



Emotional Processes for Cyber-Physical Systems Resilience

Eskandar Kouicem, Clément Raïevsky, Michel Occello

► To cite this version:

Eskandar Kouicem, Clément Raïevsky, Michel Occello. Emotional Processes for Cyber-Physical Systems Resilience. IEEE International Conference on Systems, Man, and Cybernetics (SMC), Oct 2021, Melbourne, Australia. hal-03405713

HAL Id: hal-03405713

<https://hal.univ-grenoble-alpes.fr/hal-03405713>

Submitted on 27 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Emotional Processes for Cyber-Physical Systems Resilience*

Eskandar Kouicem¹, Clément Raïevsky¹ and Michel Occello¹

Abstract—The cyber-physical systems deployment takes up an ever-increasing number of fields of application. The variety of platforms used for implementing cyber-physical systems and the dynamic interconnection of their components make the resilience issue of these systems very important. Our work aims to use the knowledge from human and social sciences, particularly to inspire from the emotional processes for creating an agent architecture that increases the cyber-physical systems resilience. Emotional inspiration occurs in the individual decision-making processes as well as in the social coordination mechanisms. In addition, one of the principal hypotheses in our research is that the multi-agent paradigm is suitable for integrating emotional processes into cyber-physical systems.

I. INTRODUCTION

Cyber-physical systems (CPS) - composed by interconnected sub-systems which are partly interacting with the physical world - are increasingly used in many fields (smart building, smart grids...). The evolution of interoperability protocols, the decrease in hardware costs, and the simplicity of the interconnection between a variety of hardware components, provide the designers with a considerable number of configurations and combinations of components, which they don't always have the total operating control. In addition, these systems are often designed to be "open". As a consequence, it becomes impossible for the designers to predict all the situations that the system will encounter. This then raises the issue of the resilience of these systems, i.e. their ability to detect, manage and adapt to particular or abnormal situations that designers may not have anticipated [1]. To address this issue, we chose to draw inspiration from human emotions which are an integral part of our ability to adapt ourselves to particular, abnormal or unknown situations, both as individuals and as a group [2]. In order to exploit the knowledge of emotions in psychology and sociology to improve the resilience of CPSs, we have decided to adopt the multi-agent paradigm [3].

After a brief definition of resilience, we present our approach and show the relationships between resilience properties, emotional processes, and multi-agent systems (MAS). We then relate our approach to existing agent architectures. Finally, we illustrate our approach in some experiments.

II. RESILIENCE

Resilience has many different definitions: in ecology, resilience is the capacity of an ecosystem or species to

return to their normal behaviour after being subjected to a trauma; in psychology, resilience is defined as the ability of a person or a group to develop and continue to project themselves into the future despite the disruptive life events, difficult living conditions and sometimes severe trauma; in systems engineering, resilience depends on how fast the system can be restored after an incident; the computer networking community defines it as the combination of reliability and tolerance; the computer science community has defined resilience as the persistence of service delivery and the availability of functionalities [4].

For our work, we have adopted the definition of Woods [1], from the resilience engineering field:

"Resilience is the ability to recognise and adapt to handle unanticipated perturbations that call into question the model of competence, and demand a shift of processes, strategies and coordination."

This definition better matches with CPSs. According to the author, resilience requires process change and coordination, which can be found in the "cyber" and in the "configuration" levels of the 5C architecture for the implementation of CPS [5]. Also, this definition has to be distinguished from robustness, which is often confused with resilience in the field of artificial systems. For us, robustness does not include the adaptation of the system's behaviour following an adverse event but resilience include it according to the author. A robust system is able to "resist" to this kind of situation because of its design, by keeping its performance without modifying its functionality. Conversely, a resilient system will be able to detect abnormal situations and change its operations to maintain its critical functions.

A. Resilience profile

To further clarify our definition of resilience, we extended the work of Linkov and Kott [6] who decompose resilience in temporal phases. We obtained a resilience *profile* which includes four phases: **detection**, **absorption**, **recovery**, and **adaptation**. While all four phases are important in resilience, we will focus our study on the first two phases in this paper. To define these phases, we will relate them to the main features of resilience identified by Wei and Ji [7], Francis and Bekera [8] and Connelly *et al* [9].

The **detection** phase includes three resilience features:

- 1) *critical functions preservation*: the ability of the CPS to ensure that the functions identified as essential by the designers will be the last to be lost in the case of a disruptive event;

This work is funded by the Trust research chair of the Grenoble-INP Foundation and the Auvergne-Rhône-Alpes region.

¹Eskandar Kouicem, Clément Raïevsky and Michel Occello are with Univ. Grenoble Alpes, Grenoble INP, LCIS, 26000 Valence, France. firstname.lastname@univ-grenoble-alpes.fr

- 2) *anticipation*, which consists of using a model to predict the occurrence of some particular situations requiring a change in the operation of the CPS;
- 3) *planning*, which is ensured by establishing the foundations to maintain the services' availability and the correct equipment functioning during a disruptive event.

The **absorption** phase also includes three resilience features:

- 1) *reactivity*, which allows the system to react quickly to the solicitations of its environment;
- 2) *anomaly sensitivity*, which indicate if the system is capable of absorbing a disturbance and if the recovery time is appropriate;
- 3) *flexibility*: to increase resilience, it is essential to identify the main architectural elements which are preserved during some particular situations and which provide the flexibility required to be able to adapt over long periods.

We evaluated the ability of the multi-agent system approach and emotional processes to enhance CPSs with resilience-related properties. Based on this study, we estimated the correspondence between resilience phases features with MAS properties and resilience phases with emotional processes as in [9], [10]. The decision-making architecture presented in Section IV is based on this evaluation.

III. PROPOSED APPROACH

The main objective of this work is to provide cyber-physical systems with functionalities and processes that make them more resilient. To do this, we have adopted the multi-agent paradigm and aim at transposing knowledge from the psychology of emotions to artificial systems [10].

The multi-agent approach has been chosen because it intrinsically takes into consideration the distributed nature of CPSs. This approach is particularly well suited to resolve some problems related to traditional approaches of resilience, such as “*single point of failure*”. The autonomy of the agents, their capacity for self-organisation, and their capacity to dynamically change their role in the system provide solutions to communication loss and scaling-up problems.

This choice distinguish our approach from classical resilience solutions, which are often based on redundancy or centralised solutions, properties which are expensive and impractical in the CPS context [11], [12]. Agent-based design is also interesting because of its ability to integrate knowledge from social sciences such as emotions. It seems more practical to integrate this type of knowledge, which concerns the cognitive processes of individuals and groups, into the decision-making mechanisms of agents than into a centralised system.

Emotions are an essential part of our ability, as humans, to adapt to some particular, abnormal or unknown situations, both as an individual and as a group [2]. Hence our interest in transposing the emotion-related cognitive processes identified by researchers in psychology into the decision-making architecture of autonomous agents. The particularity of our

use of artificial emotions in agents compared to existing work consists in its generation process, and especially in the method of its initiation. The triggering of an emotion in existing approaches is based on a specific analysis of the situation, highly depending on the task of the system and the knowledge that the designer has about the situations the system will encounter. These assumptions in existing work intrinsically limits their resilience. Furthermore, most existing approaches trigger emotional episodes using symbolic information about the situation. However, this type of information is seldom available in the simple systems encountered in CPSs. To provide the cyber-physical systems with the adaptive capabilities of emotions without undermining their resilience, we aim at identifying generic, non-task-specific mechanisms to initiate artificial equivalent of an emotional episode. Among all the emotion-related cognitive functions, their signal function and social structuring role, which make consensus in cognitive science and psychology, underpin two major objectives of our work. We thus aim at providing artificial systems with a decision-making process that allows them to detect situations requiring an adaptive response at the individual level. We also work towards reproducing the social structuring role of emotions within a group to improve its adaptability and resilience. These processes are integrated into an agent architecture, which has been used for experimentation and validation of our hypotheses.

IV. DECISION-MAKING ARCHITECTURE

To organise the different decision-making processes of our agents, we chose to integrate them in a layered architecture [13], [14]. This type of decision-making architecture combines the simplicity, low algorithmic complexity, and fault tolerance of reactive architectures [15] with the ability to use non-local information, learning capabilities, and social interactions offered by more cognitive architectures. This kind of architecture poses however the challenge to manage the interactions between the different layers to obtain the desired behaviour [13].

The decomposition of cognitive processes into layers allows us to implement simple, perception-related, emotion eliciting mechanisms in a reactive layer while allowing higher level cognitive processes, potentially based on symbolic information and reasoning, to unfold without interrupting processes supporting critical functions.

The figure 1 illustrates the proposed layered architecture that integrates emotional processes to improve the resilience of CPSs. This architecture is composed of two main parts, a part constituted by the processes and behaviours of the agent, on the left, and a second part that manages the knowledge of the agent, on the right. Each side of the architecture is separated in three: the *reactive*, *proactive-deliberative*, and *social* layers.

The reactive layer (RL) is the only layer in direct interaction with the environment, for each perceived data or event, the reactive layer uses the knowledge of the perception grid (PG), this grid is inspired from the human mechanisms [16].

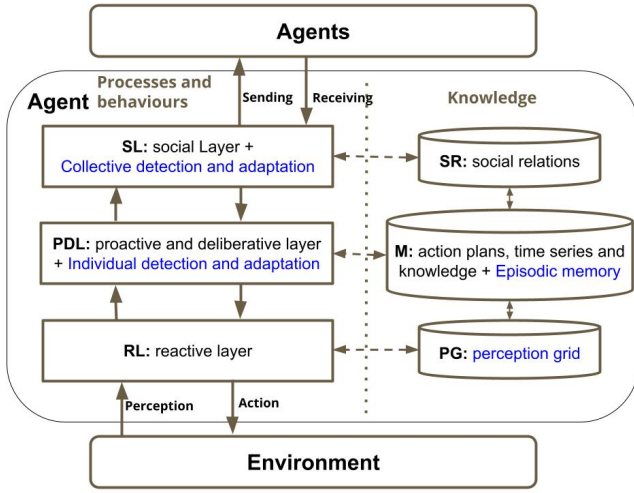


Fig. 1. The proposed architecture.

A situation in itself does not have any significance; each individual has his or her own perception grid which determines whether the situation is good or bad. In this context, the perception grid is an essential element in the first stage of the evaluation of the situation “the appraisal” - in psychology, the appraisal is the process that extracts emotions from our evaluations of events that cause specific reactions in different people [17]-.

In our architecture, the agent, especially its reactive layer, uses the perception grid to analyse the perceived data and to detect significant events and potentially abnormal situations. The events are initially neutral, but may raise doubts in the agent’s decision-making process. The knowledge in this grid is represented as a set of matrices. Each cell of a matrix corresponds to a data or a perceived event combined with a situation, which contains a Boolean value (false = no doubt and true = doubt). So, this knowledge affects the agent’s perception of its environment. At least, a part of it will be learned, which means that this knowledge is specific for each agent depending on its previous experience and its objectives. This grid is used as a first anomaly detector, and it will be used also for triggering the anomaly detection process in the upper layer.

In the proactive-deliberative layer (PDL), the agent takes decisions according to its data, its knowledge, and its action plans (M), if the reactive layer has doubts about the perceived data, an anomaly detection process will be launched at this layer, which is iterative and incremental and can even go up to the social layer (SL) that uses knowledge about social relations (SR) to interact with the other system agents as well as the diffusion and the negotiation of the detected situations.

The social relations layer contains the organisational structures of the system’s agents, and it is also used by the social layer for the diffusion of the individual adaptations to the detected situations in order to adapt collectively the agents’ behaviours and to auto-organise the agents groups. These processes will be illustrated by an example in the next section. We note that the layer (M) of our architecture integrates

an episodic memory [18], [19]. In cognitive psychology, the episodic memory refers to the process by which the human remembers previous events accompanied by their context (date, place, emotional state) [20]. We use this memory as a storage support of the abnormal situations encountered by the agent and the necessary adaptations of the reactive, the proactive and the deliberative behaviours to these situations. The episodic memory will also serve to update the perception grid and the action plans.

In our architecture, the actions are not necessarily initiated by a perception. The proactive and deliberative layer can trigger an action based on an internal decision, and the social layer can also trigger an action based on a message from another agent.

A. Resilience processes

The resilience processes in our architecture corresponding to tow phases presented in Section II-A:

a) *Detection*: this phase is considered as a monitoring of statistical anomalies in the perceived data time-series. It consists of several sub-phases, which belong to an incremental process:

Anomaly monitoring (out-of-domain values, unusual sequences or long repetition of similar data) in the sensor data. Each sensor evokes this sub-phase individually at the beginning, and then correlates it with other sensors. For a temperature sensor agent, its perception grid contains matrices that link temperature intervals with days of the year. For example, this agent perceives 30° in the morning of April 30, it takes the April matrix to see if the perceived value is normal or raises a doubt.

If an anomaly is detected, the time-series generated by the sensor agents will be analysed with more accurate mathematical tools. These two sub-phases correspond to **the occurrence of a situation** which triggers the emotional episode.

If the analysis confirms the anomaly, the sensor agent will consider the external (perceived data) and the internal (the perception grid, the episodic memory, the result of the analysis and the current goals) data to confirm or reject the abnormality of the situation. This sub-phase corresponds to the process of **evaluating the situation “appraisal”** of an emotional episode.

b) *Absorption*: the aim of this phase is to maintain the availability of the CPS’s most critical functionalities while isolating the perturbation. Increasing sampling frequency and power consumption will be required for system control and decision making. The sensor and actuator agents involved in the detected situation will communicate it to the other system agents. Following these actions, some goals will be prioritised in order to maintain the most critical functionalities.

There has been a modification of some of the system’s operating parameters in this phase, which leads to a modification of the readiness to action, which refers to **the activation of a behavioural script** in the emotional processes (but not its realisation) which aims to change the relationship between

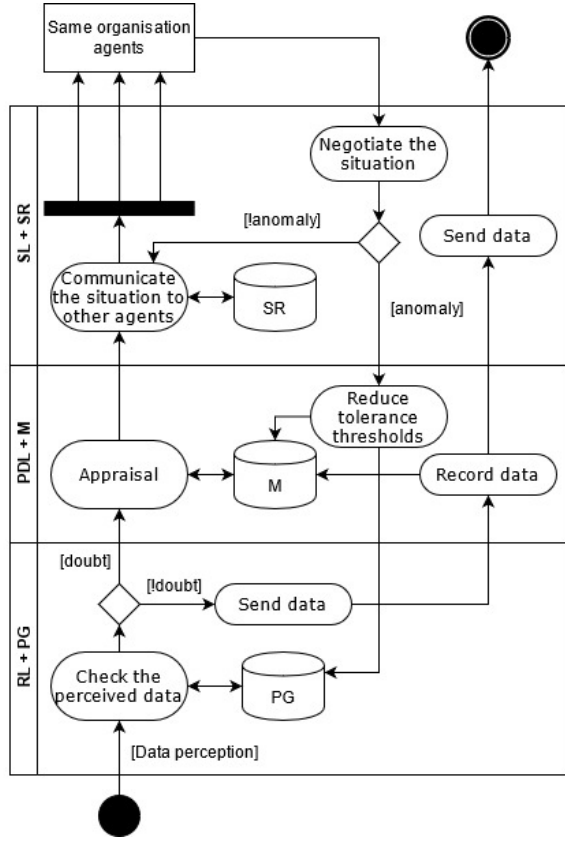


Fig. 2. Overview of detection and absorption phases.

the individual and his environment and focuses his attention on more important things.

The figure 2 illustrates how these processes are carried out using the proposed architecture.

V. EXPERIMENTS

A. System description

We chose a building's thermal regulation system (CPS_{btr}) to demonstrate the functioning of our agent architecture. This CPS is used to maintain a required temperatures in a building. To do this, it logs the physical data collected using temperature and luminosity sensors, and acts on its environment (the rooms) via actuators: heaters, air conditioners, rolling shutters and fans (see figure 3).

In CPS_{btr} , all components (sensors or actuators) are considered as agents, which are able to communicate, have a certain autonomy and have their own information. We will have several groups of agents in our multi-agent system, each group is composed by the agents of one room. Another type of agents organisation is taken into account in our system, which is the organisation of similar neighbouring agents. The agents interact via the *message exchange protocol*. The message format between two agents is: $msg = (messageID, senderID, receiverID, type, content, processed)$.

Examples of message types: *accept, command, demand, inform, modify, propose, refine, reject, request, anomaly, check, confirm*. The interactions between the layers of our




Fig. 3. Components in a room.

architecture are processed similarly to the communication between the agents.

In this case, the propagation of perceived data will be done automatically by the sensor agents. The user interacts with the system through the thermostat. The agents in the same room belong to the same organisation and have a direct communication between each other, so the actuators provide the other agents with their operating status and the sensors provide the environment's information (temperature, presence of the sun).

For the out-of-room communication, the same type components of adjacent rooms belong to the same organisation, and they are able to communicate directly with each other.

Each sensor agent records the recent perceived data as a time-series into its own memory (M), and the older time-series are shared into the actuators' memories of its own organisation. Initially, the agents have a pre-defined tolerance thresholds in terms of the stored data (M + PG).

Resilience is defined here as the ability of the CPS_{btr} to adapt in case of failure and to recover its nominal behaviour. Furthermore, the resilience in this system is that it continues to ensure the temperatures chosen by the users in case of problems.

B. Implementation

After conceiving CPS_{btr} using the multi-agent method "DIAMOND" [21], we have implemented it using Java. Both the agents and the environment are implemented as Threads which run permanently.

The reactive layer consists mainly of the following methods: "*perceive()*, *checkPerceptions()* and *act()*", which are responsible for perceiving information from the environment or messages from the proactive layer, checking the perceived data and forwarding it to the proactive layer or to change a part of the environment (e.g. change the state of an actuator).

The proactive layer records and transfers data between the two other layers. In case of a doubt raised by the reactive

layer, or a request to check a doubt from another agent, this layer uses the *appraisal()* method to evaluate the situation through a more rigorous anomaly detection process than the reactive layer's *checkPerceptions()* did.

The social layer checks permanently if there is any message to send or a new message from another agent. So "*checkMessages()*" checks both the local messages (received from the proactive layer) and the external messages (see figure 4).

To synchronise the layers, the agents and the environment we used a scheduler (ScheduledThreadPoolExecutor) to control their refresh rates. We used a logger to monitor the simulations. A JSON file containing the initial parameters of the environment is used to set up the simulations: the external temperature, the number of rooms and their temperatures.

As we said in Section II-A, we will focus our experiments on the first two phases of resilience which are: *detection* and *absorption*.

In our simulations, we used three adjacent rooms which means the temperature sensors can directly communicate between each others.

Once a temperature sensor agent's reactive layer perceives data from its environment, it stores it with the recently perceived data and checks whether it is an anomaly or not, using the perception grid. If the perceived data doesn't raise a doubt in the reactive layer, then a message with the type *request* and the perceived data as content will be sent to the proactive layer.

The proactive layer stores the received data as time-series in its memory, if the data doesn't contain any anomalies, it will build its episodic memory using this data sets. The agents tolerance thresholds are based on the collected data. Then, if the message type is *request*, this layer transfers the message to the upper layer, the social layer. The social layer is responsible for sending the messages which contain the perceived data by the sensor agent's reactive layer to the actuators of its organisation.

The following scenario illustrates the reactions of the agents in the case of anomalies. During the simulation, the environment randomly sends -10° to the temperature sensor in room 1 for several times to simulate the occurrence of abnormal situations. We logged the temperatures taken into

consideration by the agent after the appraisal and the total number of exchanged messages.

a) *Detection*: the detection is always active before every critical event until the absorption phase begins. Once a sensor agent's reactive layer has a suspicion on the perceived data, it will triggers a doubt based on the perception grid (*checkPerceptions()*). The sensor agent increases its sampling frequency (*decreaseStepsToSendData()*) and the reactive layer requests an intervention from the proactive layer via an *anomaly* message in order to access to the memory (M) and compare the perceived data with a larger time-series (*appraisal()*). The appraisal uses also the episodic memory in order to check if the agent has already experienced the current situation, it can directly confirm or affirm the doubt without negotiating it with other agents. If the doubt is confirmed, the *anomaly* message goes up to the social layer which communicates it as a *check* message to the same type agents in its organisation in order to confirm or affirm the doubt.

b) *Absorption*: for the other agents, if the situation is confirmed as abnormal by their proactive layers (*appraisal()*), they reduce their tolerance thresholds *decreaseThreshold()*, and communicate the result of the appraisal to the agent which initially detected the anomaly as a *confirm* message, or *reject* in the opposite case. This agent updates its episodic memory by creating an episode that reflects the negotiation result and then updates its tolerance thresholds and perception grid. Through this update, the agent will better maintain the system's functionalities by detecting and isolating the disturbance faster in the next occurrence of a similar situation.

C. Results and evaluation

We collected the logs of several simulations using our processes and without using them. We obtained the results in the figures 5 and 6. For evaluating our approach we chose some qualitative assessments. The qualitative assessments provide us a guidance for the development of cyber resilience measures [22]. Our assessment measures are mainly based on the features of resilience mentioned in the section II-A. As we can see in the figure 5, without the processes, the system can't preserve its *critical functions* which is the maintain of the desired temperature. Every time that the environment sends an abnormal value, the system tries to recover the desired temperature but and it seems not possible. Using our approach, the abnormal situations haven't impacted the system's functionalities. In this case, the system has detected the abnormal values and it was able to *anticipate* them.

Our processes allowed the system to get an advantage in terms of *reactivity* and *flexibility*. The reason for this is that once the desired temperature is reached, its value is not perturbed by the environmental factors as well as by the anomalies injected in the simulation.

In term of performance, our approach allowed the agents more autonomy, as we can see in the figure 6, the total number of exchanged messages between the agents and their layers has slightly decreased by using our approach. So we

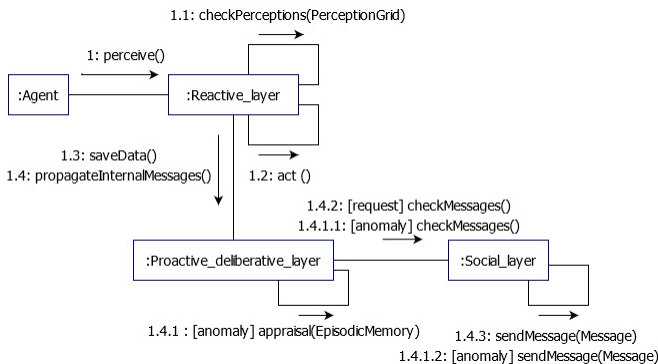


Fig. 4. Agent functioning.



Fig. 5. Temperature history for a single day simulation.



Fig. 6. Number of exchanged messages between the agents and their layers.

got a dual benefit, in terms of resilience on the one hand and in terms of autonomy and performance on the other.

VI. CONCLUSIONS AND PERSPECTIVES

In this paper, we presented our work on the use of the multi-agent paradigm for integrating knowledge from the psychology of emotions to improve the resilience of cyber-physical systems. We presented some resilience definitions and also our proposal and choices concerning the inspiration of human emotions as well as the choice of the multi-agent paradigm.

Our objective is to design and implement an agent architecture based on emotional processes, which integrates individual and decentralised processes for detecting abnormal situations. In addition, the architecture contains processes for adapting the individual behaviour according to detected situations in order to increase the resilience of the cyber-physical system. The individual detection and adaptation trigger a collective processes that use the information elaborated by the individual mechanisms in the system. The individual behavioural adaptation impacts on the group of agents controlling the cyber-physical system, which will trigger a self-organisation to adapt the collective behaviour for dealing with this abnormal situation.

The resilience detection and absorption phases were implemented and tested by simulation using the proposed agent architecture. The first results were interesting and we hope to get more relevant results in our future work. So, we plan to improve these two phases and to continue the implementation of the recovery and adaptation phases. We also plan to use our architecture in real cyber-physical systems to measure its resilience enhancements.

REFERENCES

- [1] D. D. Woods, "Essential characteristics of resilience," in *Resilience engineering*, pp. 21–34, CRC Press, 2017.
- [2] N. H. Frijda, *The emotions*. Cambridge University Press, 1986.
- [3] J. Ferber and G. Weiss, *Multi-agent systems: an introduction to distributed artificial intelligence*, vol. 1. Addison-Wesley Reading, 1999.
- [4] K. S. Trivedi, D. S. Kim, and R. Ghosh, "Resilience in computer systems and networks," in *Proceedings of the 2009 International Conference on Computer-Aided Design*, pp. 74–77, ACM, 2009.
- [5] J. Lee, B. Bagheri, and H.-A. Kao, "A cyber-physical systems architecture for industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 3, pp. 18–23, 2015.
- [6] I. Linkov and A. Kott, "Fundamental concepts of cyber resilience: Introduction and overview," in *Cyber resilience of systems and networks*, pp. 1–25, Springer, 2019.
- [7] D. Wei and K. Ji, "Resilient industrial control system (rics): Concepts, formulation, metrics, and insights," in *2010 3rd International symposium on resilient control systems*, pp. 15–22, IEEE, 2010.
- [8] R. Francis and B. Bekera, "A metric and frameworks for resilience analysis of engineered and infrastructure systems," *Reliability Engineering & System Safety*, vol. 121, pp. 90–103, 2014.
- [9] E. B. Connelly, C. R. Allen, K. Hatfield, J. M. Palma-Oliveira, D. D. Woods, and I. Linkov, "Features of resilience," *Environment systems and decisions*, vol. 37, no. 1, pp. 46–50, 2017.
- [10] E. Kouicem, C. Raěvsky, and M. Occello, "Towards a Cyber-physical Systems Resilience Approach based on Artificial Emotions and Multi-agent Systems," in *Proceedings of the 12th International Conference on Agents and Artificial Intelligence (ICAART 2020) - Volume 1*, pp. 327–334, 2020.
- [11] S. Hosseini, K. Barker, and J. E. Ramirez-Marquez, "A review of definitions and measures of system resilience," *Reliability Engineering & System Safety*, vol. 145, pp. 47–61, 2016.
- [12] S. Colabianchi, F. Costantino, G. Di Gravio, F. Nonino, and R. Patriarca, "Discussing resilience in the context of cyber physical systems," *Computers & Industrial Engineering*, p. 107534, 2021.
- [13] J. P. Müller and M. Pischel, "The agent architecture interrapp: Concept and application," 1993.
- [14] R. H. Bordini, A. El Fallah Seghrouchni, K. Hindriks, B. Logan, and A. Ricci, "Agent programming in the cognitive era," *Autonomous Agents and Multi-Agent Systems*, vol. 34, pp. 1–31, 2020.
- [15] R. Brooks, "A robust layered control system for a mobile robot," *IEEE journal on robotics and automation*, vol. 2, no. 1, pp. 14–23, 1986.
- [16] J. R. Anderson, *Cognitive psychology and its implications*. Worth publishers, 2000.
- [17] K. R. Scherer, A. Schorr, and T. Johnstone, *Appraisal processes in emotion: Theory, methods, research*. Oxford University Press, 2001.
- [18] T. Deutsch, A. Gruber, R. Lang, and R. Velik, "Episodic memory for autonomous agents," in *2008 conference on human system interactions*, pp. 621–626, IEEE, 2008.
- [19] D. Menager and D. Choi, "A robust implementation of episodic memory for a cognitive architecture," in *CogSci*, 2016.
- [20] E. Tulving *et al.*, "Episodic and semantic memory," *Organization of memory*, vol. 1, pp. 381–403, 1972.
- [21] J.-P. Jamont and M. Occello, "Meeting the challenges of decentralised embedded applications using multi-agent systems," *International Journal of Agent-Oriented Software Engineering*, vol. 5, no. 1, pp. 22–68, 2015.
- [22] A. Ligo, A. Kott, and I. Linkov, "How to measure cyber-resilience of a system with autonomous agents: Approaches and challenges," *IEEE Engineering Management Review*, 2021.