source:www.nsf.gov

CPS Cloud Computing Report Writing Workshop
Virginia Tech, 900 North Glebe Road, Arlington, VA

Flexible, Resilient and Rigorous CPS at scale

A Community Report
July, 2013

CC4CPS Report - TBD

TBD

December 27, 2013

Contents

1 Executive Summary

1.1 Making CPS flexible, resilient and rigourous at scale

1.2 Foundational Research Needs for CPS clouds at scale

- timeliness, multi-physics aware clouds (networking included)
- formal and mathematical models analysis, composition, optimization, verification (addresses security , integrity)
- resource allocation, scheduling, control programming abstractions, methods, patterns, languages
- coordination and negotiation models, mechanisms

1.3 Experimental research needs

- Discovering interactions failure modes, assumptions
- validation of theory and its assumptions.
- evaluation calibration benchmarking of infrastructure
- parameterized inter-dependent assume guarantee service level agreements for large scale CPS scenarios
- Risk analysis forensics, attribution for large scale CPS
- Education

1.4 Research Infrastructure needs

- monitoring large scale CPS in the real world
- Open experimentation platform/test beds
- Education is a strong driver for research infrastructure.
- Frameworks/platforms for architecture experimentation
- Safe Sandboxes for attack/fault injection in large scale CPS

1.5 The way forward towards cloud for large scale CPS

1.6 summary

2 Introduction

Large scale cyber-physical systems (CPS) in areas of vital national interest such as transportation, health care, energy, agriculture, defense, homeland security, and manufacturing, are becoming increasingly interconnected and interdependent. Large scale CPS are unique in their need to combine rigorous control over timing and physical properties as well as functional ones, while operating dynamically, reliably and affordably over significant scales of distribution, resource consumption, and utilization. As large scale CPS continue to evolve, they will impose significant and novel requirements for a new kind of cloud computing that simply cannot be met by current technologies.

However, current research on next generation networking, cloud computing, and other potentially relevant technologies, is not addressing the specific challenges posed by large scale CPS. In particular, the combination of (1) geographic distribution, (2) dynamic demand for resources, and (3) rigorous behavioral requirements spanning diverse temporal and physical scales, frames a unique and compelling constellation of research challenges that must be pursued together to realize new foundations for cloud computing that can meet the needs of large scale CPS.

To pursue these challenges, new research is needed to establish required real-time computing, communication, and control foundations rigorously at scale. In addition, experimental research is needed to apply these foundations to real-world large-scale CPS challenge problems. To support both foundational and experimental research, a new generation of research infrastructure and testbeds for experimentation and evaluation relevant to real-world large scale CPS applications also needs to be designed, developed, and evaluated. This community report identifies challenges, opportunities, and benefits for this research and research infrastructure, and for the large scale CPS they target.

3 Foundational Research

3.1 Intro

3.2 Future Needs

As of the time of this writing the term cloud taht is in wide use will be referred to as a *Classic Cloud*. NIST SP800-145 defines Cloud Computing as "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

When cloud and CPS are merged together, two concepts are revealed; 1) classic clouds supporting CPS, and 2) CPS built as a cloud. The first concept uses classic clouds to provide cyber resources (processing and data) to a CPS system. The second concept provides the components of a CPS (sensors, actuators, data, processing, and complex combinations of these) as cloud services (ie. having the characteristics of classic cloud computing). In the NISTIR 7951 NIST defines a Cyber-Physical Cloud Comupting framework as "a system environment that can rapidly build, modify and provision cyber-physical systems composed of a set of cloud computing based sensor, processing, control, and data services." While this is a good starting poitn, many questions still remain. Such as; What are the fundamental properties of a CPS Cloud? How is a CPS Cloud differentiated from a Classic Cloud. Are there applications that are CPS-specific, those that are Cloud-Specific, and those that are CPS-Cloud Specific? How do CPS APIs relate to Classic Cloud APIs. What functional and non-functional requirements do CPS Clouds have?

CPS Cloud Research needs to address several system facets that are essential for CPS Cloud including 1) Timeliness, 2) Locality. 3) Uncertianty, 4) Reliability 5) Security/Privacy

1. Timeliness - How to support different timing requirements of the different applications within the CPS cloud? How to support a variety of time scales? How to support hard real-time vs

soft realtime? What are the timing requirements of specific applications? What is the time scale needed? Does a CPS Cloud require hard real-time or soft real-time performance?

2. Locality - Since CPS Cloud interacts with the physical world the knoweldge of where in the world a particular sensor or acturator is located is important. When trying to build an application using CPCS Cloud

3. Uncertainty - Again, since CPS Cloud interacts with the physical world the concept of uncertaiinties and tolerances is crucial. As all measurements are approximations of properties that exist in the real world, the uncertaintiy of the approximation is essential in understanding what the measurement means.

4. Reliability is the ability of a system to perform its functions in routine and unexpected circumstances. For CPS clouds to work the appropriate level of relaibility must be proved for each application. For instance, should a CPS cloud be a dedicated private cloud with rigid QoS guarantees or is a public cloud with relaxed QoS enough? By contrast, a big issue that emerged in the workshop was the term ”constrained elasticity”. The contrast with the current view of CPS applications, which are hard real-time over owned resource networks, is that use of a CPS Cloud may require a relaxation of constraints. How does multitenancy, or multiple users in the cloud effect QoS?

5. Security and Privacy are current hot topics within the Classic Cloud environment. While the CPS Cloud has an obvious set of parallels to the Classic Cloud in confidentiality, integrity, availability, and privacy, it refines the models of trust. In a CPS Cloud, there exists another dimension of attack vectors in the physical system; the input to sensors can be spoofed. The actions taken by a CPS CLodu application also may have serious implications. As CPS is used for critical infrastructure and other areas where the effects of an attack may cause injury or death, the level of security needed may be great. Along with the physical aspect of the CPS cloud comes a drawback; there is an inherent information flow from the cyber decision making the results in observable physical actions. Thus, there is a need to isolate decisions from actions within the confidentiality and privacy context. Security, however, is in general not compositional so there is a need to understand capabilities and limitations of multiple layer cloud architectures. Traceability and attestation of actions emerges as a fundamental research need; not only why did an action occur, but what caused it. As such, security policies for the CPS Cloud need to be developed that go beyond just cyber security, but are truely cyber-physical policies. These policies must be aware of the emerging area of cyber-physical threats in which both the physical system and the cyber system can be attack vectors to the entire CPS Cloud.

Following are some more notes on CPS Clouds

- It is important to note that the requirements of a particular CPS Cloud applications within each of the five facets is different. While one application might require hard real time with a very small time scale, another application might require soft real time with a large time scale, likewise locality, uncertainty, reliability and security requriements all vary according to application. The challenge with the CPS Cloud infrastructure is having system compenents that can provide their capabilities in a well characterized manner so the application manager can chose the proper combination of compoents and then ensure those capabilities are provided for the particular application. This elasticity is fundemental to CPS Clouds. How to implement it is not well understood.

- The question of constrained elasticity uncovered the need for new types of service level agreements (SLAs) that go beyond just time and account for CPS aspects such as state estimation. The CPS Cloud needs to provide mutual reflection; (the application needs to know what is happening in the cloud and the cloud needs to know what the applications are doing so both can benefit in terms of managing resources).

- The relationship between the abstractions in the cloud and the ground truth formed by the physical elements of a CPS uncovered needs of multi-physics aware systems, a layer-independent set of analytical metrics and a consistent abstraction (such as time or application performance). A challenge is to find monitoring facilities that can understand and relate the system at and through each layer of abstraction; do lower layers have more stringent requirements than high layers, or is the situation reversed?

- Cloud services and layers, themselves, came under scrutiny. In a CPS Cloud, there is a need to extend the notion of resources to beyond just Cyber resources, but understand that these can also be physical resources (such as energy storage devices in a power system that allow more decision making flexibility in the cyber world). These physical devices are connected to the cyber devices through sensors ad actuators. Likewise, there is a need to guarantee the integrity of the response but it is not clear that this is significantly more important in a CPS Cloud than for a Classic Cloud. For CPS applications, the CPS Cloud needs to provide end-to-end security mechanisms, where the endpoints can, conceivably be between any layers.

- The tension between security and performance impacts any system. In the CPS Cloud, different parts of the cloud provide differing levels of performance and security. There is a fundamental need to understand and model the composition. This tension is already manifested in HW/SW virtualization architecture for the Classic Cloud, but the real-time aspects of CPS application with different real-time and security requirements increases the complexity of the problem.

- Improved resilience forms a potential benefit of the CPS Cloud. There is a need to determine if cloud scalability to support monitoring, diagnosis, and adaptation enhances system survivability or is a detriment to it? Does leveraging data and resources across multiple private clouds improve resiliency?

- Testbeds form a common call to action. CPS Clouds are not well understood, and an MRI-like program of investigation and study that involves multiple stakeholders is needed to incrementally define and understand the CPS Cloud.

Stuff I don't know what to do with or is just standard comments about everything

- (includes security) Formal and mathematical Models, analysis, composition, optimization, verification

- Resource allocation, scheduling, and control (meaning scalable control models and methods)

- Programming abstractions , methods, patterns, languages

- Coordination and negotiation models, protocols, mechanisms

- Leverage cloud-based social media applications (e.g., Twitter, FaceBook) for CPS **(I don't understand this)**

- Interoperability Standardization of Interfaces
- need to create better abstractions, and tools for designing, analyzing and measuring future systems
- understand pervasiveness, people-centric sensing, Big Data stream processing, informatics
- understand how CPS data is time indexed, inferences that are made affect semantic complexity, how to reason about them, e.g., sequencing, quantitative time, etc; address relationship of Big Data to CPS Cloud
- tie things to other disciplines too (beyond CISE)
- Need to understand the overlap, if any, with the XPS solicitation (foundations, scalable distributed arch, cross layer approaches – rethinking the stack, domain-specific design i.e., exploit domain knowledge for programmability and performance; "Architecture as Infrastructure"
- Do we need an Open Cloud Platform with pluggable apps and support for QoS guarantees

3.3 Barriers

- Insufficiency of cloud SLAs
  - cloud elasticity does not account for physical characteristics of the CPS
  - coarse grained with no hard bounds on latency or availability
  - time is fundamental to CPS, and not considered here
  - state estimation is not considered
  - does not consider multi-tenency, and aggregate elastic requirements
- CPS load can be very bursty
  - cloud economics based on aggregate load
  - this is a challenge to the applicability of clouds to CPS
- CPSes are always-on, high-criticality; cloud is mostly available
  - how do we reconcile this difficulty
- Mismatch between CPS control frequency and cloud/network latency
  - a fundamental physical barrier
  - limits what synergy is possible
  - network congestion impacts end-to-end service
- Criticality of CPSes is not matched by provisions of the cloud
  - certification of airplanes vs. cloud contracts
- Mismatch between CPS programming abstractions/methods and distributed system/cloud programming abstractions

  - Engineers don't have good abstractions for programming distributed systems
  - CPS focus on control systems and hard real-time
  - cloud focused on map-reduce and service-level programming

- Security/privacy in the cloud are defined based on current computational assumptions
  - CPSes

- Frequency of events in global CPSes much higher than current clouds

- Physical resources are managed and modeled completely separate from the computational cloud
  - makes co-management difficult

- Current cloud architectures impose unavoidable latencies
  - EC2 with Xen, or app engine contain software that is not CPS-optimized
  - might prevent fine-granularity guarantees being made to CPSes

- Cloud processing of requests, aggregation of data is not geography- and physics-aware
  - how can CPS-aware processing be done, if it the cloud doesn't understand which CPS directly physically interact?

- Safety is not considered as a fundamental computational aspect of current clouds

- Inadequate mainstream programming models
  - Complicated & obtrusive APIs
  - Doesnt use hardware effectively & scalably

- Inadequate knowledge of real-time, concurrency, & networking e.g., high probability of race conditions, deadlocks, priority inversion, & missed deadlines

- Inadequate mechanisms to migrate seamlessly from multi- to distributed-core environments

- Inadequate quality-of-service (QoS) support at scale e.g., lack of control over key QoS impacting resource usage & end-to-end data deliver semantics

- Elasticity of clouds is necessary but not sufficient

3.4 Enabling technologies

Consider the following examples.

- air traffic control

- smart grids

- earthquake monitoring

- medical robotics
- medical devices
- Automotive

3.5 Research Challenges

- Privacy is an essential issue whenever people are affected by or involved with systems - with cloud CPS the challenge is that information that is necessary may not be appropriate to release without approval, obfuscation, or other transformation that can integrate it *correctly* within the CPS cloud.
- Privacy is context sensitive, a matter of policy, and is piggybacked in non-trivial ways atop most other info.
- Notion of informed consent needs to be generalized based on the semantic integration that occurs within CPS clouds - no way to enumerate all the cases, but principles for how data will or will not be used need to be articulated as policy and enforced via appropriate mechanisms. What needs to be known where and when - how much do boundaries of trust play into this?
- Placeholder for the paragraph(s) that will summarize a vision from future needs sub-section above and relate it to the research challenge areas.
- Placeholder for the paragraph(s) that will summarize the enabling technologies from the subsection above and relate them to the research challenge areas.
- Placeholder for a potential classification for the research challenge areas (will be done after mining of the presentations/notes and going through the following bullet items).
- Fundamental challenges with enabling Cloud for CPS: timeliness and multiphysics aware clouds
- Funamental challenges with resilient C-P cloud design and operation: formal models, analysis, composition, optimization, verification; resource allocation, scheduling, control; programming abstractions, methods, pattern languages; coordination and negotiation models, protocols, mechanisms; identity, role, access, interdiction.
- Multi-scale control in time and space. Multiple levels of control loops. Control itself needs to evolve in a flexible way; it might be the case that control algo is designed with a given set of resources made available. But what happens when resources are elastic - goes up or down, then how control evolves? and can it be done in a geographically distributed fashion? Huge set of interlocking issues.
- Heterogeneous data (sources and types), resources, and domains.
- Programming model/abstraction. Need to have map-reduce like operators but with CPS descriptors. This was a topic suggested on the basis of differences in programming models. How do we map current processing that takes place on physical elements to now occur within the cloud? i.e., what does it mean to adapt to MapReduce and such programming abstractions. Note we must be aware of the jitter and variance in parallel cloud activities.

Programming against the cloud environment (distributed) is unfamiliar to application engineers. Self-awareness within the cloud is also a new algorithmic technique. That is some sort of mutual reflection (application needs to know what is happening in the cloud, cloud needs to know what the applns are doing so both can benefit in terms of managing resources, managing properties like deadlines, etc.). Analytical benchmarking metric/consistent abstraction (time, application performance?) independent of layers.

- Context-based semantic tagging of data.
- Elasticity. Constrained elasticity (i.e., disallowing arbitrary elastic demands otherwise others might suffer paticularly because the CPS cloud will be multitenant); decisions are parametrized and so analysis is easier.
- Elasticitydemand response, in a more general sense. Technical complexities and safety-criticality of CPS will make it daunting for public Cloud services to be provided.
- SLAs are risk driven need derivatives on safety, criticality and cost. How to audit SLAs. current ones are insufficient virtual contract between user and provider not just limited to time but things like state estimation. The cloud itself may be cyber-physical; maybe there is a real-time backbone shared by many application, but then above that there are a varying set of controls, things happening at different time scales, etc. The notion of virtual contract with these additional things now makes the abstractions and interfaces. We need compositional assurances (when properties are composed and patched). We might need symmetric properties here i.e., we can do both composition and decomposition. What granularity can we provide? SLA auditing (particularly if they are virtual).
- In security, what is the adversary, where are the boundaries of trust? Can we distinguish attacks? A denial of service attack from soft failures - the underlying hypervisor may delay execution within the cloud? The physical system may give feedback. A cloud must interact across the network, so delays due to attacks or congestion impact the end-to-end service. Establishing a model of trust.
- Humans in the loop.
- Where the application resides (in the cloud or locally) is a not “one size fits all”
- Influencing CPS applications. Define patterns, frameworks, and guidance for public and private clouds. Distributed control algorithms adaptable to weak QoS guarantees. Application architectures for cloud-based CPS components.
- Influencing cloud infrastructures. Identify and characterize CPS QoS requirements. Identify infrastructure abstractions and service models. Multi-property QoS optimization framework. Application profiles (resources and services).

4 Experimental Research

4.1 Intro

4.2 Future Needs

4.3 Barriers

4.4 Enabling technologies

4.5 Research Challenges

- Discovering interactions failure modes, assumptions
- validation of theory and its assumptions.
- evaluation calibration benchmarking of infrastructure
- parameterized inter-dependent assume guarantee service level agreements for large scale CPS scenarios
- Risk analysis forensics, attribution for large scale CPS
- Push-pull model with different levels of QoS
- hardware software architectures that enable security and performance
- Configuration and Deployment Modeling Paradigms
- How to describe and configure SLA
- Hoe to audit and monitor SLA
  - Synchronous two-way pushes
  - Asynchronous pushes and pulls
  - Caching interfaces
- "Moving clouds"
  - Up and down the hierarchy
  - Based on location and time (both time-driven & event-driven)
- Spatial and temporal fencing
- Adaptation interfaces
- How to monitor and manage faults.
- Failure recovery interfaces
- Time-varying summarization and digests
- Models for physical dynamics

- How to trust the originator?
- Mechanisms to manage end-system resources, e.g., CPU scheduling, file systems, memory management, & IPC.
- Native OS mechanisms to create reusable network programming components.
- Higher-level distributed programming models whose reusable APIs & components automate & extend native OS capabilities.
- Higher-level domain-independent services that focus on programming “business logic”
  - Behavioral logic
  - Rules of safe operation
  - Control logic
  - Abstractions and semantics
  - Decision procedures
- Tailored to the requirements of particular domains, such as SCADA, C4ISR, avionics, vectronics, air traffic management, aerospace, etc.
- Build on standard services like DDS
- Dynamic semantic gap between what the infrastructure provides and what the CPS needs?

5 Research and Education Infrastructure

5.1 Intro

5.2 Future Needs

Virtual testbeds
CPS cloud testbeds that can be duplicated

5.3 Barriers

5.4 Enabling Technologies and Proofs of Concept

- GENI `http://www.geni.net/` - also discuss the Real-Time GENI report. GENI enables programmable network infrastructure distributed across US.
- EMULAB - An emulated experiment allows you to specify an arbitrary network topology, giving you a controllable, predictable, and repeatable environment, including PC nodes on which you have full ”root” access, running an operating system of your choice.
- RT-XEN (real-time virtualization.). - This is an extension of XEN hypervisor. The Xen Hypervisor provides the ability to virtualize the computing nodes at a low level.
- experimental Arm -virtualization in XEN. - Addition to Xen. This will target arm, which is the most common embedded platform.

- OpenStack - it is an open source cloud computing platform. It allows to create and manage small to medium cloud installations. it does not support geographically distributed installation.
- PlanetLab `http://www.planet-lab.org` – is an open platform to conduct networking and related research and experimentation at planetary scale.
- cloud-based social media applications (e.g., Twitter, FaceBook)
- CPS-VO `http://www.cps-vo.org` – is a NSF-funded web portal for fostering collaborations among CPS professions in academia, government and industry.
- SPRUCE `https://www.sprucecommunity.org` – is a web portal funded by the Air Force Research Laboratory that offers a collaborative environment to demonstrate, evaluate, and document the ability of tools, methods, techniques, and technologies to demonstrate software producibility advances for software-intensive systems.
- NEES `http://nees.org/` – The Network for Earthquake Engineering Simulation (NEES) is a NSF-funded organization whose focus is on creating a community of researchers and practicing engineers who are involved in novel research aimed at reducing the impact of seismic disasters.

5.5 Infrastructure Challenges

CPSforge - integrated VO and experimental testbed, submit experiment specification which is checked and if feasible scheduled provisioned and run. Results are archived along with specification as a CPSforge entry, which is noted in the publication of the results - others can resubmit (possibly with tweaks) to validated, confirm, compare, extend, etc.

5.5.1 Infrastructure Challenges for Foundational Research

[Cross-reference grand challenges defined in foundational research]

5.5.2 Infrastructure Challenges for Experimental Research

[Cross-reference grand challenges defined in experimental research]

5.5.3 Infrastructure Challenges for Education

Distributed CPS education: e.g., K-12 STEM classes nationwide submit experiments to a competition, from which selected experiments are run on a dedicated testbed with classrooms linked in remotely via skype.

6 Illustrative Scenarios

To illustrate the potential for cyber-physical system clouds (CPS clouds) to revolutionize different applications, we describe scenarios in each of several domains of vital iterest: medical cyber-physical systems for health care; cyber-physical systems for energy generation, storage, transmission and management; cyber-physical transportation systems; cyber-physical manufacturing systems; and cyber-physical civil infrastructure. We conclude with scenarios that cross-cut those domains, in the context of disaster management and recovery.

6.1 Medical CPS

In the domain of medical cyber-physical systems (medical CPS), CPS clouds have significant potential for improving how chronic conditions are managed in modern healthcare settings. For example, type I diabetes is a chronic condition with an onset that is often early in life, which requires long-term management of blood glucose levels through monitoring, diet, and regular doses of insulin according to established standards of care.

Using cyber-physical clouds to connect devices that can monitor physiological factors and administer medications automatically may offer entirely new ways of managing such diseases, through integration of (1) closed-loop control within the devices themselves, (2) on-line collection, dissemination, and storage of data with appropriate guarantees of timely availability and transmission of such data, and (3) on-line analytics that can detect or discover essential patterns in the data to improve management of the disease.

A key issue for such integration, which further motivates the use of cyber-physical clouds in medical CPS, is the significant variation in clinical, physiological, and compliance factors that impact the efficacy of treatment. For example, multiple protocols are available for how and when to administer doses of insulin and how much insulin to administer in each dose, relative to the timing and content of a patient's meals.

Individual variation in physiological response (e.g., in terms of insulin sensitivity, baseline metabolic activity, and other factors) may suggest that particular protocols may be efficacious, but even then must be adjusted according to parameters measured for each individual. Furthermore, control over the timing of meals and insulin doses may vary significantly between clinical settings, which can have a significant effect on patients' glucose-insulin dynamics - for example in a hospital the timing may be closely regulated but at home such precision is unusual. Finally, patients' non-compliance with designated protocols is widespread, in particular during teenage years when significant social, developmental, and other external factors have a major influence on the degree and form of non-compliance in each particular case.

Given such variability, devloping, refining, and applying detailed models of each patient's response to treatment in terms of all the relevant dimensions affecting patient outcomes is essential. For example, treatment plans list objectives that can be tracked within measures of adherence, which are being applied within the context of ongoing research towards eventual evidence based identification of patterns of compliance and classification of compliance-related behaviors.

In this context, cyber-physical clouds offer revolutionary potential to integrate those dimensions effectively, both for each patient individually and across cohorts of patients, to improve care. For example, simply recording and storing time-stamped data for each patient within each medical device, and then propagating streams of data for storage within a CPS cloud offers new opportunities to record data faithfully (and to detect and potentially correct cases where it has not been) and then updating treatment plans more effectively over time according to recorded observations, tuning of protocol parameters, and other adaptive measures.

Such rigorous data collection and the availability of such data both immediately over short intervals and on-demand over longer intervals can motivate and enable research into medical device control system models that can tolerate uncertainty in parameters that are not yet perfectly calibrated, perform on-line learning of model parameters or modes, operate open-loop for certain intervals, and still provide strong guarantees of stability and safety.

Such integration in turn will require that the CPS cloud itself offer guarantees to the control systems about availability and latency of access to data stored in the cloud, capacity for storage, and potential feedback from analytic services that can be used to tune parameters or other features of the control model itself.

For example, hybrid control models that sense glucose levels, receive input data regarding meals, and inject insulin accordingly must be parameterized with mode changes and (potentially stochastic) parameters in order to remain predictive over reasonably short intervals (minutes, which is an appropriate time-scale for glucose-related physiology) and may need to be tuned regularly over longer intervals (hours and days). To achieve this, learning and control methods need to be integrated, and analysis techniques need to be extended to provide appropriate guarantees (e.g., probably-approximately-correct bounds etc.) on the quality of control. Finally, the inherent variability described above may offer novel opportunities to leverage classic explore/exploit trade-offs in learning as it is integrated with stochastic and adaptive control on-line.

Within the cloud, the control model also can be augmented through additional psychological and social sensors that combine to give a more complete picture of factors affecting adherence, etc. Social sensors that can detect different modes of behavior are of particular interest in the context of diabetes management, since social factors such as stress and inter-personal interactions can have significant effects on patients' decision making, and since transitions from clinical normalcy into decline may be abrupt as well as chronic.

In addition, details of the control model, such as how often it needs certain inputs to close the loop, and at what time scales other inputs can be integrated effectively, can serve to define timing requirements within the cloud itself - i.e., what combinations of information and its timing must be available to maintain control stability and safety, and what guarantees of privacy and other constraints must be enforced by the cloud. This in turn allows (and benefits from) significant flexibility for elastic multi-tenant allocation of cloud resources to diverse but potentially interoperable feedback loops needed by different patients whose data streams are managed in the cloud.

In addition to such "small data" feedback loops, integrating "big data" analytics within clouds for medical CPS, across data sets for multiple patients and within the entire data set for each patient, offers further potential advances in improved care. Mining of correllations between clinical, social, physiological, and other factors with different degrees and forms of adherence to protocols, may offer new insights into efficacy of treatment under different conditions or in different contexts, in practice. This in turn will require attention to standardized or inter-operable data collection, integration, and storage protocols, so that individual streams of data can be aggregated easily and effectively.

Analyses that can cluster dynamic factors such as sequences and durations of events may offer new insights and potential improvements to the current standards of care. For example, clustering events across shorter time scales (meals and insulin doses), intermediate time scales (sleep and stress), and longer time scales (age and social context) may be essential to modeling diverse time-varying phenomena influencing clinical outcomes.

Analytics that integrate social sensing with other data also offer significant potential for privacy-preserving extraction of modes of behavior and other important parameters beyond those available solely within the context of medical device control. Classification of patients according to those factors may in turn offer opportunities to select treatment protocols more appropriately, as well as to enable switching or adaptation of protocols over the course of treatment.

In summary, medical CPS clouds help both to stratify and integrate learning and decision making, which in turn appears to offer a natural structure for integrating "small data" services for multiple patients within a common cloud. Giving individuals complete and precise closed- and open-loop information about their own individual physiology through a CPS cloud already would be an important advance in the state of the art, to help manage their conditions more effectively. The aggregation of appropriately curated (e.g., secure and restricted or anonymized) individual data streams within a CPS cloud would enable further advances in repurposing data and devices to detect and manage other conditions, observe patterns across cohorts of patients, or improve practices across clinical environments.

6.2 Energy CPS

6.3 Transportation CPS

6.4 Manufacturing CPS

6.5 Civil Infrastructure CPS

Small data appears in this context, e.g., to monitor energy usage throughout a home, a business, etc.

6.6 Scenarios Cross-Cutting CPS Domains

We consider disaster recovery as an examplar of the kind of application that is likely to cross cut scenarios involving other applications of CPS clouds.

6.7 Summary

Control loops and decision loops immediately affecting the system's stability can be given higher priority than loops performing less time-critical or safety-critical functions, but when possible resources would be shifted elastically in a mixed-criticality manner. This kind of resource management approach, which exploits the unique capabilities of CPS clouds, could in turn enable exploration of scientific hypotheses, analyses for process improvement, and other longer term goals with potential to improve the current standard of care even as it is being faithfully followed.

Trends in resource usage and demand profiles could in turn be captured and analyzed to allow advance provisioning of additional resources, allowing CPS clouds to be scaled according to anticipated demand. This is particularly important for the less elastic resources such as individual medical devices, high-availability data caches, and other resources that are more difficult to virtualize (and less fungible in their use), which are likely to be an essential part of CPS clouds (in contrast to traditional computational clouds). Advance planning to allow federation of existing clouds at least temporarily, to increase scalability and elasticity in the face of temporary overloads or unanticipated repurposing of cloud services, also would be enabled by the unique self-monitoring and analysis capabilities offered by CPS clouds.

7 Implications to Stakeholders

8 Strategic Plan for Cloud for CPS

9 Concluding Remarks

A Contributors

B Workshop Working Groups and Writing Teams

C Organizers

D Workshop program

E Acronyms and Definitions