Research Statement

Youngki Lee

School of Information Systems, Singapore Management University

Tel: (65) 6808-5421; Email: youngkilee@smu.edu.sg

19 Jan. 2015

Background

Computing services have been rapidly immersed into every facet of our daily lives including child education, elderly support, daily well-being, smart transportation, automated homes and buildings. Especially, *life-immersive sensing applications* are emerging on top of commodity mobile, wearable, sensor devices, accelerating this unprecedented wave. Such applications are clearly distinguished from conventional mobile applications, in that they provide *proactive* services fitting to situational contexts of mobile users; in contrast, conventional applications provide rather passive services reactively upon user requests. Such *proactive-ness* will be extremely important in future mobile service design since mobile users cannot pay full attention to the applications as they perform other primary activities. Moreover, sensing applications will evolve to capture and use high-level *human-centric contexts* (such as intention, engagement, emotion, attention, fatigue, anxiety, depression, distractibility, mindfulness, etc.) beyond low-level physical contexts such as location and activity [6]. Given such capabilities, a future context-aware advertisement, for example, will not only target users based on their location and activity, but also more precisely target people who has buying intention to purchase certain products. Moreover, discount coupons can be sent out when the user is bored or has free attentive capacity.

The core of my research is to enable highly-enriched and always-available context awareness as a common underlying platform. This new platform will accelerate development and deployment of useful life-immersive sensing applications; despite their high practical value, their spread has been extremely slow due to unprecedented challenges to accurately capture rich user contexts in dynamic real-life situations. The resource limitations of mobile and wearable computing devices aggravate these challenges significantly. I have pursued this vision even before smartphones are deeply penetrated into our lives and set meaningful milestones to address key challenges to realize the goal. More specifically, I have been continuously working on the following specific research topics to bring this future mobile platform into reality and make breakthrough the current state-of-the-art techniques in related research fields. Moreover, I have been building various life-immersive applications on top of my platform, collaborating with a number of experts in various application domains.

Building Context Monitoring Platform: with Focus on Resource Scarcity

My main research thread has been designing and implementing novel mobile systems to enable continuous monitoring of rich user contexts as an underlying platform [1][3][4][7][14]. A core challenge results from continuous execution of highly complex, multi-step operations to infer context (i.e., sensing, data cleaning, feature extraction, classification), on top of resource-scarce smartphones and sensor devices.

I proposed systems and techniques to address the challenges in several papers. A key underlying approach behind the proposed systems is the *translation-based approach* [3][4][14]. This approach has significant advantages for both application developers and the system; it allows developers to easily specify their contexts of interests with semantic-level queries (e.g., let me know if a user is shopping with friends), while providing the systems with flexibility to translate the queries into diverse low-level processing details with suitable resource optimization techniques. Amongst various optimizations we have done, the method of incremental context processing and the notion of the essential sensor set are particularly noticeable [14]. Behind the incremental processing technique, we had an interesting and useful observation of context continuity, that is, contexts tend to change rather gradually, which had been unfolded at the first time to the research community.

The concept of context continuity have been inspiring the researchers who design a system for continuous context monitoring, being employed as a versatile and important heuristic.

Coordinating Multiple Sensing Applications. The resource management for mobile context monitoring becomes much more challenging as users run multiple sensing applications simultaneously; note that such applications provide proactive services without requiring users' attention, and thus, users are likely to run multiple applications in the background. In my papers, I pointed out that resource coordination for concurrent sensing workload will be a key system functionality of future mobile systems. Specifically, Orchestrator [3][4] and SymPhoney [7] delved into the challenges of severe resource contention by concurrent sensing applications, and proposed resource planning mechanisms to prevent skewness in applications' battery and CPU use. They are recognized as one of the first systems to support resource coordination among concurrent sensing applications, whereas most previous sensing systems assume a single outstanding application running in a system.

In more detail, Orchestrator focuses on a future personal sensor network environment, which is composed of a smartphone and multiple heterogeneous wearable sensor devices. It generates processing plans to distribute concurrent sensing tasks in a balanced way, newly utilizing diverse sensor fusion methods to monitor contexts. Through the planning, it prevents the situation that a specific sensor is particularly overloaded while others are not. On the other hand, SymPhoney coordinates sensing applications' resource use on a smartphone. It newly identifies internal processing structure of sensing dataflow, and leverages its unique characteristics to control the processing intervals and scheduling orders of concurrent applications. Moreover, SymPhoney is notable in terms of its programming model, as it newly adopts and supports a dataflow model to facilitate specification of highly-customized and complex sensing applications.

Bringing Social Awareness into the Platform.

Another important research direction I pursue is to bring in-situ social awareness into the mobile platform. Despite prevalent smartphone uses, it is still at an early dawn that everyday gathering and social activities are deeply considered in the design of mobile applications and systems. Importantly, I envision that co-presenting smartphones serve as a newly emerging computing substrate, opening new research opportunities. First, it could potentially accommodate completely new social sensing applications, for example, group-context-aware mobile advertisement in commercial complexes, in-situ social networking among shoppers, and a new mobile gaming among crowds in a metro/bus. Second, such new substrate could enable significant improvement in previous mobile systems by sharing the larger, more capable union of computing devices and resources.

I took a meaningful first step towards the latter direction, and built a cooperative platform for context monitoring, CoMon [1]. CoMon leverages in-situ cooperation among nearby mobile users, and opens a new possibility to addresses an inherent battery problem in continuous sensing. Specifically, for the design of CoMon, we made inspiring observations from mobility data analysis: (1) people spend a significant portion of their time with acquaintances, ensuring enough cooperation time (>8 hours a day in average), (2) meetings often last long (>50 minutes per meeting in average), enabling stable cooperation, (3) The power consumption for sensing and context inference often exceeds the overhead to obtain inferred context results from nearby users. These observations are not limited to CoMon but provide important design guidelines to build other cooperative mobile systems in the future. CoMon also incorporates two key techniques to achieve mutual energy saving among users, i.e., continuity-aware cooperator detection and a benefit-aware negotiation mechanism, which shows an exemplary way to leverage the above observations.

Toward the former direction, I proposed SocioPhone, a mobile platform for face-to-face interaction monitoring [2]. Face-to-face interaction, especially conversation, is a fundamental part of everyday life. In mobile and pervasive computing domain, face-to-face interaction is often simply abstracted as a contact in proximity; however, to design more useful applications for interaction situations, much deeper contextual understanding is required. Through SocioPhone design, we identify useful meta-linguistic contexts of conversation, such as turn-takings, prosodic features, a dominant participant, and pace. SocioPhone abstracts such useful meta-linguistic contexts as a set of intuitive APIs. Its runtime efficiently monitors registered contexts during in-progress conversations and notifies applications on-the-fly- note that recognizing such delicate conversational contexts with commodity mobile devices involves severe resource problems. Importantly, we have noticed that online turn monitoring is the key building block for extracting diverse meta-linguistic contexts, and devised a novel volume-topography-based method that enables highly accurate and energy-efficient turn monitoring.

Taking a step further, we have been building useful interaction-aware applications on top of SocioPhone. One outstanding application is TalkBetter [11] that is designed for the parents who have a child with language delay issue. Language delay is a developmental problem of children who do not acquire language as expected for their chronological ages. Without timely intervention, language delay can act as a lifelong risk factor. Speech-language pathologists highlight that effective parent participation in everyday parent-child conversation is important to treat children’s language delay. For effective roles, however, parents need to alter their own lifelong-established conversation habits, requiring extensive period of conscious effort and staying alert. TalkBetter provides in-situ intervention service to parents in daily parent-child conversation through real-time meta-linguistic analysis of ongoing conversations. We have been doing extensive field studies with speech-language pathologists and parents, and study multilateral motivations and implications of TalkBetter.

Developing innovative Life-immersive Sensing Applications.

On top of my above introduced sensing platform, I have been continuously designing new life-immersive sensing applications to innovate daily life of people, extend human cognition, and bring fun and convenience. I have been actively collaborating with domain experts in various application domains such as child education, advertisement, healthcare, city planning, and published interesting initial results. I below describe a few example applications.

Childcare: Childcare and kindergarten education are very suitable application domains for sensing technology. This is mainly because children naturally require continuous observation for their safety and education. Moreover, in many cases, parents and teachers (especially teachers who need to take care of many children at the same time) may not have sufficient cognitive power to take care of children’s safety while finding educationally meaningful behaviors to help their development. We have been working on various applications to help parents and teachers for more effective childcare. TalkBetter [11] introduced above is a nice example application in this direction. Also, we built a useful application to find out unique behavior of a child with respect to his/her classmates during Kindergarten field trips [16]. Such identified behavior can help kindergarten teachers to better figure out the talent of a child, otherwise would have been overlooked easily.

Daily healthcare and well-being: I have also been working on daily well-being and healthcare applications. As a first step to this direction, we proposed Sinabro [12], an opportunistic and unobtrusive mobile electrocardiogram (ECG) monitoring system. The uniqueness of Sinabro is that it monitors users’ ECG *opportunistically* during daily smartphone use. Such unobtrusive monitoring will open up an unprecedented opportunity for pervasive healthcare applications since the obtrusiveness was a core reason why daily ECG monitoring has still not become reality. It will enable the daily detection and prevention of heart problems and also allow inferences on stress, emotion, and even caffeine consumption and sleep quality. We conducted basic feasibility studies on this approach, and build a prototype ECG sensor and smartphone middleware to support diverse ECG-driven contexts. We are further extending Sinabro to make the system more practically usable. Toward this daily healthcare direction, we have also been building a food consumption monitoring system using various sensors on wearable smartwatch [13]. This can serve a cornerstone to address prevalent obesity problem.

Persuasive games: I have been working on platforms to support design and implementation persuasive games with real-life context inputs to help people change their undesirable lifestyle through gaming [17]. Especially, I collaborated to develop social exercise games and platform, namely ExerLink [15]. With heterogeneous sensor-enabled exercise devices, ExerLink effectively convert an isolated and boring exercise to an exciting and enjoyable social activity, which could potentially persuade people who gave up individual exercises like treadmill running and jump roping due to their boring nature.

Besides the above-mentioned areas, I believe a huge research opportunity is still open to design creative sensing applications in many application domains. I am continuously working on designing and building new applications that could help people in their daily lives.

Enabling Human-Centric Contexts: Beyond Location and Activity

One of my key research directions is to expand context-awareness towards *high-level human-centric contexts* (intention, engagement, emotion, attention, fatigue, anxiety, depression, distractibility, mindfulness, etc.) beyond currently available physical contexts such as location and activity [6]. Our experiences show that context-aware applications are forced to infer useful human-centric contexts from low-level inputs due to lack of system support, and they usually do a bad job of high level inference yet. Instead, it would be better if these low level inputs were already combined (with online profile data and other data sources), in smart ways, to provide applications with user states (such as intention or emotion) that are more useful for intervention purposes. For example, a future context-aware advertisement will not only target users based on their location and activity, but also more precisely target people who has buying intention to purchase certain products. Also, discount coupons can be sent out when the user is bored or has free attentive capacity.

Moreover, a flexible combination of real-time and historical data should be available for querying as this allows the discovery of deeper insights. For example to send out effective coffee shop coupons, the application will need a historic understanding of the user's coffee drinking behavior. For example, to figure out if she even likes coffee and if she does, what time(s) she usually drinks coffee, what brands she likes, etc. The real-time capability will then be used to adjust the historical preferences to match the current situation. For example, if she has just passed by three coffee shops that she normally stops in, this might be a bad time to suggest having a coffee. To enable human-centric context-awareness and analytics, I am taking the initial steps to build a new mobile service, which can be used by applications to accurately, efficiently, and effectively discover the human-centric contexts of smartphone users.

Large-scale Mobile Testbed and Urban Computing

Another thread of my vision is to enable advanced urban computing services by harnessing vast contextual information collected from the whole fleet of sensing devices deployed over people, buildings, and other infrastructure. Such services will provide fine-granule and real-time understanding on a highly complex and dynamic city, which has hardly been explored yet. Furthermore, they will significantly enrich urban life experiences of citizens and supports highly interactive operations of city services, e.g., a highly flexible transportation scheduling and a crowd-aware public library service.

From a different perspective, one of the biggest hurdles in mobile computing research community is the inability to test novel mobile applications, services, or usage patterns in scale with real users. This is particularly challenging for life-immersive sensing services; while a large body of recent research has shown the ability to use mobile sensing to infer a rich set of context, we have little understanding of the challenges of deploying such advanced applications at-scale, in the real world. Usually, researchers are limited to controlled user studies, often with a small number of users or restricted to specific campus/office environments, that attempt to capture reality as best as they can. Empirical researchers understand that such controlled user studies do not provide the accuracy and "proof" that only real-world deployments can.

Real-world Deployment and Testing at Scale. To overcome these limitations and enable the vision of urban computing, we have been building the LiveLabs Urban LifeStyle Innovation Platform, commonly referred to as LiveLabs [10]. LiveLabs seeks to turn four real-world public spaces in Singapore into a spatially-distributed testbed, where mobile applications, strategies and interventions can be tested on real people, via their own mobile devices, while they are engaged in regular lifestyle-driven activities. As an interesting case to show the LiveLabs capability, we are talking initiatives to investigate indoor localization problem in a lot more practical settings (e.g., in crowded environments, with heterogeneous phones of real users), which has been remained as a practically challenging problem even after a decade of active research.

Collective Analytics for Urban Computing Services: LiveLabs architecture involves active data collection from a large participant pool. We leverage such various sensor data from a large participant pool to perform various real-time collective analytics. As a first step, we designed and built two systems: (1) GruMon [8] to detect groups and their rich contexts, and (2) QueueVadis [9] to detect queuing behavior of people in dense urban places.

GruMon addresses the performance criteria of precise group detection at low latencies by overcoming two critical challenges of practical urban spaces, namely (a) the high density of crowds, and (b) the imprecise location information available indoors. Using a host of novel features extracted from commodity smartphone sensors, GruMon can detect over 80% of the groups, with 97% precision, using 10 minutes latency windows, even in venues with limited or no location information. Moreover, in venues where location information is available, GruMon improves the detection latency by up to 20% using semantic information and additional sensors to complement traditional spatio-temporal clustering approaches. We evaluated GruMon on data collected from shopping episodes from 154 real participants, in two large shopping complexes in Korea and Singapore. We also tested GruMon on a large-scale dataset from an international airport (containing ≈37K+ unlabelled location traces per day), and showed both GruMon's potential performance at scale and various scalability challenges for real-world dense environment deployments. Now, we have been applied the system on our university to obtain various insights over students' physical group activities.

QueueVadis addresses the problem of estimating, in real-time, the properties of queues at commonplace urban locations, such as coffee shops, taxi stands and movie theaters. Abjuring the use of any queuing-specific infrastructure sensors, QueueVadis uses participatory mobile sensing to detect both (i) the individual-level queuing episodes for any arbitrarily-shaped queue and (ii) the aggregate-level queue properties such as expected wait or service times via appropriate statistical aggregation of multi-person data. Moreover, for venues where multiple queues are too close to be separated via location estimates, QueueVadis also uses a novel disambiguation technique to separate users into multiple distinct queues. User studies, performed with 138 cumulative total users observed at 23 different real-world queues across Singapore and Japan, show that QueueVadis is able to (a) identify all individual queuing episodes, (b) predict service and wait times fairly accurately (with median estimation errors in the 10%-20% range), independent of the queue's shape, (c) separate users in multiple proximate queues with close to 80% accuracy and (d) provide reasonable estimates when the participation rate (the fraction of QueueVadis-equipped people in the queue) is modest.

Large-scale location data processing. I have also taking data-driven approach to build high-performance and scalable monitoring platforms for large-scale urban-context-aware services. In this thread, I developed a scalable monitoring framework of *Activity-Travel Patterns (ATPs)* [4] of city residents. The concept of ATP has been recognized as one of the most useful tools to understand city residents' everyday lives in domains such as urban planning, geography, and transportation. However, studies in those domains have mainly focused on developing logical models of ATP from offline analysis of survey data, requiring huge amounts of time, labor and expense. My work is the first, interdisciplinary attempt to introduce automated ATP monitoring in a large-scale city environment. To enable ATP monitoring, we first develop a new computational model of ATP; it newly abstracts the urban activities with regard to the spatio-temporal occupancy of urban resources by the residents. The space-time abstraction provides a basis for a macroscopic urban activity model, upon which a rich composition of diverse patterns can be enabled. More important, I developed highly efficient and scalable processing techniques for ATP monitoring; they process multiple queries in a shared manner based on the new observation that many urban places are commonly interested by multiple queries. Based on the system, I also worked on building a practical application, *AdNext* [18], which provides mobile advertisement based on users' store visit patterns in a large commercial complex. This has been deployed and tested in the COEX mall, a largest commercial complex in Korea.

Summary

I have so far grounded my research in mobile and pervasive computing fields and published important results in premier conferences and journals in the fields (*ACM MobiSys*, *ACM SenSys*, *ACM UbiComp*, *ACM CSCW*, *IEEE PerCom*, *Pervasive*, *IEEE TMC,* and *Communications of the ACM*). I believe that the concepts, technologies, and systems introduced in my works have important implications, which will serve as guiding references for subsequent progress in future mobile platforms. Especially, I think the new concept of the proactive services and the context monitoring platform will be integrated as a core functionality of a future mobile OS, which is beyond what the state-of-the-art mobile platforms such as Apple iOS and Google Android could do. Also, I believe my research direction will be a core of computer science research since computing is continuously expanding its horizon and making fundamental changes to cover highly diverse life situations, new applications and workloads.

In a bigger perspective, my overall research vision is to explore essential needs of people, abstract the requirements build creative systems and tackle first-of-the-kind interdisciplinary technical challenges. So far, I have mainly worked in the fields of mobile and pervasive computing, but I am open to collaborate with people in the field of HCI and core systems like OS and networking. Through active collaboration, I would like to add new perspectives to my on-going research and extend my viewpoint to new significant research problems.

Selected Publications

1. **Youngki Lee**, Younghyun Ju, Chulhong Min, Seungwoo Kang, Inseok Hwang, and Junehwa Song, "CoMon: Cooperative Ambience Monitoring Platform with Continuity and Benefit Awareness", In *Proc. of ACM International Conference on Mobile Systems, Applications and Services (MobiSys), June 2012*
2. **Youngki Lee**, Chulhong Min, Chanyou Hwang, Jaeung Lee, Inseok Hwang, Younghyun Ju, Chungkuk Yoo, Miri Moon, Uichin Lee, and Junehwa Song, "SocioPhone: Everyday Face-to-Face Interaction Monitoring Platform Using Multi-Phone Sensor Fusion", in *Proc. of ACM International Conference on Mobile Systems, Applications and Services (MobiSys)*, June 2013
3. **Youngki Lee**, S.S. Iyengar, Chulhong Min, Younghyun Ju, Taiwoo Park, Seungwoo Kang, Jinwon Lee, Yunseok Rhee, and Junehwa Song, "MobiCon: Mobile Context Monitoring Platform", *in Communications of the ACM (CACM), March 2012*
4. **Youngki Lee**, Chulhong Min, Younghyun Ju, Seungwoo Kang, Yunseok Rhee, and Junehwa Song, "Orchestrator: An Active Resource Orchestration Framework for PAN-scale Sensor-rich Mobile Environments", *IEEE Transactions on Mobile Computing (TMC), March 2014*
5. **Youngki Lee**, SangJeong Lee, Byoungjip Kim, Jungwoo Kim, Yunseok Rhee, and Junehwa Song, "Scalable Activity-Travel Pattern Monitoring Framework for Large-scale City Environment", *in IEEE Transactions on Mobile Computing (TMC), Volume 11, Issue 4, April 2012*
6. **Youngki Lee** and Rajesh Krishna Balan, "The case for Human-Centric Personal Analytics", in *Proc. of MobiSys workshop on Physical Analytics, June 2014*
7. Younghyun Ju, **Youngki Lee**, Jihyun Yu, Chulhong Min, Insik Shin, Junehwa Song, "SymPhoney: Coordinated Dataflow Execution Engine for Concurrent Mobile Sensing Applications", in *Proc. of ACM Conference on Embedded Networked Sensor Systems (SenSys), November 2012*
8. Rijurekha Sen, **Youngki Lee**, Kasthuri Jayarajah, Archan Misra, and Rajesh Krishna Balan, "GruMon: Fast and Accurate Group Monitoring for Heterogeneous Urban Spaces", in *Proc. of ACM Conference on Embedded Networked Sensor Systems (SenSys), November 2014*.
9. Tadashi Okoshi, Vu Lu, Chetna Vig, **Youngki Lee**, Rajesh Krishna Balan, and Archan Misra, QueueVadis: Queuing Analytics Using Smartphones, accepted to *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), 2015*
10. Rajesh Krishna Balan, Archan Misra, and **Youngki Lee**, "LiveLabs: Building An In-Situ Real-Time Mobile Experimentation Testbed", in *Proc. of ACM International Workshop on Mobile Computing Systems and Applications (HotMobile), February 2014*
11. Inseok Hwang, Chungkuk Yoo, Chanyou Hwang, Dongsun Yim, **Youngki Lee**, Chulhong Min, John Kim, and Junehwa Song. "TalkBetter: Family-driven Mobile Intervention Care for Children with

Language Delay". In *Proc. of ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW), February 2014* (Won Best Paper Award)

12. Seungwoo Kang, Sungjun Kwon, Sangwon Seo, Chungkuk Yoo, Kwangsuk Park, Junehwa Song, and **Youngki Lee**, "Sinabro: Opportunistic and Unobtrusive Mobile ECG Monitoring System", in *Proc. of ACM International Workshop on Mobile Computing Systems and Applications (HotMobile), February 2014*
13. Sougata Sen, Vigneshwaran Subbaraju, Archan Misra, Rajesh Balan, and **Youngki Lee**, "The Case for Smartwatch-based Diet Monitoring", to *appear in WristSense (PerCom workshop), March 2015*.
14. Seungwoo Kang, Jinwon Lee, Hyukjae Jang, Hyonik Lee, **Youngki Lee**, Souneil Park, Taiwoo Park, and Junehwa Song, "SeeMon: Scalable and Energy-efficient Context Monitoring Framework for Sensor-rich Mobile Environments", in *Proc. of International Conference on Mobile Systems, Applications, and Services* (*MobiSys*), *June 2008*
15. Taiwoo Park, Inseok Hwang, Uichin Lee, Sunghoon Ivan Lee, Chungkuk Yoo, **Youngki Lee**, Hyukjae Jang, Sungwon Peter Choe, Souneil Park, and Junehwa Song, "ExerLink: Enabling Pervasive Social Exergames with Heterogeneous Exercise Devices", in *Proc. of International Conference on Mobile Systems, Applications, and Services* (*MobiSys*), *June 2012*
16. Inseok Hwang, HyukJae Jang, Taiwoo Park, Aram Choi, **Youngki Lee**, Chanyou Hwang, Yanggui Choi, Lama Nachman, and Junehwa Song, "Leveraging Children's Behavioral Distribution and Singularities in New Interactive Environments: Study in Kindergarten Field Trips", in *Proc. of International Conference on Pervasive Computing (Pervasive)*, June 2012
17. Inseok Hwang, **Youngki Lee**, Taiwoo Park, and Junehwa Song, "Toward a Mobile Platform for Pervasive Games" in *Proc. of SIGCOMM workshop on Mobile Gaming (MobiGame), August 2012*
18. Byoungjip Kim, Jin-Young Ha, SangJeong Lee, Seungwoo Kang, **Youngki Lee**, Yunseok Rhee, Lama Nachman, and Junehwa Song, "AdNext: A Visit-Pattern-Aware Mobile Advertising System for Urban Commercial Complexes", in *Proc. of ACM International Workshop on Mobile Computing Systems and Applications (HotMobile), March 2011*