Engineering Micro-intelligence at the Edge of CPCS: Design Guidelines

Roberta Calegari Giovanni Ciatto Enrico Denti Andrea Omicini

{roberta.calegari, giovanni.ciatto, enrico.denti, andrea.omicini}@unibo.it

Dipartimento di Informatica, Scienza e Ingegneria—Università di Bologna

Talk @ *12th International Conference on Internet and Distributed Computing Systems (IDCS 2019)*
June 11, 2019, Napoli, Italy

Next in Line. . .

Context

Era of digital CPS Revolution [Rauch et al., ress, Um, 2019, Waschull et al., ress]

- reality based on smart technologies
- discovering, storing, processing, learning, analysing, and predicting from huge amounts of environment-collected data
- spread of intelligence in distributed system

Edge computing as a new paradigm

[Chen et al., 2019, Pace et al., 2019, Tang et al., 2017]

- distributing intelligence between the Cloud and the Edge
- addressing CPCS issues such as context-awareness and real-time planning

Motivation & Contribution I

- Research work do not provide for a general software engineering standpoint [Ananthanarayanan et al., 2017, Hu et al., 2019, Tang et al., 2017]
- Criteria, guidelines, and methodologies on *how to design and distribute intelligence* are mostly missing [Cicirelli et al., 2019]

$\Downarrow$

Guidelines for the design of distributed intelligence in CPCS

- symbolic and sub-symbolic combined to produce intelligent behaviour
- deal with the diverse requirements of CPCS
- intelligent environment to complement individuals cognitive process

Motivation & Contribution II

- *micro-intelligence* as the *situated intelligence* of Edge computing
- *Situated Logic Programming* as delivering framework for micro-intelligence in CPCS
- *Logic Programming as a Service* [Calegari et al., 2018b] as technological embodiment

Next in Line...

Design Guidelines for Distributing Intelligence

<table>
<tr><th rowspan="2">intelligence</th><th colspan="2">which</th><th colspan="2">where</th><th colspan="3">who</th></tr>
<tr><th>symbolic</th><th>sub-symbolic</th><th>Cloud</th><th>Edge</th><th>individual</th><th>environment</th><th>e-institution</th></tr>
<tr><td>no requirement</td><td colspan="2">✓</td><td>✓</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>real-time req.</td><td></td><td>(✓)</td><td></td><td>✓</td><td colspan="2">✓</td><td></td></tr>
<tr><td>no requirement</td><td colspan="2" rowspan="2">✓</td><td>✓</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>fault-tolerance req.</td><td></td><td>✓</td><td colspan="2">✓</td><td></td></tr>
<tr><td>no requirement</td><td colspan="2">✓</td><td>✓</td><td></td><td></td><td colspan="2">✓</td></tr>
<tr><td>safety critical req.</td><td>✓</td><td></td><td></td><td>✓</td><td>✓</td><td></td><td></td></tr>
<tr><td>big data</td><td></td><td>✓</td><td>✓</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>small data</td><td>✓</td><td></td><td></td><td>✓</td><td colspan="2">✓</td><td></td></tr>
<tr><td>macro scale</td><td></td><td>✓</td><td>✓</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>micro scale</td><td>✓</td><td></td><td></td><td>✓</td><td colspan="2">✓</td><td></td></tr>
<tr><td>commonsense KB</td><td colspan="2" rowspan="2">✓</td><td colspan="2">✓</td><td colspan="2">(✓)</td><td></td></tr>
<tr><td>contextual KB</td><td></td><td>✓</td><td></td><td colspan="2">✓</td></tr>
<tr><td>disembodied</td><td colspan="2" rowspan="2">✓</td><td>(✓)</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>embodied</td><td></td><td>✓</td><td colspan="2">✓</td><td></td></tr>
<tr><td>no requirement</td><td colspan="2">✓</td><td>✓</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>situatedness req.</td><td>(✓)</td><td></td><td></td><td>✓</td><td colspan="2">✓</td><td></td></tr>
</table>

Design choices *(columns)* and design criteria *(rows)*

Focus on. . .

Real-time & Fault-tolerance requirements

intelligence	which		where		who		
	symb	sub-symb	Cloud	Edge	individual	env	institution
no requirement	✓		✓				✓
real-time		(✓)		✓	✓		
no requirement	✓		✓				✓
fault-tolerance				✓	✓		

CPS strict *real-time* requirements

- preferred sub-symbolic approaches at the Edge (individuals/resources)
- e-institutions enforcing norms and monitoring long-term evolution

fault-tolerance not strictly related to intelligence

- enhanced by *decentralised* approaches of Edge computing
- Cloud is intrinsically a fault-tolerant computational environment but also a bottleneck as well as a single point of failure

Focus on. . .

Safety-critical & Big data requirements

intelligence	which		where		who		
	symb	sub-symb	Cloud	Edge	individual	env	institution
no requirement		✓	✓				✓
safety critical	✓			✓	✓		
big data		✓	✓				✓
small data	✓			✓	✓		

CPS usually *safety-critical* system

- preferred symbolic approaches at the Edge (individual responsibility)
- sub-symbolic approaches cannot guarantee no errors
- symbolic techniques can encode strict safety rules to apply no matter what

Big data available?

- sub-symbolic techniques to find relevant patterns (Cloud horsepower)
- exploited by e-institution for long-term planning and predictive adaptation

Focus on. . .

Scale of interest & Nature of knowledge

<table>
<tr><th rowspan="2">intelligence</th><th colspan="2">which</th><th colspan="2">where</th><th colspan="3">who</th></tr>
<tr><th>symb</th><th>sub-symb</th><th>Cloud</th><th>Edge</th><th>individual</th><th>env</th><th>institution</th></tr>
<tr><td>macro scale</td><td></td><td>✓</td><td>✓</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>micro scale</td><td>✓</td><td></td><td></td><td>✓</td><td>✓</td><td></td><td></td></tr>
<tr><td>commonsense</td><td colspan="2" rowspan="2">✓</td><td colspan="2">✓</td><td>(✓)</td><td></td><td></td></tr>
<tr><td>contextual KB</td><td></td><td>✓</td><td></td><td></td><td>✓</td></tr>
</table>

Scale of interest: *micro* vs *macro*

- big data best suits Cloud and institutions to power sub-symbolic techniques
  - → focus on the macro scale of CPCS
- symbolic approaches at the Edge effective for small data analysis
  - → individuals and intelligent environment perform local inference locally

Nature of *knowledge*

- contextual knowledge usually acquired *dynamically* at the Edge of the CPCS
- commonsense knowledge more static → Cloud/Edge depending on criticality

Focus on. . .

Embodied computation & Situatedness requirement

<table>
<tr><th rowspan="2">intelligence</th><th colspan="2">which</th><th colspan="2">where</th><th colspan="3">who</th></tr>
<tr><th>symb</th><th>sub-symb</th><th>Cloud</th><th>Edge</th><th>individual</th><th>env</th><th>institution</th></tr>
<tr><td>disembodied</td><td colspan="2" rowspan="2">✓</td><td>(✓)</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>embodied</td><td></td><td>✓</td><td>✓</td><td></td><td></td></tr>
<tr><td>no requirement</td><td colspan="2">✓</td><td>✓</td><td></td><td></td><td></td><td>✓</td></tr>
<tr><td>situatedness req.</td><td>(✓)</td><td></td><td></td><td>✓</td><td>✓</td><td></td><td></td></tr>
</table>

embodiment of computation: bound to system *physical nature*

- embodied — responsibility of individuals / environment (Edge)
- disembodied — e-institutions for encapsulate all the norms (Cloud)

situatedness requirement

- stronger prevalence of symbolic approaches when situatedness *is required*
- by definition, it may correspond to rare occurrences of events out of the normal system operations

Next in Line. . .

Micro-intelligence in CPCS with LPaaS

→ micro-intelligence approach to reify the guidelines in the LPaaS technology

→ completing the tile of intelligence in CPCS by acting synergically with sub-symbolic techniques

micro-intelligence [Calegari, 2018]

- scattering small chunks of machine intelligence all over a distributed, situated system, capable of enabling the individual intelligence of any kind of devices
- can be interpreted as the *externalised rationality* of cognitive agents (individual), complementing their own in the sense of Clark and Chambers' *active externalism* [Clark and Chalmers, 1998], and under Hutchins' *distributed cognition* perspective [Hollan et al., 2000]

LPaaS I

- suitable reference architecture for delivering micro-intelligence to the CPCS edge [Calegari et al., 2018a]
- designed according SOA paradigm [Erl, 2005]
- multiple logical theories, scattered in the environment, to represent the local, possibly partial/conflicting KB
- inference capabilities are admissible and available over this knowledge (extending the local KB)

LPaaS II

Guidelines reification

scale of interest intelligence at the *micro* scale of the system
→ enabling local (symbolic) reasoning complementing global (sub-symbolic)

Big / small data synergistically exploit symbolic and sub-symbolic approaches
→ ML for general KB generation (scattered in the Edge) then LPaaS for local reification

fault-tolerance microservices LPaaS architecture to achieve fault tolerance

knowledge nature micro-intelligence rooted in the notion of situatedness
→ LPaaS resolution process intrinsically bound to the computational context

embodied computation micro-intelligence rooted in IoT and inner physical nature
→ LPaaS *ad hoc* API for dealing with sensors and streams of data

situatedness LPaaS as an effective embodiment of the situated LP paradigm
→ LP inference process is itself situated in the spatio-temporal context
→ solve operations dedicated to spatio-temporal issues (timeout, neighbourhood exploration)

Next in Line...

Conclusion

CPS revolution: decentralisation (Cloud/Edge) and need to exploit intelligence on low-level devices for

- distributed processing
- situated deliberation
- fault tolerance
- better scalability
- low latency
- ...

Guidelines for distributing intelligence in CPCS (Cloud/Edge)

Reification of the methodology –providing feasibility– into:

- *micro-intelligence* as the way to distribute chunks of symbolic intelligence at the Edge of CPCS
- *LPaaS* to empower the Edge with knowledge and inference capabilities of computational logic

References I

Ananthanarayanan, G., Bahl, V., Bodik, P., Chintalapudi, K., Philipose, M., Sivalingam, L. R., and Sinha, S. (2017).
Real-time video analytics: The killer app for edge computing.
IEEE Computer, 50(10):58–67.

Calegari, R. (2018).
Micro-Intelligence for the IoT: Logic-based Models and Technologies.
PhD thesis, Alma Mater Studiorum—Università di Bologna, Bologna, Italy.

Calegari, R., Ciatto, G., Mariani, S., Denti, E., and Omicini, A. (2018a).
LPaaS as micro-intelligence: Enhancing IoT with symbolic reasoning.
Big Data and Cognitive Computing, 2(3).

Calegari, R., Denti, E., Mariani, S., and Omicini, A. (2018b).
Logic programming as a service.
Theory and Practice of Logic Programming, 18(3-4):1–28.

Chen, M., Li, W., Fortino, G., Hao, Y., Hu, L., and Humar, I. (2019).
A dynamic service migration mechanism in edge cognitive computing.
ACM Transactions on Internet Technology, 19(2).

References II

Cicirelli, F., Guerrieri, A., Mercuri, A., Spezzano, G., and Vinci, A. (2019).
ITEMa: A methodological approach for cognitive edge computing iot ecosystems.
Future Generation Computer Systems, 92:189–197.

Clark, A. and Chalmers, D. J. (1998).
The extended mind.
Analysis, 58(1):7–19.

Erl, T. (2005).
Service-Oriented Architecture: Concepts, Technology, and Design.
Prentice Hall / Pearson Education International, Upper Saddle River, NJ, USA.

Hollan, J., Hutchins, E., and Kirsh, D. (2000).
Distributed cognition: Toward a new foundation for human-computer interaction research.
ACM Transactions on Computer-Human Interaction, 7(2):174–196.

Hu, L., Miao, Y., Wu, G., Hassan, M. M., and Humar, I. (2019).
iRobot-Factory: An intelligent robot factory based on cognitive manufacturing and edge computing.
Future Generation Computer Systems, 90:569–577.

References III

Pace, P., Aloi, G., Gravina, R., Caliciuri, G., Fortino, G., and Liotta, A. (2019).
An Edge-based architecture to support efficient applications for healthcare Industry 4.0.
IEEE Transactions on Industrial Informatics, 15(1):481–489.

Rauch, E., Linder, C., and Dallasega, P. (In press).
Anthropocentric perspective of production before and within industry 4.0.
Computers & Industrial Engineering.

Tang, B., Chen, Z., Hefferman, G., Pei, S., Wei, T., He, H., and Yang, Q. (2017).
Incorporating intelligence in fog computing for big data analysis in smart cities.
IEEE Transactions on Industrial Informatics, 13(5):2140–2150.

Um, J.-S. (2019).
Futurology and future prospect of drone CPS.
In *Drones as Cyber-Physical Systems*, chapter 8, pages 257–274. Springer Singapore.

Waschull, S., Bokhorst, J., Molleman, E., and Wortmann, J. (In press).
Work design in future industrial production: Transforming towards cyber-physical systems.
Computers & Industrial Engineering.

Engineering Micro-intelligence at the Edge of CPCS: Design Guidelines

Roberta Calegari Giovanni Ciatto Enrico Denti Andrea Omicini

{roberta.calegari, giovanni.ciatto, enrico.denti, andrea.omicini}@unibo.it

Dipartimento di Informatica, Scienza e Ingegneria—Università di Bologna

Talk @ *12th International Conference on Internet and Distributed Computing Systems (IDCS 2019)*

June 11, 2019, Napoli, Italy