomous and distributed AI systems
- Experimental and model-driven cognitive systems
- Machine learning for the detection and prognosis of patterns in data
- Neural Computing, especially Deep Learning and Reinforcement Learning
- Scalable, efficient and resilient analysis algorithms and learning methods

The Bundeswehr University Munich is looking for a professor who, in addition to having outstanding scientific qualifications, will also contribute actively to the CODE research institute.

Besides excellent research work, the new professor is expected to develop demanding lectures, tutorials, and seminars for the master's program in Cyber Security and to provide excellent teaching in her or his respective specialist area.

The applicant is also expected to carry out teaching in the bachelor's programs in computer science and business computer science, and to work closely with the other departments at the Bundeswehr University Munich.

The professorship will be provided with superbly equipped laboratories housed in a new building that will be built in the near future. The candidate must have an excellent scientific track record, as demonstrated by a habilitation or equivalent scientific achievement, as well as relevant publications in academic journals. Proven teaching experience in her or his respective specialist area is highly desired. The new professor should have an international perspective, e.g., based on participation in international research projects and experience in acquiring third-party funding. The duties will also include participation in the academic self-administration of the university. Further, the candidate will be expected to assume a gender equality based leadership role.

The Bundeswehr University Munich offers academic programs directed primarily at officer candidates and officers, who can obtain bachelor's and master's degrees in a trimester system. Study programs are complemented by an integrated program entitled "studium plus" which consists of interdisciplinary seminars, tutorials and training in key professional qualifications.

The recruitment requirements and the employment status of professors are governed by the Federal Civil Service Act (Bundesbeamtengesetz). To become an appointed civil servant (Beamter) candidates must be no older than 50 on the date of their appointment.

The university seeks to increase the number of female professors and thus explicitly invites women to submit applications. Severely disabled candidates with equal qualifications will receive preferential consideration.

Please submit your application documents marked as Confidential Personnel Matter to the *Dean of the Department of Computer Science, Professor Dr. Oliver Rose, Bundeswehr University Munich, D-85577 Neubiberg, via email to dekanat.inf@unibw.de by 30th of September, 2019.*

DHV's HALLMARK — FAIR AND TRANSPARENT APPOINTMENT NEGOTIATIONS
CHARTA FAMILIE IN DER HOCHSCHULE unterzeichnet 2015
charta der vielfalt UNTERZEICHNET

NUS
National University of Singapore

Sung Kah Kay Assistant Professor in All Areas of Computer Science

The Department of Computer Science at the National University of Singapore (NUS) invites applications for the Sung Kah Kay Assistant Professorship. Applicants can be in any area of computer science. This prestigious chair position was set up by the family and friends of the late Assistant Professor Sung Kah Kay after his untimely demise early in his career at NUS. Candidates should be early in their academic careers and yet demonstrate outstanding research potential, and a strong commitment to teaching.

The Department enjoys ample research funding, moderate teaching loads, excellent facilities, and extensive international collaborations. We have a full range of faculty covering all major research areas in computer science and boasts a thriving PhD program that attracts the brightest students from the region and beyond. More information is available at www.comp.nus.edu.sg/careers.

NUS is an equal opportunity employer that offers highly competitive salaries, and is situated in Singapore, an English-speaking cosmopolitan city that is a melting pot of many cultures, both the east and the west. Singapore offers high-quality education and healthcare at all levels, as well as very low tax rates.

Application Details:

Submit the following documents (in a single PDF) online via: https://faces.comp.nus.edu.sg

- A cover letter that indicates the position applied for and the main research interests
- Curriculum Vitae
- A teaching statement
- A research statement

Provide the contact information of 3 referees when submitting your online application, or, arrange for at least 3 references to be sent directly to csrec@comp.nus.edu.sg

Application reviews will commence immediately. We hope to fill the position by January 2020

If you have further enquiries, please contact the Search Committee Chair, Weng-Fai Wong, at csrec@comp.nus.edu.sg.

FAU

COLLEGE OF ENGINEERING & COMPUTER SCIENCE
Florida Atlantic University

The College of Engineering and Computer Science (COECS) at Florida Atlantic University (FAU) is seeking a dynamic individual to serve as the Chair of the Department of Computer & Electrical Engineering and Computer Science (CEECS) - http://www.ceecs.fau.edu/. A visionary leader with demonstrated collaboration and communication skills is sought to further advance the department's mission.

The department, consisting of 39 faculty and 8 educational and research staff members offers, baccalaureate (800 students), masters and PhD degrees (280 students) in computer engineering, electrical engineering and computer science. A multidisciplinary masters of bioengineering is also offered in the department. The department collaborates closely with FAU's ISENSE pillar (http://isense.fau.edu/). The department's current active funding portfolio is over $7.0M, with support from NSF, NIH, DoD, DoE, and Industry.

Applicants should have strong research and scholarly portfolios and will: provide leadership for academic and research programs; promote the reputation of the department at the national and international levels; create cooperation opportunities with industry partners; attract world-class faculty, high-caliber students and staff; promote collaborative research within and outside FAU; and mentor faculty at different stages of their careers.

Minimum Qualification:

- An earned doctorate from an accredited institution in electrical engineering, computer engineering, computer science or a closely-related field by the time of application required.
- Must have demonstrated effective leadership, management and communication with faculty, staff, students, and other administrators.
- Must have commitment to diversity and inclusion as it applies to faculty, staff and students.
- Must have commitment to shared governance, collegiality, and professional development for faculty and staff.

All applicants must apply electronically to the currently posted position (Chairperson and Professor) on the Office of Human Resources' job website (https://fau.edu/jobs) by completing the required online employment application and submitting the related documents. When completing the online application, please upload the following: a cover letter detailing professional experience, curriculum vitae, a statement describing the vision for the department, leadership and mentoring philosophy, teaching and research interests (max 3 pages), contact information for six professional references and copies of official transcripts scanned into an electronic format. Review of applications will begin in August 2019. All applications received by July 31, 2019 will receive full consideration.

A background check will be required for the candidate selected for this position. This position is subject to funding. For more information and to apply, visit www.fau.edu/jobs and go to Apply Now (REQ06164).

Florida Atlantic University is an equal opportunity/affirmative action/equal access institution and all qualified applicants will receive consideration without regard to race, color, religion, sex, sexual orientation, gender identity, national origin, disability status, protected veterans status or other protected status. Individuals with disabilities requiring accommodation, please call 561-297-3057. 711.

FAU is committed to the principles of engaged teaching, research and service. All persons aspiring to achieve excellence in the practice of these principles are encouraged to apply.

Conference Calendar

Questions? Contact conferences@computer.org

IEEE Computer Society conferences are valuable forums for learning on broad and dynamically shifting topics from within the computing profession. With over 200 conferences featuring leading experts and thought leaders, we have an event that is right for you.

Find a region:	Africa ■	Australia ◆	North America ◗
	Asia ▲	Europe ●	South America ★

AUGUST

1 August

- CSE (IEEE Int'l Conf. on Computational Science and Eng.) ◗
- EUC (IEEE Int'l Conf. on Embedded and Ubiquitous Computing) ◗

5 August

- TrustCom (IEEE Int'l Conf. on Trust, Security and Privacy in Computing and Communications) ◆

8 August

- 2019 Cloud Summit ◗

9 August

- SmartIoT (IEEE Int'l Conf. on Smart Internet of Things) ▲

10 August

- HPCC (IEEE 21st Int'l Conf. on High Performance Computing and Communications) ▲
- SmartCity (IEEE 17th Int'l Conf. on Smart City) ▲
- DSS (IEEE 5th Int'l Conf. on Data Science and Systems) ▲

15 August

- NAS (IEEE Int'l Conf. on Networking, Architecture and Storage) ▲

18 August

- HCS (IEEE Hot Chips 31 Symposium) ◗
- RTCSA (IEEE 25th Int'l Conf. on Embedded and Real-Time Computing Systems and Applications) ▲

27 August

- ASONAM (IEEE/ACM Int'l Conf. on Advances in Social Networks Analysis and Mining) ◗

SEPTEMBER

13 September

- EWDTS (IEEE East-West Design & Test Symposium) ●

19 September

- AVSS (16th IEEE Int'l Conf. on Advanced Video and Signal Based Surveillance) ▲
- ESEM (ACM/IEEE Int'l Symposium on Empirical Software Eng. and Measurement) ★ ●