

Article

# Real-Time Compliant Stream Processing Agents for Physical Rehabilitation

Davide Calvaresi <sup>1</sup>  and Jean-Paul Calbimonte <sup>1</sup> \*

<sup>1</sup> University of Applied Sciences and Arts Western Switzerland HES-SO; name.surname@hevs.ch

\* Correspondence: jean-paul.calbimonte@hevs.ch; Institute of Information Systems HES-SO Valais-Wallis, TechnoPole 3, CH-3960 Sierre, Switzerland.

Version February 4, 2020 submitted to *Sensors*

**Abstract:** Digital rehabilitation is a novel concept that integrates state-of-the-art technologies for motion sensing and monitoring, with personalized patient-centric methodologies emerging from the field of physiotherapy. Thanks to the advances in wearable and portable sensing technologies, it is possible to provide patients with accurate monitoring devices, which simplifies the tracking of performance and effectiveness of physical exercises and treatments. Employing these approaches in everyday practice has enormous potential. Besides facilitating and improving the quality of care provided by physiotherapists, the usage of these technologies also promotes the personalization of treatments, thanks to data analytics and patient profiling (e.g., performance and behavior). However, achieving such goals implies tackling both technical and methodological challenges. In particular, (i) the capability of undertaking autonomous behaviors must comply with strict real-time constraints (e.g., scheduling, communication, and negotiation), (ii) plug-and-play sensors must seamlessly manage data and functional heterogeneity, and finally (iii) multi-device coordination must enable flexible and scalable sensor interactions. Beyond traditional top-down and best-effort solutions, unsuitable for safety-critical scenarios, we propose a novel approach for decentralized real-time compliant semantic agents. In particular, these agents can autonomously coordinate with each other, schedule sensing and data delivery tasks (complying with strict real-time constraints), while relying on ontology-based models to cope with data heterogeneity. Moreover, we present a model that represents sensors as autonomous agents able to schedule tasks and ensure interactions and negotiations compliant with strict timing constraints. Furthermore, to show the feasibility of the proposal, we present a practical study on upper and lower-limb digital rehabilitation scenarios, simulated on the MAXIM-GPRT environment for real-time compliance. Finally, we conduct an extensive evaluation over an implementation of the stream processing multi-agent architecture, which relies on existing RDF stream processing engines.

**Keywords:** Stream reasoning; real-time multi-agents; RDF stream processing; stream processing agents; digital rehabilitation; real-time sensors.)

## 1. Introduction

The demographic changes in our society, including the lengthening of life expectancy, entails several challenges regarding healthcare support and assistance for older adults. Chronic diseases and health issues affect an increasing amount of people among this population, often leading to a decrease in the quality of life. Some of the key factors related to this decline are the diminution of physical activity and the consequent reduction of mobility [1,2]. Many different clinical conditions can be co-factors originating impairment situations. Nevertheless, factors such as falls, progressive loss of mobility, and lack of exercise have also shown to have a high impact [3,4]. Numerous studies

34 have demonstrated the effectiveness of physical therapy for improving life conditions in different  
35 situations, including post-operative interventions for hip replacement, cancer survival treatments,  
36 back pain, stroke rehabilitation, among many others [5–7]. However, to maximize the efficacy of  
37 these rehabilitation strategies, providing personalized treatment, feedback, and support is essential.  
38 Moreover, several treatments (especially due to chronic conditions) require continuous monitoring  
39 over extended periods, often including home-based exercises and periodic check-ups. Furthermore,  
40 given the wide variety of physical rehabilitation procedures with healthcare professionals, patients  
41 tend either to abandon the treatments quickly or to decrease the adherence to the proposed exercises  
42 progressively.

43 Digital rehabilitation emerges as a promising approach that consists of leveraging information  
44 and communication technologies for boosting the efficacy of physical rehabilitation interventions [8].  
45 These technologies include, among others: (i) the usage of sensing devices with monitoring capabilities  
46 (e.g., able to capture motion and physiological data) (ii) applications based on machine learning and  
47 other data analytics methods; (iii) the deployment of distributed intelligent systems with streaming  
48 capabilities and *in-time* feedback.

49 In the last decade, the progress of the Internet of Things (IoT) and wearable sensors has paved  
50 the way for the implementation of integrated systems, which produce data in the form of continuous  
51 streams that require dynamic processing techniques [9]. In addition to the *velocity* and *volume*  
52 dimensions related to data streams, the *variety* of the data produced is particularly relevant when  
53 dealing with distributed heterogeneous sensors. This variety has been mitigated through the usage of  
54 machine-readable semantic data models (i.e., ontologies and knowledge graphs) in the areas of *stream*  
55 *reasoning* [10] and *RDF stream processing* [11]. Nevertheless, in application domains such as physical  
56 rehabilitation, *in-time* feedback is needed, as otherwise, a delay might jeopardize a measurement or  
57 endanger the safety of the patient [12]. Therefore, in such cases, traditional streaming data processing  
58 algorithms are not sufficient. In these scenarios, to guarantee strict processing deadlines and time  
59 restrictions is essential. Although such constraints have been long studied in the discipline of *real-time*  
60 systems, currently available (semantic) IoT architectures do not provide support for real-time compliant  
61 policies and protocols [13].

62 Recently, real-time techniques have been explored in the context of autonomous intelligent devices  
63 and multi-agent systems (MAS), with results that may have a potential impact in digital rehabilitation  
64 scenarios [12]. However, these works still have not explored the integration of real-time techniques  
65 within semantically-enabled stream processing systems. Thus, to support practical applications  
66 demanding for distributed intelligent entities interacting dynamically via the exchange of semantic data  
67 streams and negotiating/executing time-critical tasks, this paper addresses the following challenges:

- 68 • The inclusion of strict real-time compliance in RDF stream processing systems,
- 69 • The combination of autonomous real-time agents and semantic stream processing, and
- 70 • The application of real-time stream processing agents in digital physiotherapy scenarios.

71 With respect to the challenges mentioned above, the contribution of this paper can be summarized  
72 as follows:

- 73 • The definition of a model for real-time compliant stream processing agents, constrained by strict  
74 deadlines for interactions and negotiation among participating agents.
- 75 • The study and implementation of agent-based stream processing entities, based on the RDF  
76 stream processing paradigm.
- 77 • The simulation and analysis of both real-time and general-purpose MAS, considering different  
78 scenarios for digital rehabilitation using motion sensors.

79 These contributions constitute a concrete advancement with respect to the state of the art,  
80 especially regarding the inclusion of real-time strict constraints on stream processing agents. In this  
81 context, the main objectives of this work are: (i) to formally define real time constraints in RDF stream  
82 processing agents; (ii) to propose an architecture of real-time multi-agent RDF stream processing

83 system; (iii) to show the feasibility of the stream processing agents approach, through a concrete  
84 implementation based on existing RDF stream processing engines; and (iv) to demonstrate the effect  
85 and consequences of real-time constraints in simulated scenarios of a multi-agent digital rehabilitation  
86 application.

87  
88 The remainder of the paper can be summarized as follows. In Section 2 we describe the field of  
89 digital rehabilitation and how multi agent systems and semantic streaming data have been applied  
90 in recent years. In Section 3 we present the main challenges of real-time autonomous systems for  
91 digital rehabilitation, followed by the description of our model and architecture in Sections 4 and 5.  
92 Section 6 describes the simulation scenarios and their results, while Section 7 provides details on the  
93 implementation and experimentation on our architecture for RDF stream agents. Finally, Section 8  
94 concludes the paper.

## 95 2. Digital rehabilitation

96 Conventional rehabilitation practices are characterized by unilateral, interactive bilateral, and  
97 cooperative bilateral [12] interactions. In the framework of such interactions, for either physical  
98 or cognitive rehabilitation, the main activities that need to be addressed are training, counseling,  
99 monitoring, and assessment. The transition towards *digital* rehabilitation is mainly driven by the will  
100 of fastening the follow-up, enhancing the healing process, shortening the hospitalization, lowering the  
101 costs for both patients and health structures, enabling continuous monitoring, providing equitable  
102 access to rehabilitation services and finally supporting technological advancements in telemedicine [14].  
103 In recent years, the increasingly broad range of available technologies for physical activity monitoring  
104 enabled a face-paced advancement of these approaches. Application domains in the scope of digital  
105 rehabilitation include the usage of technologies ranging from: video analysis [15,16], wearable  
106 technologies [17], robotics [18,19], distributed sensing [20], and gamification [21,22]. This work focuses  
107 on physical rehabilitation. Therefore, the next section dives quickly into wearable-based rehabilitation  
108 systems.

### 109 2.1. Wearable-based rehabilitation systems

110 Wearable technologies constitute one of the key enablers for digital rehabilitation, and are expected  
111 to further provide improvements in both preventive and rehabilitation approaches. Although there are  
112 still concerns about the potentially invasive characteristics of wearable-based systems, a recent study  
113 (targeting patients in an older adult-care facility) revealed that 93% of the patients accepted body-worn  
114 sensor systems [23]. Concerned about the possible reluctance in using wearable-based systems (i.e.,  
115 *stigmatized*-like), Bergmann et al. [24] reported a surprisingly positive assessment of the aesthetics of  
116 those systems. Nevertheless, major concerns still arise with respect to restricted recording time (e.g.,  
117 due to limited storage capacity or limited battery-life), wearability, and reliable real-time feedback.

118 Wearable sensors employed in digital rehabilitation systems can span from being “simple”  
119 micro-sensors (e.g., capturing inertial movements or biomedical information by using small, intelligent,  
120 and low-energy active devices) to be “complex smart” (e.g., capturing data via thin and flexible sensors,  
121 compatible with textiles or made of textile technologies with specific mechanical, electrical or optical  
122 properties). Examples of rehabilitation systems based on this type of sensors have been applied to  
123 different use cases including exercise assessment in knee osteoarthritis [25], upper-limb motor training  
124 for stroke patients [26], or classification of motor activities for COPD patients [27]. In all these systems  
125 the focus is on ensuring the accuracy and efficiency of the sensing process and the results they produce  
126 in order to address the specific use-case needs. However, in all these systems the autonomy of sensing  
127 devices and the compliance to strict time constraints is not considered.

## 128 2.2. Real-time feedback and agents in digital rehabilitation

129 The Multi-agent system (MAS) approach is a comprehensive paradigm for modeling complex  
130 (distributed) systems and their dynamics. An agent is a (partially) autonomous entity, embodying  
131 a program, a sensor, a robot, or even a human being operating (sensing and actuating) in a given  
132 environment.

133 In the area of digital rehabilitation, we can count a plethora of contributions coming from the MAS  
134 community. For example, Rodriguez et al. [28] proposed a system assisting upper limb rehabilitation.  
135 Agents in this system perform abstract tasks such as (i) recording the movements while the patient is  
136 executing the exercise, (ii) receiving specific inputs (e.g., BPM, skin conductance), and defining the  
137 level of stress/fatigue, and (iii) behaving like a “virtual therapist”, adapting the therapy according  
138 to the current level of stress of the patient. Felisberto et al. [29] developed a MAS able to recognize  
139 patients’ movements and postures and to detect possibly harmful activities. Exploiting a wireless  
140 body area network (WBAN) as underlying system, an intelligent agent analyzes cyclically possible  
141 variations in the values received. The ultimate goal is to identify physical/posture deterioration. The  
142 consumption of medical drugs is a common eventuality in rehabilitation scenarios. Indeed, Mutingi et  
143 al. [30] proposed an agent-based system to cope with decision-making processes in drug delivery.

144 Summarizing, concerning the “lower” layers of wearable-based systems, they are able to perceive  
145 and pre-elaborate in-loco kinematic and biomedical parameters, and no one is powered by MAS.  
146 However, if further analysis is required, proprietary (often closed) solutions have to be involved.  
147 Conversely, solutions employing MAS can provide sophisticated, extensible, and scalable analysis.  
148 However, it has to be highlighted that it is challenging to deploy MAS on wearable sensors. The  
149 reasons hampering MAS from pervading wearable and embedded sensors stem from both technical  
150 and technological limitations [12]. In particular, MAS lack of technological means (e.g., frameworks)  
151 to be seamlessly deployed in distributed/wearable sensors characterized by scarce computational  
152 resources. Moreover, from the technical perspective, traditional MAS algorithms are incapable of  
153 dealing with strict timing constraints – crucial characteristic to deliver *in-time* feedback (typical of  
154 physical rehabilitation). Indeed, a recent study identified and formalized the challenges for MAS  
155 to be compliant with strict-timing constraints [13]. In turn, Calvaresi et al. [12] elaborated on the  
156 challenges related to the compliance of MAS with strict deadlines and embedded architectures.  
157 Moreover, the authors proposed and detailed a viable solution to bridge the gap between MAS  
158 and real-world wearable-based rehabilitation systems. Finally, an ultimate element still needs to be  
159 taken into consideration in the overall picture: the capability of processing streaming data in MAS for  
160 rehabilitation purposes [9].

## 161 2.3. Semantic data stream processing

162 Considering the prominent role of sensors and wearable devices for digital rehabilitation, it  
163 becomes a necessity to: (i) provide the means for managing data streams from these sensors through  
164 rich semantic models; (ii) enable the processing of querying and reasoning of these semantic streams,  
165 and (iii) provide real-time mechanisms for decentralized and autonomous interactions among stream  
166 processors. A number of RDF Stream Processing (RSP) systems have been developed in the last decade,  
167 focusing on the processing aspects of semantic streams, including incremental reasoning, continuous  
168 querying, complex event processing, among others [31–35]. However, most of these RDF stream  
169 processors rely on centralized architectures for interaction and organization, even if at least in theory  
170 they rely on Web standards.

171 In the literature, early attempts to provide REST-ful service interfaces for streaming data were  
172 explored in [36,37]. These exploratory proposals introduced the usage of Web infrastructure, which

173 evolved towards prototype developments such as the RSP Service Interface<sup>1</sup>, which further develops  
174 the ideas presented in [36], providing a generic implementable programming API for continuous query  
175 engines. The distributed execution of RDF stream workflows over a network has further been explored  
176 in the SLD Revolution framework [38], which optimizes interactions using lazy-transformation  
177 techniques so that non-optimized RDF formats are used only when necessary.

178 Regarding the publication of streams from distributed producers, TripleWave [39] enables the  
179 transformation of non-RDF data streams, as well as time-annotated RDF, into RDF streams. TripleWave  
180 allows the publication of these RDF streams so that they can be directly consumed or connected with  
181 applications that process them. The concept of TripleWave is further expanded in WeSP<sup>2</sup>, a conceptual  
182 model for producing and consuming RDF streams on the Web.

183 Even though the notion of decentralization is a fundamental aspect of the Web, in most of the  
184 streaming data and processing models the deployment of autonomous stream processors effectively  
185 remains challenging. The vision introduced by the *Semantic Web* [40], and its related initiatives,  
186 considered agents as primary actors for the generation and consumption of data on the Web. However,  
187 most implementations of the Semantic Web have focused on ontology modeling, reasoning engines,  
188 Linked Data, or RDF data querying, but have relegated agents to a marginal position. This also applies  
189 to semantic-aware streaming data processing (i.e., RSP and stream reasoning approaches). Although  
190 the introduction of agent and multi-agent approaches has been introduced recently [9,41], it is still  
191 required to provide concrete models, specifications, and implementations of viable agent-based stream  
192 processing and reasoning systems. Finally, regarding the inclusion of real-time constraints for RDF  
193 stream processing, so far most approaches have focused on *best-effort* stream processing, without the  
194 notion of compliance to strict deadline constraints.

### 195 3. Autonomous real-time streaming agents for digital rehabilitation

196 Digital rehabilitation can be applied to a wide variety of practices. Let us focus on the use case of  
197 post-operative knee physiotherapy. After surgical intervention, and the subsequent hospitalization,  
198 patients are traditionally left alone at the time of discharge, with a prescription that includes the series  
199 of exercises (therapy) that they are supposed to do. Such exercises may include simple strength and  
200 flexibility conditioning programs, e.g., *heel cord stretch*, *standing quadriceps stretch*, or *supine hamstring*  
201 *stretch*.


202 Commercial solutions in the market for monitoring this type of exercise have a number of  
203 significant limitations, especially regarding precision and usability. For example, devices such as the  
204 Kinetec [42] are used for performing passive and continuous knee movements during rehabilitation.  
205 This device is typically used during the acute phase, allowing to control among other parameters  
206 the angle of the knee movement. However, the angle of the machine does not precisely correspond  
207 to the angle of the knee itself, mostly because of structural reasons, limb misplacement, or attempts  
208 to compensate the movement performed by the patient trying to reduce an undetectable pain [12].  
209 This limitation leads to an inadequate assessment of pain, muscular resistance, and evolution of the  
210 treatment, added to other potentially misleading information due to the unsupervised usage of the  
211 Kinetec device.

212 As opposed to conventional physiotherapy with episodic encounters with health professionals to  
213 keep track of the patient's progress, the digital rehabilitation approach can rely on sensing devices  
214 for accurately capturing the movements of the patient during the exercises. These include triaxial  
215 accelerometer measurements, which can be used to calculate the flexion angle of the knee, number  
216 of repetitions, coordination of movements, duration of the stretching episodes, etc. However, this  
217 scenario requires the sensors to autonomously coordinate and organize their interactions and data

---

<sup>1</sup> <http://streamreasoning.org/resources/rsp-services>

<sup>2</sup> <http://w3id.org/wesp/web-data-streams>




**Figure 1.** Digital rehabilitation scenario: knee motion sensors for exercise monitoring and real-time feedback. In this scenario the knee sensors interact in real-time with the patient's application, which in turn provides immediate feedback, and provides information to the health professional about the exercise performance.

218 exchange within *strict* time constraints. For instance, two knee sensors computing flexion angles  
 219 during physical exercises must synchronize their interactions in real-time (Figure 1). If these real-time  
 220 deadline constraints are violated, the misalignment in the combined measurements jeopardizes the  
 221 accuracy of the angle calculations, resulting in inaccurate monitoring of the exercise performance.  
 222 Hence, once the angles have been calculated, the patient may need to get *in-time* feedback about the  
 223 exercise (e.g., through visual/audio indicators in her tablet or mobile device). Again, this feedback  
 224 requires to be bound in time. Otherwise, it would provide entirely useless information to the patient.  
 225 In particular, in the case of coaching, late and jeopardized feedback can be potentially dangerous:  
 226 e.g., a late *stop* feedback can generate an over-extension exceeding the safety range established for  
 227 that given patient. Finally, the physiotherapist and other physicians (or healthcare providers) may  
 228 need to monitor the exercise live, potentially on a tele-rehabilitation scenario. This last interaction  
 229 could require real-time feedback (in case of immediate response from the physiotherapist) or standard  
 230 communication requirements (in case of asynchronous feedback or off-line monitoring).

231 Extending this one-patient scenario to a setting where multiple patients are followed throughout  
 232 their treatment entails additional challenges (e.g., managing the heterogeneity of the different sensors,  
 233 increasing the potential activities to be monitored, and scaling to real-time multi-patient digital  
 234 rehabilitation monitoring). Indeed, each patient might follow a specific therapy, therefore having  
 235 different treatment prescriptions and body-parts to rehabilitate (thus possibly requiring various sensors  
 236 – in number and capabilities). Moreover, the evolution of the patient can also be different in terms of  
 237 compliance, time, and efficacy. For all these reasons, it is essential that beyond the real-time interactions  
 238 within the sensor environment, a digital rehabilitation system must also rely on semantic models for  
 239 representing information across patients and healthcare providers, using technology standards such as  
 240 RDF (Figure 2).

241 As it has been described in [12], a number of functionalities are needed in this type of digital  
 242 rehabilitation scenarios, including tools and methodologies for supporting: (i) adherence to the  
 243 treatment; (ii) monitoring of performance and correctness of the movements; (iii) enabling adjustments,  
 244 management of errors and compensations; (iv) coaching, encouraging and motivation of the patient; (v)  
 245 providing motivation, commitment and fatigue measurements; and (vi) incorporating practice-specific  
 246 parameters. Multi-agent systems have provided relevant contributions to all of them [43], although  
 247 breaching distributed sensing in real-world applications is still an open challenge. The usage  
 248 of affordable wearable sensors is essential to allow monitoring and *in-time* feedback in digital  
 249 rehabilitation scenarios. Indeed, a previous work [12] elaborated on the employments of MAS for  
 250 telerehabilitation, proposing a theoretical model and relevant future steps to undertake. Undoubtedly  
 251 the success of digital rehabilitation will also require the integration of diverse sensing technologies  
 252 and semantic compatibility. Although the question of sensor data heterogeneity has been addressed in



**Figure 2.** Digital rehabilitation scenario with multiple patients: decentralized and autonomous stream data management for heterogeneous sensors. Expanding from the previous scenario, real-time stream processing agents coordinate to provide per-user monitoring and feedback, while RDF streams are used for semantically-aware data exchange within a group of patients. The health-care professional can then monitor and provide personalized feedback according to individual performance.

253 past years, with promising results regarding the inclusion of ontology-based approaches for stream  
 254 processing in IoT environments, the integration of these methods with real-time compliant technology  
 255 remains mostly unexplored. Essential challenges in this respect include the addition of real-time  
 256 constraints in existing models such as RDF stream processing, or the representation of behaviors and  
 257 negotiation using semantic technologies.

258 Having described the context of digital rehabilitation and the need for autonomous systems for  
 259 managing *in-time* interactions among sensor devices, we elaborate the following set of challenges,  
 260 which need to be addressed:

- 261 • **Autonomous sensor interactions.** Beyond traditional IoT deployments configured following  
 262 top-down paradigms; digital rehabilitation often requires autonomy on the configuration of the  
 263 devices, as well as their synchronization and negotiation over data and services.
- 264 • **Real-time guarantees in sensor processing.** Given the necessity of complying with deadlines  
 265 and strict constraints on data execution, negotiation and delivery, autonomous sensors must  
 266 incorporate scheduling mechanisms to ensure these real-time guarantees.
- 267 • **Standard and extensible messaging & metadata.** Sensing devices should be able to exchange  
 268 data, in different formats and representations, potentially using Web standards for representing  
 269 metadata. Possible information to be specified are time constraints, performatives, conditions,  
 270 and negotiation protocols.
- 271 • **Asynchronous and distributed communication.** Sensors should be able to send and receive  
 272 messages, as well as coordinating among them without the need of a central entity that governs  
 273 their interaction flow.
- 274 • **Semantic stream data management.** Semantic representations should be employed to allow  
 275 sensors to understand and act accordingly to a given stream of data. These representations should  
 276 align with Web standards (e.g., OWL, RDF), and allow extensibility and high expressiveness.

#### 277 4. Real-time stream processing agents model

278 Having enumerated the main challenges related to real-time compliance in stream processing  
 279 agents for applications in the digital rehabilitation domain, in this section we introduce an agent-based  
 280 model that addresses those issues. First, we describe the RDF stream processing formalization used  
 281 throughout the paper. Then, we propose a multi-agent model in which behaviors are linked to

282 time-bound constraints. Finally, we describe how interactions among these agents are established,  
283 using semantic standards.

#### 284 4.1. RSP Data Model

285 The RSP model describes how data streams can be represented as potentially infinite sequences of  
286 RDF triples annotated with timestamps [44]. According to [45], we define an RDF triple as a tuple:

$$(s, p, o) \in (I \cup B) \times I \times (I \cup B \cup L)$$

287 where  $s$  is the *subject*,  $p$  is the *predicate*, and  $o$  is the *object*; and  $I$ ,  $B$ , and  $L$  are the disjoint infinite  
288 sets of IRIs, blank nodes, and literals, respectively. A set of RDF triples is called an *RDF Graph*.  
289 Moreover, collections of RDF Graphs can be represented as *RDF datasets*. A *named* RDF graph is a pair:

$$(n, G), \text{ with } n \in (I \cup B), \text{ and } G \text{ is an RDF graph}$$

290 Then, an RDF dataset is a set  $D$  defined as:

$$D = \{G_0, (n_1, G_1), (n_2, G_2), \dots, (n_i, G_i)\}$$

291 The RDF graph  $G_0$  is called the default graph for  $D$ , and  $(n_j, G_j)$  are named graphs, with  $j \in$   
292  $\{1, 2, \dots, i\}$ .

293 These time-agnostic definitions can be extended to support the notion of RDF streams, seen as a  
294 sequence of special RDF datasets, or *RDF stream elements*, each of which is formalized as follows:

$$s = (G_t, (n, G))$$

295  $G_t$  carries the timestamp of the stream element, including a triple of the form  $(n, p, \tau)$ , where  $p$  is  
296 a predicate that describes a timestamp annotation (e.g. `schema:observationDate`, or `ssn:resultTime`  
297 from well-known ontologies such as SSN, SOSA, or Schema.org), and  $\tau$  is the timestamp. For simplicity,  
298 we represent the timestamp of a stream element as  $\tau_s$ . The named graph  $G$  contains the payload of the  
299 stream element. Then, a stream  $S$  is defined as an unbounded ordered sequence of stream elements:

$$S = (\dots, s_i, \dots, s_j, \dots)$$

300 In this sequence, for every  $s_i$  and  $s_j$ , the order is established by their timestamps, i.e.  $\tau_{s_i} < \tau_{s_j}$ ,  
301 where  $\tau_{s_k}$  is the timestamp in  $G_t$  of the stream element  $s_k$ . Furthermore each stream  $S$  may be identified  
302 by an IRI  $n$ , in a tuple  $(n, S)$ .

303 Given the unbounded nature of streams, windows are defined as a way to refer to portions of the  
304 stream over time, so that processing and querying operators can be applied. A window  $W$  applied  
305 over a stream  $S$  consists of a finite set of stream elements from  $S$ . Different strategies can be used to  
306 extract this finite subset of elements from the stream. In this work, we focus on time-based windows,  
307 defined as follows:

$$W_{u,v}(S) = \{s \mid s \in S \text{ and } u \leq \tau_s < v\}$$

308 In this time window,  $W_{u,v}$  the time parameters  $u, v$  represent an interval that delimits the contents  
309 of the window based on timestamps of the stream elements. Notice that a window length could be  
310 specified instead of the interval upper bound. To illustrate the usage of windows, we present an  
311 example of a CQELS [33] query (Listing 1) that obtains heartbeat measurements from a specific sensor,  
312 under a window of 2 seconds. In this case, the stream is identified through a URI, over which the  
313 window boundaries are specified. Notice that different strategies can be applied as to how and when  
314 the contents of the window are filled (e.g., when the window closes, or as soon as the window content  
315 changes).



```

316
317 PREFIX sosa: <http://www.w3.org/ns/sosa/>
318 PREFIX m3lite: <http://purl.org/iot/vocab/m3-lite#>
319 SELECT ?heartbeat
320 WHERE {
321   STREAM <http://hevs.ch/streams/user1/kneesensor> [RANGE 2s]
322   {
323     ?obs sosa:madeBySensor ?hbSensor .
324     ?hbSensor a m3lite:HeartBeatSensor .
325     ?obs sosa:simpleResult ?heartbeat .
326   }
327 }
328

```

Listing 1: CQELS continous query over a stream of heartbeat observations.

#### 330 4.2. Stream processing agent model

331 A stream processing agent  $\alpha$  can be formalized as a tuple with the following structure:

$$\alpha = (K, G, B, F)$$

332 Where  $K$  represents the agent's beliefs,  $G$  corresponds to the goals,  $B$  represents the behaviors, and  $F$ , a selection function.

333 The set of beliefs  $K$  is represented as RDF statements on a dataset. In the case of data streams,  $K$  may also include stream elements from one or more streams, available through the application of a time window, as explained above. As an example, consider Listing 2, which represents a sequence of RDF stream elements on a stream in JSON-LD format. This specific example shows heartbeat observations, each of which is contained on a timestamped graph, as indicated in the model introduced previously. The agent may use this information as part of its beliefs, although as the stream is potentially infinite, it may only keep the latest heartbeat values, typically bounded through a time window.

```

341
342 {
343   "@context": {
344     "prov": "http://www.w3.org/ns/prov#",
345     "sosa": "http://www.w3.org/ns/sosa/",
346     "ex": "http://example.org#"
347   },
348   "@graph": [
349     { "prov:generatedAtTime": "2019-09-22T05:00:00.000Z",
350       "@id": "ex:Graph1",
351       "@graph": [
352         { "@id": "ex:heartbeatObs1",
353           "sosa:simpleResult": 34.5 } ]
354     },
355     { "prov:generatedAtTime": "2019-09-22T05:00:05.000Z",
356       "@id": "ex:Graph2",
357       "@graph": [
358         { "@id": "ex:heartbeatObs2",
359           "sosa:simpleResult": 44.5 } ]
360     }
361   ]
362 }
363

```

Listing 2: RDF stream elements containing heartbeat observations in JSON-LD. This example shows how stream elements can be represented as time-annotated graphs and containing RDF triples that represent the stream contents (e.g. sensor observations).

362  $G$  is a set of goals, which account for restrictions and targets defined for the agent. These may include time-bounded restrictions over the desired goals, which may have an impact on the real-time strict scheduling mechanisms, as we will see later.

363  $B$  represents the set of *behaviors* for the agent. Each behavior  $b \in B$  represents a strategy or procedure to solve a certain problem. Each behavior is characterized by the following elements  $b = (a, sel_a)$ .  $a \in A$  is an action from the set  $A$  of actions that the agent can execute, while  $sel_a : K \times (D, T) \rightarrow A$  is the function that defines which action to take depending on the current belief  $k \in K$ ,

369 and the temporal restrictions: deadline ( $D$ ) and period ( $T$ ). The deadline is the maximum elapsed time  
 370 to execute the action, while the period is the frequency at which the agent may launch the action.

371 Finally,  $F$  is a selection function that specifies which behavior the agent will take, according to the  
 372 current goals in  $G$  and beliefs in  $K$ .

373 In this paper, for simplicity, we consider the case of simple behaviors, referred to as *tasks*. One  
 374 particular task model is named *periodic tasks* (i.e., tasks whose execution recur endlessly after a  
 375 fixed period  $T$ ) [46]. Besides the period  $T$ , and relative deadline  $D$  this model of task also includes  
 376 information such as the release time  $R$ , and computation time  $C$  (typically a worst-case estimation). To  
 377 guarantee compliance with strict-timing constraints, the agent local scheduling algorithm evaluates  
 378 the feasibility of its task-set based on the related *utilization factor*. In particular, the utilization factor  $U_b$   
 379 of a given task (periodic behavior)  $b$  in a task-set  $\Gamma$  is computed dividing its computation time  $C_b$  by  
 380 its period  $T_b$  (when period is equal to deadline:  $T = D$ ). Therefore, the utilization factor of a given  
 381 agent  $\alpha$  at a given time  $t$  is defined by

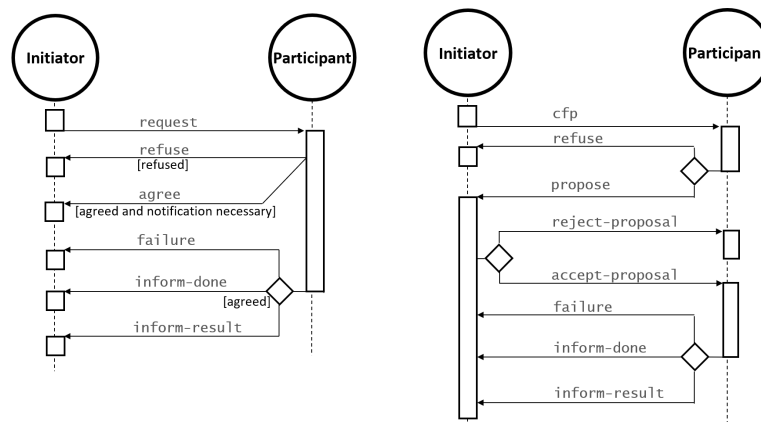
$$U^\alpha(t) = \sum_{b \in \Gamma(t)} U \text{ with } U = \frac{C}{T} \quad (1)$$

382 Then, the task-set is considered feasible if its utilization factor is less (or equal) than the *least upper*  
 383 *bound* ( $\leq U_{lub}$ ) of its scheduling algorithm [46].

#### 384 4.3. Negotiation among RSP Agents

385 RSP Agents interact with each other in a structured manner, through protocols based on existing  
 386 standards for multi-agent cooperation, negotiation, and delegation. Among these protocols, we focus  
 387 on those defined by FIPA [47]. In Figure 3 depicts two examples of standard FIPA negotiation protocols.  
 388 Figure 3 (a) represents a request for a given task/data from a given *initiator* to a given *participant*. This  
 389 latter can accept or refuse the request. If the answer is positive and the initiation awards the task  
 390 execution, the participant will start performing it.

391 Figure 3 (b) depicts a protocol in which the initiator launches a *call-for-proposals* (cfp), asking other  
 392 agents to bid for the execution of a task. The participant then makes a proposal (or called *bid*), which  
 393 can be accepted or rejected by the initiator.




(a) FIPA: request protocol.

(b) FIPA: call for proposals.

**Figure 3.** Examples of FIPA interaction protocols. The standard protocols establish the way in which agents interact, providing an abstraction of generic message exchanges.

394 Using primitive negotiation protocols of this type as building blocks, RSP agents can develop  
 395 more complex interactions among them. As an example, consider limb sensors producing continuous  
 396 streams of movement observations in Figure 4. The first agent requires raw data from the second agent  
 397 and thus performs a request, in the form of an RSP query. After agreeing, the second agent will inform  
 398 the results of the request, but unlike a traditional request, it will stream the sensor data back. There are

399 different possible ways of actually implementing data delivery, as we will see in Section 5, but in a  
 400 general sense, we can abstract the stream as a sequence of messages delivered in a specific order.



**Figure 4.** Simple request interaction between two sensor agents: continuous RSP query. This interaction can be embedded into a FIPA request protocol, with the addition of a continuous set of final responses instead of a single one. The continuous query will expect multiple responses until the query is canceled or its execution expires.

401 As another example of interaction among agents in digital physiotherapy in Figure 5, we illustrate  
 402 a scenario in which the agent organizes and coordinates their work autonomously, on behalf of the  
 403 sensors and devices they represent. First, two body sensors publish the tasks or behaviors they can  
 404 provide, using a *register* agent as an intermediary. These tasks can refer, for example, to the provision  
 405 of RDF streams of movement computation data, or physiological aggregated data. Then, the agent  
 406 acting on behalf of the patient monitoring system asks the registry for metadata about those agents  
 407 providing a specific type of task. Having the list of sensors that provide these features, it emits a *cfp*  
 408 request, to which the sensor agents may bid, and get awarded the execution of the job. Once the bid is  
 409 granted, the exchange of RDF streams can be established using a specific channel (e.g., MQTT, HTTP,  
 410 and WebSocket).



**Figure 5.** Call-for-proposal interaction among stream processing agents. This interaction can be represented as a FIPA call-for proposals, in which sensor agents publish the tasks they can provide (e.g. sensing and monitoring) in a register agent. Then an application agent can request for agents who are capable of a certain tasks, and emit a *cfp*. After a bid is performed, the task can be scheduled and executed.

411 The protocol described above shows the capacity of RDF stream agents to establish collaboration  
 412 strategies in a decentralized manner. However, to allow the compatibility and understanding of  
 413 the protocols, it is essential to represent these interactions through semantic and machine-readable  
 414 standards. In Listing 3 we show an excerpt of a FIPA *call-for-proposals* in JSON-LD, which can be used  
 415 as a message exchanged among RSP agents, e.g. for a scenario such as the one in Figure 5.

416  
 417

```

418 {
419   "prov:generatedAtTime": "2017-09-14T04:00:00.000Z",
420   "@id": "ex:callForProposals1",
421   "@graph": [
422     { "@id": "ex:cfp1",
423       "ag:permormative": "ag:CallForProposals",
424       "ag:sender": "ex:agent1",
425       "ag:protocol": "ag:ContractNet",
426       "ag:ontology": "http://example.org/health0ntology#",
427       "ag:content": "..."} ],
428   "@context": {
429     "prov": "http://www.w3.org/ns/prov#",
430     "ex": "http://example.org#",
431     "ag": "https://w3id.org/rdf-agents/msg#"
432   }
433 }

```

Listing 3: FIPA call for proposal message represented in JSON-LD. This example shows an excerpt of a message emitted by an agent soliciting a service or task to which other agents may bid for. The message itself is represented in RDF and is exchanged through the RSP agent interfaces.

434 The proposed model presented in this section essentially contributes on three main novel aspects  
435 with respect to the state of the art: (i) the addition of real-time constraints in the RDF stream processing  
436 model; (ii) the modelling of RDF stream processors (or reasoners) as decentralized agents; and (iii) the  
437 alignment of RSP agent interactions as extensions to standard multi-agent system protocols, with a  
438 particular accent on sensor-based application such as digital rehabilitation.

## 439 5. An agent-based architecture for decentralized RDF Stream Processing

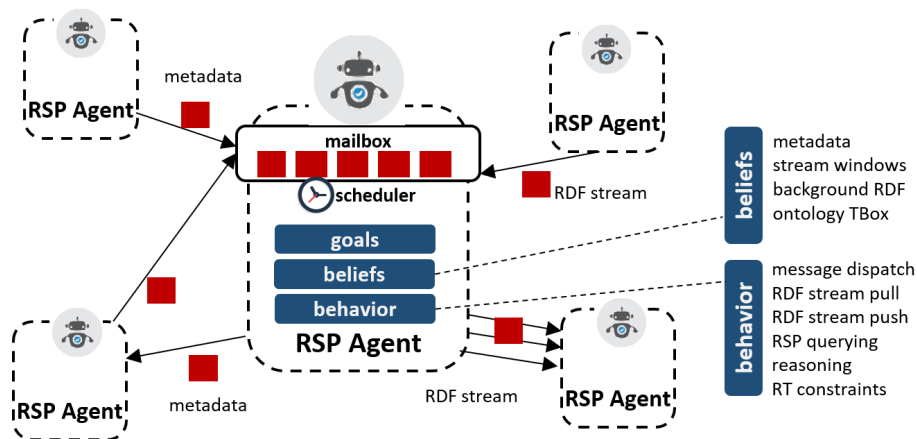
440 The model described in the previous section introduces the notion of stream processing agents,  
441 capable of acting according to beliefs and goals (e.g., knowledge) represented as RDF streams, and  
442 including real-time constraints. The streaming nature of the data exchanged by these agents makes it  
443 necessary to change certain paradigms regarding the delivery and processing of data, leaning towards  
444 continuous data processing mechanisms. The Web provides the necessary infrastructure and standards  
445 for enabling decentralized communication and interaction among software agents and constitutes an  
446 appropriate foundation layer for the implementation of the model in Section 4. To successfully make  
447 RDF streams available on the Web, Dell’Aglia [11] outlined a set of requirements partially derived from  
448 the more general guidelines for stream processors [48]. These can be summarized as (1) prioritizing  
449 active paradigms for data exchange, (2) combination of streaming and stored data, (3) availability,  
450 distribution, and scalability, (4) wide range of stream operations, (5) availability of stream metadata,  
451 (6) support for a variety of streams, and (7) reuse of existing protocols and standards.

452 While we endorse these requirements in general, this paper emphasizes the need for guaranteeing  
453 that RSP engines interact with each other in a decentralized manner, following the nature of the Web.  
454 This implies a departure from the usual setting in previous RSP approaches, where a server-centric  
455 paradigm governs continuous query processing and data flows among entities. An example of such  
456 an approach is reflected in the interaction patterns of a continuous query workflow, where the entire  
457 focus is solely on the query engine server.

458 We propose an architecture centered on *RSP agents* (i.e., autonomous agents that can be deployed  
459 in a distributed fashion and that are able to communicate and exchange RDF streams) and their  
460 corresponding metadata (Figure 6) through their *inbox*.

461 As described in the model, each RSP agent encapsulates a set of beliefs, goals, and behaviors, and  
462 interacts with other agents exchanging messages and streams of data through its inbox. Each agent  
463 may act as a sender or receiver of two main types of messages:

- 464 • *RDF stream elements*: these are RDF triples or graphs from a given RDF stream, as defined in  
465 Section 4. The stream delivery of these messages, either pulled or pushed, can be applied in  
466 different scenarios (e.g., feeding a stream, delivering query answers, and pushing reasoning  
467 entailments).



**Figure 6.** RSP Agents architecture. Each agent encapsulates beliefs, goals and behavior, and manages incoming messages through its mailbox. The execution of real-time constrained tasks is governed by an internal scheduler, and the constraints themselves are included in the agent behavior. Ontologies and vocabularies are included as part of the agent beliefs or knowledge, as well as the contents of dynamic stream that the produce/consume. Interactions among RSP agents happen through RDF message exchange of either metadata or continuous RDF streams.

- 468 • *RDF stream metadata*: these are essentially metadata messages required to perform tasks such as  
469 retrieving a stream description, declaring and RDF stream, filter a set of stream endpoints, and  
470 declaring a query.

471 Each agent is linked to a unique identifier that can be used to locate it, and it can have a set of  
472 endpoints, which can be used to reach the RSP agent resources (i.e., its streams and metadata). The  
473 resources of each RSP agent includes the metadata of the RDF streams it manages, as well as other  
474 information relative to them (i.e., background RDF datasets, RDF stream buffers, ontology TBoxes,  
475 and RDF constraint rules). Not all of these resources need to be accessible to other agents. The goals  
476 and beliefs of the agent are typically private, and may evolve over time, as they may be affected by  
477 changes in incoming RDF streams. The behavior of each agent defines how it proceeds at the arrival  
478 of incoming messages. In particular, it typically implements internal processing mechanisms such as  
479 continuous query processing, complex event processing, and stream reasoning. To do so, the agent  
480 may emit new messages (e.g., response to a query), create new agents (e.g., a pushing emitter, or a  
481 subscriber handler), or schedule other actions (see Figure 6).

482 Regarding the requirements mentioned above, this architecture addresses them in the following  
483 ways. For requirement 1, it natively supports asynchronous message passing, including the ability to  
484 push streams of messages if necessary. Concerning the combination of streams and stored data, the  
485 model takes a no-sharing approach for state information, so, in principle, any stored RDF data is only  
486 locally accessible and modifiable. The solely allowed procedure to exchange it is through message  
487 delivery, which is fundamental also to guarantee scalability and distribution (requirement 3). The  
488 behavior of each RSP agent allows sufficient freedom and flexibility to implement different types of  
489 operators and processing mechanisms (requirement 4), while the explicit definition of RDF stream  
490 metadata covers (requirement 5). The variety of streams is not restricted by the model (requirement 6),  
491 and the usage of well-established standards is also advocated.

492 Finally, the execution of behaviors is governed by the *scheduler* component, which may include  
493 different strategies [49] depending on the type of tasks and the constraints it may have. In particular,  
494 in this paper we refer to real-time compliance to strict deadlines for task execution, as described in  
495 Section 4.

## 5.1. Messages and Notifications in RSP Agents

The RSP agent architecture attributes particular importance to the exchange of messages, as they are the basic way to share information and coordinate interactions. The architecture adopts specific considerations about how the messages (also called notifications) are handled, taking into account the differences in dealing with streaming vs. stored RDF data.

**Format & vocabularies.** Messages in RSP agents fall under two fundamental categories: RDF stream elements, and RDF stream metadata. In both cases, RDF is the underlying data model used to represent the information that is exchanged, although there are minimal expectations. For RDF stream elements, these are expected to conform to the general RDF stream abstract model, as described by the W3C RSP Community Group<sup>3</sup> and as in Section 4. However, this abstract model provides high flexibility concerning the use of a particular vocabulary (e.g., using the SSN ontology for representing sensor streams, the Event ontology for streams of events, or the PROV ontology for provenance descriptions). Concerning the metadata, it would be advisable to provide a standard vocabulary for RDF stream descriptions, as proposed in [50], although this goes beyond the scope of this paper.

**Message storage.** RSP engines are designed in such a way that RDF stream elements flow through them, and produce continuous results. Therefore, the stream is not stored, at least not in the way it is done in a traditional database or data store. Consequently, RDF stream elements are accessed through time windows, as in the RSP model, and older messages are bound to *fade* as time passes.

**Message resolvability.** As a consequence of the previous observation, it is hard to allow stream messages to be retrievable after an RSP agent has processed them. This differs from other RDF/Linked Data use cases where data dynamics do not follow a streaming paradigm. As we will see later, this leads to the introduction of *input* and *output* RDF streams, which restrict RSP agents to either only write or read from a stream. In any case, resolving a particular stream element is of less importance in the context of RSP, than resolving the current contents of a stream, or a view over a stream.

**Push message delivery.** While on traditional Web standards, pulling is the primary method for delivering data (e.g., through HTTP GET/POST requests), it is not always the most suitable option for data streams. As described below, our proposed RSP agent architecture provides alternative delivery methods, allowing the usage of WebSocket or HTTP Server-sent-events. The abstract description of the RSP agent leaves the message delivery method open for either of the options.

**Querying.** Given the ubiquity of query access patterns in RDF stream processing, it would be natural to include explicit interaction specifications for registering standing queries, as well as accessing their results as streaming notifications. This applies not only for window-based continuous queries, but also for Complex Event Processing, and given use cases in stream reasoning.

## 5.2. Stream Receivers, Senders and Consumers

The RSP agent architecture specifies three main types of agents: *Stream Receiver*, *Stream Sender*, and *Stream Consumer*. An RSP agent may play the role of one or all of these types.

### 5.2.1. Stream Receiver

The Stream Receiver is a profile for an RSP agent that is capable of receiving and processing the following types of messages:

---

<sup>3</sup> <https://www.w3.org/community/rsp/>

- 540 • **RetrieveAllStreams**: to request metadata of all RDF streams registered in the receiver agent.
- 541 • **CreateStream**: to request the declaration of an RDF stream. This message includes the metadata
- 542 of the RDF stream to be created.
- 543 • **RetrieveStream**: to request for the metadata of a given RDF stream in the receiver agent. Its IRI
- 544 identifies the requested stream.
- 545 • **SendStreamItem**: to add a stream element to an existing RDF stream residing in the requested
- 546 receiver agent. This message includes the stream element itself, as well as the RDF stream IRI.
- 547 • **RetrieveStreamItem**: to request for a specific stream element. The message includes the IRI of the
- 548 RDF stream and the element itself. As in-stream systems and element might be volatile, in the
- 549 sense that it might not be de-referenceable after some time, this message also includes views
- 550 over stream elements (e.g., based on time recency)
- 551 • **PushStreamItems**: to request for stream items to be pushed back. The message includes the RDF
- 552 stream IRI.
- 553 • **CreateQuery**: to request for a continuous query to be registered. The query includes the reference
- 554 to the stream IRIs to be used and the IRI of the resulting stream of responses.

555 The Stream Receiver will act upon arrival of any of the above messages to its inbox. We show in  
 556 Algorithm 1 a sketch of how the agent reacts to these messages. The receive method of the agent is the  
 557 interface used to indicate what action to take in each case.

---

#### Algorithm 1 Stream Receiver: receive function

---

```

1: procedure RECEIVE(msg)
2:   sender ← msg.sender
3:   switch msg:
4:   case RetrieveAllStreams:
5:     send(getAllStreams) to sender
6:   case CreateStream:
7:     ack ← postInputStream(msg.body)
8:     send(ack) to sender
9:   case RetrieveStream:
10:    send(getStream(msg.uri)) to sender
11:  case SendStreamItem:
12:    postStreamItem(msg.uri, msg.body)
13:  case RetrieveStreamItem:
14:    send(retrieveStreamItem(r.uri)) to sender
15:  case PushStreamItems:
16:    handler ← pushStreamItems(msg.uri)
17:    handler.onReceive(data) :
18:      send(data) to sender
19:  case CreateQuery:
20:    ackgetsPostQuery(msg.body)
21:    send(ack) to sender

```

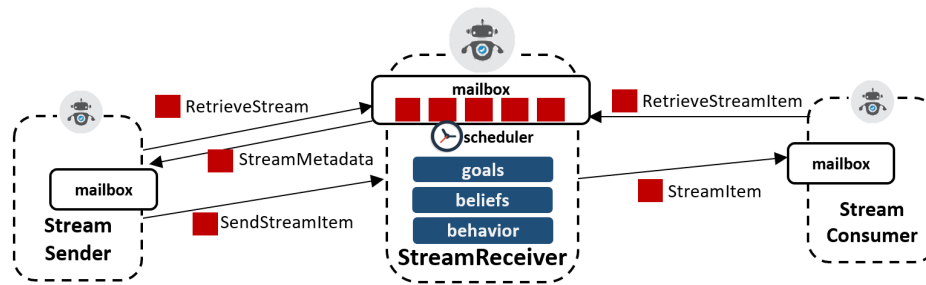
---

558 The steps taken are self-explanatory in most cases. The Stream Receiver calls internal methods,  
 559 for instance to create a stream (*postInputStream*) or to retrieve a stream item (*retrieveStreamItem*). In  
 560 practice, these methods will be implemented on top of exiting RSP, CEP, or stream reasoners, binding  
 561 their native implementations to this interface.

562 The Stream Receiver can manage several RDF streams registered within it. Streams may be of two  
 563 different kinds:

- 564 • **Input streams**: these streams are essentially meant only to receive new items, but are not intended
- 565 to be consumed by other agents other than the one that hosts it. Examples of such streams are
- 566 those used as input for RSP queries: other RSP agents can feed these streams, but the query
- 567 processor on the Stream Receiver is the only one that consumes it.
- 568 • **Output streams**: these are those RDF streams that are meant to be consumed by other RSP agents,
- 569 but fed only by the agent that hosts it. An example of such a stream is the continuous result of an
- 570 RSP query engine.

571 RDF streams can also be available as both input and output.



**Figure 7.** RSP Stream Receiver, processes incoming messages from the sender that posts RDF stream items, and sends RDF stream items to a consumer. In case of real-time constraints, the scheduler will rely on a real-time compliant strategy to satisfy the established policies.

### 572 5.2.2. Stream Sender

573 This type of RSP agent characterizes those interactions related to sending RDF metadata, as well  
574 as RDF stream contents to another agent. The sender defines the following basic operations:

- 575 • *postStream*: send RDF stream metadata to declare it on a Stream Receiver. The sender emits a  
576 *CreateStream* message through this operation.
- 577 • *postStreamItem*: send and RDF stream element to an (input) stream on a given Stream Receiver.  
578 This is typically a feed stream message.
- 579 • *postQuery*: register a query on a Stream Receiver with a *CreateQuery* message.

580 Apart from these operations, a sender must also be able to discover the Stream Receiver endpoints.  
581 For this, it has a discover operation, which for a given stream IRI, requests the endpoint or endpoints  
582 available for sending (or consuming) stream elements. The sender also has operations to retrieve  
583 the metadata of a given RDF stream, or all available RDF streams on a receiver (*getStream*, and  
584 *getAllStreams*, respectively). These operations are common to a Stream Consumer, described below.

### 585 5.2.3. Stream Consumer

586 A Stream Consumer characterizes RSP agent interactions relative to receiving RDF stream data. It  
587 essentially defines two operations:

- 588 • *getStreamItem*: requests to consume an RDF stream item. Implementations of this operation can  
589 derive different strategies for retrieving RDF stream contents. Given the dynamicity of streams,  
590 it is usually unfeasible to collect them one by one through their identifiers. Alternatively, these  
591 implementations may rather rely on stream views that may capture, for example, the latest  
592 stream items on a given window of time, or the ones complying to some filtering criteria.
- 593 • *pushStreamItems*: requests stream items to be pushed to the consumer. Conversely to the previous  
594 operation, which is essentially poll-based, this one requests the receiver to act as a sender as soon  
595 as there is an RDF stream element available for consumption.

596 For example, let us consider the Stream Receiver depicted in Figure 7. First, it receives a message  
597 on its inbox, requesting the metadata of a specific stream. The receiver dispatches back the metadata  
598 to the requester, which then post new elements to this stream. Then a consumer may also request  
599 a specific stream item to the receiver through a corresponding message. To make these interactions  
600 possible, the agents need to have first the addresses of the other agents, and if necessary, discover their  
601 endpoint locations.

602 The RSP agent architecture detailed in this section constitutes a first specification of how the stream  
603 reasoning agents vision [9] can be implemented. Beyond existing stream processing engines for RDF  
604 data and stream reasoners, which typically work based on top-down organization and communication  
605 approaches, this architecture relies on decentralized RSP agents, capable of self-organizing and  
606 coordinating processing actions and streaming data exchange.



## 607 6. Stream processing agents simulation

608 To design and assess the behavior of a real-time compliant multi-(streaming)-agent system for  
 609 digital rehabilitation, we have set up a virtual environment using the MAXIM-GPRT [51] simulator.  
 610 More precisely, the case study presented in Section 3 has been used to define scenarios in which  
 611 real-time compliance is assessed, under different conditions and parameters. The simulator enables  
 612 the design and analysis of multi-agent behaviors composed of both General-Purpose (GP) and  
 613 Real-time (RT) algorithms. Therefore, MAXIM-GPRT has been used to design, simulate, and assess the  
 614 performance of several agents' setups with respect to performance metrics such as compliance with  
 615 strict timing constraints (deadline miss ratio), utilization factors, negotiated workload, response time,  
 616 and lateness.

### 617 6.1. Simulated Scenarios and Setups

618 We studied the following scenarios in the context of digital rehabilitation:

- 619 S1: *One physiotherapist, One patient, Two sensors;*
- 620 S2: *One physiotherapist, Two patient, Two sensors;*
- 621 S3: *One physiotherapist, One patient, Five sensors.*


622 The S1 scenario can refer to a simple rehabilitation setting, in which body sensors capture  
 623 motion data from a specific part of the body (e.g., knee movement) during particular exercises. The  
 624 patient monitoring application is also represented through an agent, which can coordinate with the  
 625 physiotherapist agent. S2 refers to a similar scenario, but extended to more than one patient. The  
 626 scenario in S3 considers more motion sensors (and agents), for more complex rehabilitation treatments.  
 627 These scenarios serve as a basis for exploring the potential of the agent-based architecture and allow a  
 628 system designer to choose the most suitable configuration for a given situation.

629 These scenarios have been studied considering real-time assumptions for negotiation and  
 630 interaction among agents. In particular, according to [13], (i) Earliest Deadline First (EDF) has been  
 631 employed as agent local scheduler (coupled with the Constant Bandwidth Server - CBS - to deal with  
 632 sporadic and aperiodic tasks) [46]; (ii) it has been assumed a bounded-time delay (e.g., between 10  
 633 and 50 milliseconds) communication middleware (RTPS-like) [52,53]; and (iii) the Reservation-Based  
 634 Negotiation protocol (RBN) has been used to enforce a real-time compliant task negotiation among  
 635 the agents [54]. Moreover, to analyze and prove the unsuitability of best-effort approaches (e.g.,  
 636 general-purpose such as FIFO and RR-like) in the studied scenarios, S3 has been studied only varying  
 637 the agent local scheduler (i.e., FIFO) and the negotiation protocol (i.e., contract Net - CNET [55])  
 638 keeping unaltered inputs and communication middleware. Hereafter, we refer to S3 as S3rt (real-time  
 639 configuration) and S3gp (general-purpose configuration).

640 For the sake of clarity, MAXIM-GPRT allows several more real-time and general-purpose  
 641 configurations. However, assessing the impact of every available algorithm with respect to this  
 642 case study goes beyond the scope of this paper. The configuration mentioned above produced  
 643 sufficient results to well identify and explain the challenges, capabilities, and performances of both  
 644 general-purpose and real-time algorithms at the service of streaming-agents for digital rehabilitation  
 645 purposes. Table 1 details the general setups of the tested scenarios and Figure 8 shows the role of every  
 646 agent per scenario.

Scenario	N. of agents	Comm. delay (milliseconds)	Sim. time (seconds)	DF scheduler	DF server	Agent local scheduler	Agent server	Negotiation protocol	Contractor heuristic	Award heuristic
S1	4	10~50	200	EDF	CBS	EDF	CBS	RBN	ALL	BEST
S2	7	10~50	200	EDF	CBS	EDF	CBS	RBN	ALL	BEST
S3rt	7	10~50	200	EDF	CBS	EDF	CBS	RBN	ALL	BEST
S3gp	7	10~50	200	FIFO	-	FIFO	-	CNET	ALL	BEST

**Table 1.** General setup of the scenarios.



**Figure 8.** Simulated configurations for S1, S2 and S3. S1 represents a single patient multi-agent scenario, S2 a simultaneous two-patient scenario, and S3 a multi-agent scenario with multiple coordinating body-sensors.

647 In this study, we have used simple behaviors, which for simplicity are referred as tasks.  
 648 In particular, we have employed both periodic tasks (recurring after a constant period) and  
 649 aperiodic/sporadic task models (the arrival of the task is not predictable) [46]. Every periodic task is  
 650 characterized by a *task id*, the agent executor (*Ex*), the agent demander (*Dm*), the worst case of their  
 651 computational time (*C*), the release time (*R*), a period (*T*), a relative deadline (*D*), first activation (*f.R.*),  
 652 last activation (*l.R.*), a flag indicating if that task is public (*Pub*) — meaning that the agent is willing to  
 653 perform it on demand, and a flag indicating if such a task is just part of the agent’s knowledge or if it  
 654 part of its running task-set. Moreover, if the task is aperiodic, *T* cannot be applied. Nevertheless, it can  
 655 be indicated, if any, a number of executions (*n*) and an associated server to bound its execution [46]. In  
 656 particular, *Task id*, *Ex*, *Dm*, and *S* are integers IDs; *C* is measured in CPU-clock cycles — normalized in  
 657 instances of time measured seconds [46]; and *R*, *T*, *f.R.*, and *l.R.* are instances of time — expressed in  
 658 seconds.

659 To enable the agents to demand a task execution, we implemented the concept of *need*. Such a  
 660 need is characterized by id of the agent willing to require for its execution, a need id, a release time  
 661 (*R*), the duration of the bidding window to negotiate its execution (*W*), a starting time to begin the  
 662 task execution (*RT*) and its finishing time (*TD*), a number (*n*) of executions (if not periodic), and a  
 663 max (*maxT*) and min (*minT*) period to execute the demanded task (used only in specific conditions  
 664 specified at negotiation time). In particular, *Need id*, *Agent id*, and *task(s)* are integer IDs; and *R*, *W*,  
 665 *TR*, *TD*, *MinT*, and *MaxT* are expressed in seconds. Although the configurations and the workloads  
 666 have been changed among the scenarios studied, the semantics of the task has been kept the same (see  
 667 Table 2). The needs generation depends on the specific application scenario (e.g., the need for a specific  
 668 inertial information given a particular rehabilitating joint.

669 Fostering fairness, the server handling common aperiodic communication tasks (e.g., read and  
 670 write messages) have been kept uniform for all the agents among all the scenarios, and are characterized  
 671 as shown in Table 3

Task	Behavior
$\tau_0$	Kernel task
$\tau_1$	read message
$\tau_2$	write message
$\tau_3$	compute inertial information
$\tau_4$	compute inertial information
$\tau_5$	display graphical information
$\tau_6$	synchronization task
$\tau_7$	on-board data elaboration
$\tau_8$	MIDI signal reproduction

**Table 2.** Task descriptions.

Server id	Agent	Budget	Period	Type	Task(s) served
S100	all	1	10	CBS	$\tau_1$
S200	all	1	10	CBS	$\tau_2$

**Table 3.** Common server characterization.

## 6.1.1. Scenario S1

As shown in Figure 8(a), the mapping agent - device is the following:  $agent_0 \rightarrow$  physiotherapist device,  $agent_1 \rightarrow$  patient device,  $agent_2 \rightarrow$  femur sensor,  $agent_3 \rightarrow$  tibia sensor. In this scenario, the kernel tasks have been setup uniform among the devices of the same type. For example the kernel tasks  $\tau_0$  are the same for both physiotherapist and patient device, and the same among femur and tibia sensor. The complete characterization of the task-sets employed in this scenario is detailed in Table 4. Table 5 details the needs used to generated the dynamics represented in Figure 8(a). In particular, given the agents and tasks distribution of this scenario,  $agent_0$  needs the information from the participant (in this case, the data are made available by the  $agent_1$ , the agent executing on the tablet/smartphone of the only participating patient). In turn,  $agent_1$  needs the inertial information from the wearable sensors (to compute the aggregated plots) — thus it will ask them to  $agent_2$  and  $agent_3$ .

Agent id	Task id	Ex	Dm	C	R	T	D	n	f.R	l.R	S	Pub	Act
0	0	0	0	3	0	20	20	-1	-	-	-	X	✓
	1	0	0	1	-	-	-	-	-	-	S100	X	✓
	2	0	0	1	-	-	-	-	-	-	S200	X	✓
1	0	1	1	3	0	20	20	-	-	-	-	X	✓
	1	1	1	1	-	-	-	-	-	-	S100	X	✓
	2	1	1	1	-	-	-	-	-	-	S200	X	✓
	5	1	-	4	-	20	20	-	-	-	-	✓	X
2	0	2	2	2	0	15	15	-	-	-	-	X	✓
	1	2	2	1	-	-	-	-	-	-	S100	X	✓
	2	2	2	1	-	-	-	-	-	-	S200	X	✓
	3	2	-	4	-	20	20	-	-	-	-	✓	X
3	0	3	3	2	0	15	15	-	-	-	-	X	✓
	1	3	3	1	-	-	-	-	-	-	S100	X	✓
	2	3	3	1	-	-	-	-	-	-	S200	X	✓
	3	3	-	4	-	20	20	-	-	-	-	✓	X
	4	3	-	4	-	20	20	-	-	-	-	✓	X

Table 4. Agents' task-set for Scenario S1.

Agent id	Need id	R	W	TR	TD	n	MinT	MaxT	Task(s)
0	0	4	10	70	200	-	20	25	5
1	0	5	10	60	200	-	20	20	3
	1	5	10	60	200	-	20	20	4

Table 5. Agents' needs for Scenario S1.

It is worth to recall that a task-set  $\Gamma^j$  of a given agent  $j$  is feasible if its utilization factor is less (or equal) than the *least upper bound*<sup>4</sup> ( $\leq U_{lub}$ ) of its scheduling algorithm [46]. The utilization factor  $U_k$  of a single task  $\tau_k$  is computed dividing its computation time  $C_k$  by its period  $T_k$ <sup>5</sup>:  $U_k = \frac{C_k}{T_k}$ . Therefore, the utilization factor of a given agent  $a_j$  at a given time  $t$  is defined by

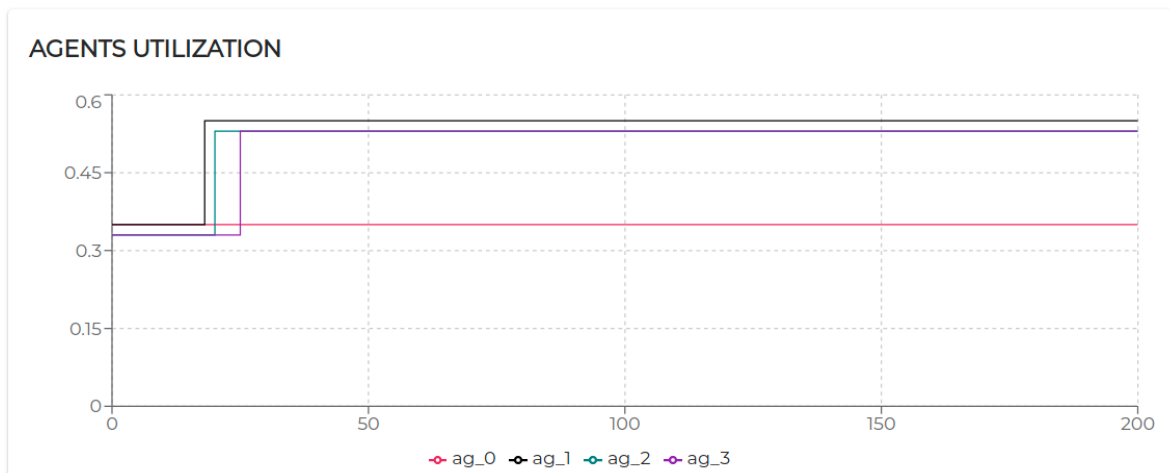
$$U^j(t) = \sum_{\tau_k \in \Gamma^j(t)} U_k \quad (2)$$

Figure 9 shows the trend of the utilization factors of all the agents taking part in S1. It is possible to notice that the utilization of each agent does not exceed the upper bound ( $U_{lub} \leq 1$ ) defined for the tested scheduling algorithm (EDF).

All the negotiated needs (see Table 5) have been accepted. Hence, as visible in Figure 9,  $agent_2$ ,  $agent_3$ , and  $agent_4$  increase their utilization  $U$  at a certain point. In particular, looking at Figure 10, it is possible to see that the basic utilization of  $agent_2$  is  $U = U_{\tau_1} + U_{S100} + U_{S200} = (3/20) + (1/10) + (1/10) = 0.35$ . At  $t = 8.06s$ ,  $agent_2$  negotiates the execution of the task  $\tau_5$ . According to the RBN

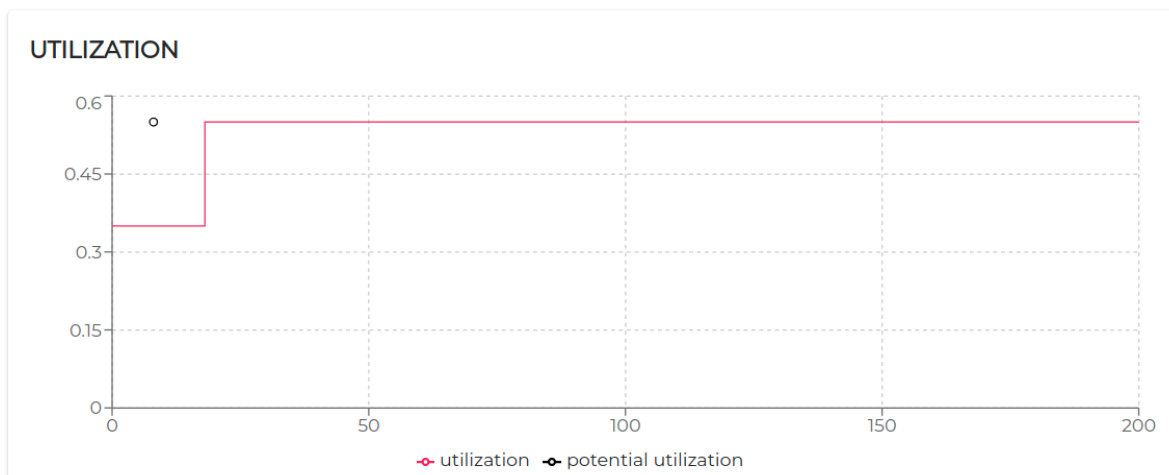
<sup>4</sup> For example, in the case of algorithms such as EDF and CBS  $U_{lub} = 1$ .

<sup>5</sup> In the case where the period  $T_k$  and deadline  $D_k$  are equal.



**Figure 9.** Agens utilization over the simulated time [0 - 200s] in scenario S1. *y-axis*: utilization (adimensional); *x-axis*: time (seconds).


694 protocol, the acceptance of a new task is subject to the schedulability test [49]. In this case, adding the  
 695  $\tau_5$  to its task-set would bring its potential utilization factor to  $U_{pot} = 0.55$  (see black circle in Figure 12).  
 696 The negotiation ends with *agent 0* awarding the execution of  $\tau_5$  to *agent 1*. When *agent 1* receives such a  
 697 communication turns its  $U_{pot}$  in actual  $U$  (see  $t = 18.06s$  in Figure 10).




**Figure 10.** Utilization factor of *agent 2* [0-200s] in scenario S1. *y-axis*: utilization (adimensional); *x-axis*: time (seconds).

698 Figure 11 shows the response time for all the tasks executed by *agent 2*. It is possible to notice that  
 699 the design tested in this scenario generates linear response time with negligible exceptions due to a  
 700 few moments of intense message exchange.

701 Moreover, starting from  $t = 67s$  (see Figure 12), it is possible to see the timely execution of  $\tau_1$   
 702 (twice, reading the inertial positions shared by *agent 2* and *agent 3*) and  $\tau_5$  plotting those information  
 703 on the screen of the used device.



**Figure 11.** Response time of tasks performed by *agent 2* in scenario S1. *y-axis*: response time (seconds); *x-axis*: time (seconds).



**Figure 12.** Graphical representation of the local scheduler of *agent 2* in scenario S1, [0-200s]. Due to the length of the selected period, the timeline is depicted in three lines. The task  $\tau_5$  (in orange) timely displays information after the inertial positions were reported from the agent sensors.

## 704 6.1.2. Scenario S2

705 To implement the scenario shown in Figure 8(b), the mapping agent - device is the following:  
 706 *agent0* → physiotherapist device, *agent1* → patient 1 device, *agent2* → femur sensor (patient 1), *agent3*  
 707 → tibia sensor (patient 1), *agent4* → patient 2 device, *agent5* → femur sensor (patient 2), and *agent6* →  
 708 tibia sensor (patient 2).

709 To replicate the same conditions per patient, even in S2, the kernel tasks have been setup uniform  
 710 among the devices of the same type (see Table 6). The complete characterization of the task-sets  
 711 employed in this scenario is detailed in Table 6. Table 7 details the needs used to generated the  
 712 dynamics represented in Figure 8(b). In particular, in this scenario, we have two patients (with two  
 713 sensors and one tablet/smartphone each) and one physiotherapist. Therefore, *agent0* will demand  
 714 aggregated information to *agent1* and *agent4*, which, in turn, will demand inertial information to  
 715 respectively *agent2* and 3 and *agent4* and 5.

Agent id	Task id	Ex	Dm	C	R	T	D	n	f.R	l.R	S	Pub	Act
0	0	0	0	3	0	20	20	-1	-	-	-	X	✓
	1	0	0	1	-	-	-	-	-	-	S100	X	✓
	2	0	0	1	-	-	-	-	-	-	S200	X	✓
1	0	1	1	3	0	20	20	-	-	-	-	X	✓
	1	1	1	1	-	-	-	-	-	-	S100	X	✓
	2	1	1	1	-	-	-	-	-	-	S200	X	✓
	5	1	-	4	-	20	20	-	-	-	-	✓	X
2	0	2	2	2	0	15	15	-	-	-	-	X	✓
	1	2	2	1	-	-	-	-	-	-	S100	X	✓
	2	2	2	1	-	-	-	-	-	-	S200	X	✓
	3	2	-	4	-	20	20	-	-	-	-	✓	X
3	0	3	3	2	0	15	15	-	-	-	-	X	✓
	1	3	3	1	-	-	-	-	-	-	S100	X	✓
	2	3	3	1	-	-	-	-	-	-	S200	X	✓
	4	3	-	4	-	20	20	-	-	-	-	✓	X
4	0	4	4	3	0	20	20	-	-	-	-	X	✓
	1	4	4	1	-	-	-	-	-	-	S100	X	✓
	2	4	4	1	-	-	-	-	-	-	S200	X	✓
	6	4	-	4	-	20	20	-	-	-	-	✓	X
5	0	5	5	2	0	15	15	-	-	-	-	X	✓
	1	5	5	1	-	-	-	-	-	-	S100	X	✓
	2	5	5	1	-	-	-	-	-	-	S200	X	✓
	7	5	-	4	-	20	20	-	-	-	-	✓	X
6	0	6	6	2	0	15	15	-	-	-	-	X	✓
	1	6	6	1	-	-	-	-	-	-	S100	X	✓
	2	6	6	1	-	-	-	-	-	-	S200	X	✓
	8	6	-	4	-	20	20	-	-	-	-	✓	X


Table 6. Agents' task-set for Scenario S2.

Agent id	Need id	R	W	TR	TD	n	MinT	MaxT	Task(s)
0	0	4	10	70	200	-	20	25	5
	1	4	10	70	200	-	20	25	6
1	0	5	10	60	200	-	20	20	3
	1	5	10	60	200	-	20	20	4
4	0	5	10	60	200	-	20	25	7
	1	5	10	60	200	-	20	25	8

Table 7. Agents' needs for Scenario S2.


716 The Scenario S2 is based on S1, doubling the number of participants (and therefore the number  
 717 of sensors). On the one hand, the behaviors characterizing patient 1 and his/her device and sensors  
 718 (*agent 1*, *agent 2*, and *agent 3*) and patient 2 and his/her device and sensors (*agent 4*, *agent 5*, and *agent 6*)  
 719 remained unaltered. On the other hand, the sole actor potentially affected by the growth of patients is  
 720 the physiotherapist, therefore *agent 0*.

721 Hence, although increasing the number of patients does not impact on the response time of the  
 722 kernel task of *agent 0* (see Figure 14b) the response time to process the incoming messages has already



**Figure 13.** Agents Utilization over the simulated time [0 - 200s] in scenario S2. *y-axis*: utilization (adimensional); *x-axis*: time (seconds).

723 been affected, in some cases doubling its value. Figure 14a shows the response time of the *agent 0* in S1  
 724 and Figure 14b shows its response time in S2. While *agent 0* is still respecting strict timing constraints,  
 725 the demand for more stringent performance may, at a certain point, require to adapt the system design  
 726 to the scale of the application domain.



(a) Response time of *agent 0*'s tasks in S1.

(b) Response time of *agent 0*'s tasks in S2.

**Figure 14.** Response times in scenarios S1 and S2. *y-axis*: response time (seconds); *x-axis*: time (seconds).

### 727 6.1.3. Scenario S3rt

728 To simulate behaviors and conditions represented in Figure 8(c), the mapping agent - device is the  
 729 following: *agent0* → physiotherapist device, *agent1* → patient device, *agent2* → right arm sensor, *agent3*  
 730 → chest sensor, *agent4* → left arm sensor, *agent5* → right femur sensor, and *agent6* → left femur sensor.  
 731 To simulate sensor heterogeneity, in S3rt, it has been assumed different workloads for the several  
 732 kernel and specific tasks. The complete characterization of the task-sets employed in this scenario is  
 733 detailed in Table 8. Table 9 details the needs used to generated the dynamics represented in Figure 8(c).  
 734 In particular, *agent0* demands for the usual aggregated data, to the only participant (*agent1*), which, in  
 735 turn, demands inertial information to all the distributed wearable sensors to compute the complex  
 736 kinematics of the motor tasks targeted in this scenario.

737 The execution of the tasks and needs listed above produces the utilization factors plotted in  
 738 Figure 15. As we can see, beside *agent 0* and *agent 1* which remained unaltered, the dynamics generated  
 739 by the release of the needs affected the rest of the agents (especially given their higher utilization  
 740 factors).

741 Indeed, in S3rt, only part of the need have been satisfied. In particular, *agent 1* and *agent 5* recorded  
 742 100% of acceptance and *agent 3* and *agent 6* accepted only 50% of the demanded tasks. To understand  
 743 such behavior, let us look at Figure 16. *Agent 3* has positively answered (bided) to the execution of  
 744 a task, which would bring its utilization to  $U = 0.53$ . Before the confirmation (award) of such a bid,

Agent id	Task id	Ex	Dm	C	R	T	D	n	f.R	l.R	S	Pub	Act
0	0	0	0	3	0	20	20	-1	-	-	-	X	✓
	1	0	0	1	-	-	-	-	-	-	S100	X	✓
	2	0	0	1	-	-	-	-	-	-	S200	X	✓
1	0	1	1	3	0	20	20	-	-	-	-	X	✓
	1	1	1	1	-	-	-	-	-	-	S100	X	✓
	2	1	1	1	-	-	-	-	-	-	S200	X	✓
	5	1	-	4	-	20	20	-	-	-	-	✓	X
2	0	2	2	2	0	15	15	-	-	-	-	X	✓
	1	2	2	1	-	-	-	-	-	-	S100	X	✓
	2	2	2	1	-	-	-	-	-	-	S200	X	✓
	3	2	-	4	-	20	20	-	-	-	-	✓	X
3	0	3	3	2	0	15	15	-	-	-	-	X	✓
	1	3	3	1	-	-	-	-	-	-	S100	X	✓
	2	3	3	1	-	-	-	-	-	-	S200	X	✓
	4	3	-	4	-	20	20	-	-	-	-	✓	X
	8	3	-	11	-	16	16	-	-	-	-	✓	X
4	0	4	4	6	0	11	11	-	-	-	-	X	✓
	1	4	4	1	-	-	-	-	-	-	S100	X	✓
	2	4	4	1	-	-	-	-	-	-	S200	X	✓
	7	4	-	15	-	16	16	-	-	-	-	✓	X
5	0	5	5	3	0	15	15	-	-	-	-	X	✓
	1	5	5	1	-	-	-	-	-	-	S100	X	✓
	2	5	5	1	-	-	-	-	-	-	S200	X	✓
	3	5	-	6	-	23	23	-	-	-	-	✓	X
6	0	6	6	4	0	15	15	-	-	-	-	X	✓
	1	6	6	1	-	-	-	-	-	-	S100	X	✓
	2	6	6	1	-	-	-	-	-	-	S200	X	✓
	6	6	-	9	-	20	20	-	-	-	-	✓	X

Table 8. Agents' task-set for Scenario S3rt.

Agent id	Need id	R	W	TR	TD	n	MinT	MaxT	Task(s)
0	0	4	10	70	200	-	20	25	5
1	0	5	35	60	200	-	20	20	3
	1	5	35	60	200	-	20	25	4
3	0	6	10	60	200	-	20	20	3
	1	10	10	80	200	-	20	20	6
4	0	15	10	80	200	-	20	25	6
5	0	25	10	100	200	-	16	16	8
6	0	35	10	90	200	-	20	20	7

Table 9. Agents' needs for Scenario S3rt.


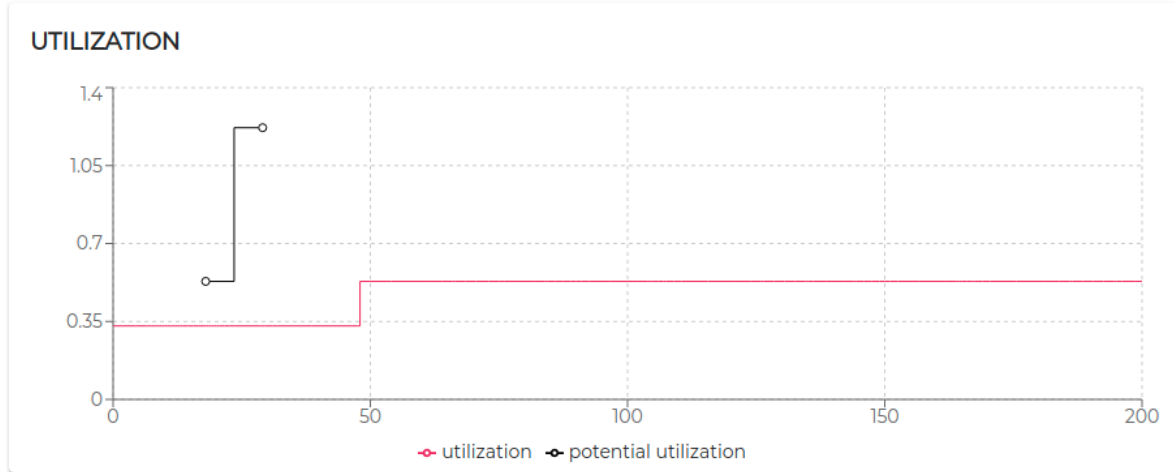


Figure 15. Agents Utilization over the simulated time [0 - 200s] in scenario S3rt. *y-axis*: utilization (adimensional); *x-axis*: time (seconds).

745 *agent 3* receives a second request. According to the RBN protocol, an agent can accept the execution  
 746 of a given task only if it can allocate it (without overcoming its maximum utilization factor). In this




747 case, the schedulability test performed at  $t = 29.06s$  to verify the possible allocation of the second  
 748 negotiated tasks brings its utilization at  $U = 1.22$ . Therefore, the bid for such a task has been negative  
 749 (rejected). In turn, at  $t = 48s$ , the first bid has been awarded. Thus, the potential utilization factor turns  
 750 into effective utilization.



**Figure 16.** Agent 3 Utilization over the simulated time [0 - 200s] in scenario S3rt. *y-axis*: utilization (adimensional); *x-axis*: time (seconds).

751 Finally, let us analyze *agent 5* to see how it is performing with the utilization is  $U = 1$ , the  
 752 theoretical maximum to still ensure predictability.



**(a)** Utilization over 200s of the agent 5 in S3rt. *y-axis*: utilization (adimensional); *x-axis*: time (seconds). **(b)** Response time over 200s of the agent 5 in S3gp. *y-axis*: response time (seconds); *x-axis*: time (seconds).

**Figure 17.** Utilization and response time of agent 5 in S3rt.


753 Initially, its utilization is  $U = 0.4$  (due to running  $\tau_0$ ,  $\tau_1$ , and  $\tau_2$ ). At  $t = 10.07s$  it receives a first  
 754 request for the execution of  $\tau_3$ , which, if accepted, would raise its utilization to  $U = 0.7$ . Right after, at  
 755  $t = 12.07s$ , it receives a second request for executing  $\tau_3$ . At  $t = 12.07s$ , *agent 5* has not received yet an  
 756 answer (award/rejection) for the previous bid. Thus, besides its actual utilization is still  $U = 0.4$ , it has  
 757 to consider its potential utilization  $U_{pot} = 0.7$  to perform the schedulability test. Bidding positively to  
 758 both the requests would bring its potential utilization to  $U_{pot} = 1$ . So, it bids positively once again. At  
 759  $t = 21.12s$  and  $t = 45.01s$ , it gets respectively awarded its two bids. Therefore, as visible in Figure 17a  
 760 its  $U_{pot}$  turns into  $U$ .

761 Although operating on the edge of its capabilities, *agent 5* does not record any deadline miss  
 762 during the entire simulation. Nevertheless, high variability in the response time can be acknowledged  
 763 (see Figure 17b). To refine the response time, adjustments on the initial design might be required. The  
 764 MAXIM-GPRT tool can be a valuable support in such a process.

#### 765 6.1.4. Scenario S3gp


766 In the scenario S3gr, it has been used the same task-set of used in S3rt (see Table 8) and the same  
 767 needs distribution and characterization (see Table 9).

768 However, considering the general-purpose nature of the underlying mechanisms (FIFO as agent local  
 769 scheduler and CNET as negotiation protocol), no timing guarantee can be enforced nor predicted [13,  
 770 49,54]. Hence, over 200 seconds of simulation, *agent 3* misses 8% and *agent 4* misses 52% of their  
 771 deadlines. Figure 18 shows that *agent 3* at  $t = 40s$  refuses to execute a task that would raise its  
 772 utilization at  $U = 1.5$  (which would have entailed unpredictable consequences).




(a) Utilization over 200s of the *agent 4* in S3rt (b) Utilization over 200s of the *agent 4* in S3gp.  
**Figure 18.** Comparison among the utilization of *agent 4* in S3rt and S3gp. *y-axes:* utilization (adimensional); *x-axes:* time (seconds).

773 The combination of FIFO and CNET lacks mechanisms typical of real-time systems crucial to  
 774 handle workloads, deadlines, and strict timing constraints. Hence, in S3gp, *agent 4* accepts to execute  
 775 the demanded tasks. As a result, both the tasks executed by *agent 4* record deadline miss. Clearly, in the  
 776 scenario S3rt, *agent 4* has a more conservative (refusing to execute the demanded task), Conversely, in  
 777 S3gp *agent 4* is more flexible (accepting the demanded task), which, however, resulted in compromising  
 778 its predictability and reliability. To better understand the visibly different response time provided by  
 779 *agent 4* in the two tested scenarios, Figure 19 shows the performances recorded.



(a) Response time over 200s of the *agent 4* in S3rt. (b) Response time over 200s of the *agent 4* in S3gp.  
**Figure 19.** Comparison among the response time of *agent 4* in S3rt and S3gp. *y-axes:* response time (seconds); *x-axes:* time (seconds).

780 Although less, the flexibility of the general-purpose algorithms cost to the *agent 3* to record 8%  
 781 of deadline miss. The comparison of the response time among scenario S3rt and S3gp is shown in  
 782 Figure 20.



(a) Response time over 200s of the *agent 3* in S3rt (b) Response time over 200s of the *agent 3* in S3gp.  
**Figure 20.** Comparison among the response time of *agent 4* in S3rt and S3gp. *y-axes:* response time (seconds); *x-axes:* time (seconds).

## 783 7. Implementation & experimentation with RSP agents

784 We have implemented the RSP agent architecture as a library available in Scala. The code is  
785 open-source, and it is available in Github<sup>6</sup>. The core of the RSP agents implementation is written  
786 using the Akka Actors library<sup>7</sup>. Akka provides the essential programming abstractions to create  
787 actors, providing message dispatching, remoting, actor hierarchies, and other features. The RSP  
788 agents implementation defines *traits* (analogous to interfaces in Java and other languages) for its  
789 main types of objects. For instance, the `StreamReceiver` trait implements the receiver agent described  
790 in Section 5.2.1. These traits are independent of the communication layer, i.e., it allows plugging  
791 different types of channel implementations, such as MQTT or WebSocket. Additional modules can be  
792 plugged into the architecture (e.g., for concrete implementations of specific RSP engines). In our initial  
793 implementation, we have focused on using CQELS as underlying RDF stream processors, although we  
794 have also implemented classes for C-SPARQL and TrOWL. To allow the integration with these existing  
795 engines, it suffices that they provide a JVM-compatible API. RSP agent traits make use of abstract  
796 methods that need to be implemented for any specific extension. For example, the `StreamReceiver`  
797 trait defines abstract methods that allow: feeding an RDF stream with graph (`consumeGraph`), register a  
798 query (`query`), push data results (`push`), and terminate push and clean resources (`terminatePush`).

799 In the remainder of this section, we present a set of experimental results of the implementation  
800 of RSP agents. The goal of these experiments is to show how the agent architecture implemented in  
801 the library performs under different configurations. By changing the number of senders/receivers,  
802 concurrent operations, rates of streaming data flow, etc., we present various scenarios which could be  
803 implemented in a digital physiotherapy use-cases. Although in single patient scenarios, the number of  
804 senders and receivers would be usually low, when dealing with larger numbers of simultaneous users,  
805 the number of required agents will also increase.

806 In the following experiments, the main metric is the throughput, measured in terms of efficiency  
807 (i.e., the rate between the actual number of RDF stream elements processed per unit of time, and  
808 the maximum ideal number of processed elements). The choice of this metric is based first on the  
809 need to assess the behavior of the platform to different conditions of the input streams (e.g., number  
810 of streams/senders, the velocity of the streams, number of parallel processors). The usage of a rate  
811 indicator is due to the fact that an absolute throughput is clearly variable depending on the input  
812 stream characteristics. Therefore the efficiency rate provides a normalized parameter. Queries and  
813 data have been adapted from SRBench [56], using an upgraded version of the datasets, using the new  
814 version of the SSN Ontology<sup>8</sup>, and a synthetic generation stream feeder. The original data consists  
815 of sensor observations extracted from the LinkedSensorData [57] initiative, based on observations  
816 collected since 2002 by 20K sensor stations. There are typically five sensors per station, i.e. a total of  
817 around 100,000 sensors in the data set. The sensors measure phenomena such as temperature, position,  
818 visibility, pressure, etc. Irrespective of the application domain, the experiments described in this  
819 section show the feasibility of the stream reasoning agents architecture in a concrete implementation  
820 All experiments were run on Ubuntu 16.04 LTS, Intel Core i7-7700U (3.60GHz, 8MB cache, Quad-Core).

821 In the first set of experiments, we measured the throughput efficiency, for different input stream  
822 rates (1,10,10, and 1000 graphs/s), and a different number of concurrent senders, and a single receiver  
823 (Figure 21a).


824 Clearly, the efficiency decreases considerably as the input stream increases. A drastic increase  
825 produces a significant drop in efficiency, either if it is by increasing the number of senders or the input  
826 rate. This is basically due to the limitations of CQELS as the underlying engine. The next experiment is  
827 set in exactly the same conditions, except that it uses five concurrent stream receiver agents instead of

---

<sup>6</sup> <https://github.com/jpcik/ldn-streams>

<sup>7</sup> <http://akka.io>


<sup>8</sup> <https://www.w3.org/TR/vocab-ssn/>



**(a)** Throughput efficiency: one receiver. **(b)** Throughput efficiency: five receivers.  
**Figure 21.** Throughput efficiency vs. input stream rates, for different sets of concurrent senders.

828 only 1 (Figure 21b). As can be seen, using more receiver agents already provides a higher throughput  
 829 efficiency for a larger number of cases.

830 The next experiment provides more information on how a set of CQELS engines running as  
 831 RSP agents can handle a total of 10K concurrent senders, each spitting one graph per second. The  
 832 experiment is set for 1, 5, 10, 20, 50, and 100 concurrent CQELS agent receivers. As can be seen  
 833 in Figure 22a, with 5, 10, and 20 concurrent senders, there is a considerable improvement in the  
 834 throughput efficiency. However, increasing even more receivers produces a sustained decrease, as the  
 835 CPU is not able to scale on its own to that number of engines. A distributed deployment would be  
 836 required to scale in that case.




**(a)** Throughput efficiency: 10K concurrent senders. **(b)** Throughput efficiency: different input stream rates.

**Figure 22.** Throughput efficiency vs. number of RSP senders.


837 The next experiment shows how RSP stream receivers (CQELS engine) respond to a different  
 838 number of senders (1, 10, 100, 1000, 1000). It shows results for input streams of 1, 10, 100, 1000 graphs/s  
 839 (Figure 22b). The graph shows the progressive efficiency decrease as the input stream (combined with  
 840 the number of senders) increases. For instance, for 10K senders, and 1K graphs/s, the input load is of  
 841 10M graphs/s, which is too much for a single receiver instance.

842 In the next experiment, we evaluate a similar scenario, but this time, adding CQELS RSP instances  
 843 (1, 2, 5, 8, and 10 concurrent engines). The results are shown for different numbers of RSP concurrent  
 844 senders (2K, 4K, 8K, 10K, and 16K). For very high input loads, the system is still capable of at least 0.5  
 845 efficiency. It is clear that at this point, a cluster deployment would be required.

846 The final experiment was performed only for RSP agents ingesting but not processing data. It  
 847 shows results for a different number of RSP agents ingesting streams (1, 5, 10, 20, 50, and 100 concurrent  
 848 agents), and for 100, 1000, and 10000 senders. As can be seen, for 1K and 10K adding additional  
 849 resources in general increases the overall efficiency (Figure 24).



**Figure 23.** Throughput efficiency vs. number of RSP senders, for different numbers of RSP receiver engines.



**Figure 24.** Throughput efficiency vs. number of RSP receivers without processing, for different sets of concurrent senders.

## 850 8. Discussion & Conclusions

851 This paper proposes a novel approach for enabling decentralized stream processing in digital  
 852 rehabilitation by proposing a stream processing agent-based architecture that enforces real-time  
 853 constraints. The idea behind these decentralized stream processing agents is that they are capable  
 854 of sharing not only streaming data, but also processing duties, using collaboration and negotiation  
 855 protocols, while relying on common vocabularies/ontologies that consider the high dynamicity of  
 856 their beliefs, state, goals, and behavior [9,58]. These features are essential to provide highly responsive  
 857 feedback and accurate data analytics, which are required in digital physiotherapy.

858

859

Compared to the state of the art, this work provides the following contributions:

860

861

862

863

864

865

866

867

868

869

870

871

872

873

874

875

- 860 • *Stream reasoning agents model.* Beyond existing approaches in stream reasoning or RDF stream  
 861 processing (as seen in Section 2), the proposed model does not only focus on complex processing  
 862 algorithms and methods over semantic streams, but also on the autonomous cooperation among  
 863 agents that produce and consume those streams, following the vision described in [9].
- 864 • *Real-time compliance for RSP agents.* Existing RSP systems provide execution models that do  
 865 not support mechanisms for compliance with strict real-time constraints. Filling this gap, our  
 866 proposed model incorporates these constraints at its core, which can be implemented using  
 867 existing strategies as in [12].
- 868 • *Real-time agent simulation results for digital rehabilitation.* The simulation environment presented in  
 869 the paper constitutes an important milestone for modelling and configuring agent-based systems  
 870 for different scenarios, considering strict real-time specifications. While in previous works on  
 871 digital rehabilitation feedback was typically provided on best-effort strategies, these simulations  
 872 provide indication of how and when real-time scheduling strategies can be helpful in order to  
 873 deal with strict timing limitations.
- 874 • *Implementation and evaluation of RSP agents.* The feasibility and behavior of the RSP agents concept  
 875 has been demonstrated in this work, through a concrete implementation that relies on an existing

876 RSP engine. This is a first implementation of the agent model for RDF stream processing, beyond  
877 the centralized systems present in the literature, as seen in Section 2.

878 The model, simulation and implementation presented in this paper also constitute an important  
879 milestone towards the adoption of agent-based technologies for real-time sensing applications. The  
880 results of this research work open several opportunities, even if there are some limitations that we  
881 need to consider, as explained below.

882

883 *Opportunities* - Based on the principles of stream reasoning and RDF stream processing, the proposed  
884 model incorporates the flexibility of autonomous organization of streaming agents, with the capability  
885 of defining *strict* deadlines for agent behaviors. This feature is a fundamental advantage for digital  
886 physiotherapy, as it enables reliable *in-time* feedback among sensors and/or eHealth applications for  
887 patient support. Moreover, sharing common ontology models and underlying RT-MAS mechanisms,  
888 additional sensors and devices can be easily plugged. The approach proposed in this paper  
889 relies upon and extends previous works on stream reasoning, also including the representation of  
890 heterogeneous data streams as dynamic knowledge graphs on which complex-event processing (CEP)  
891 and inductive/deductive reasoning can be applied. This feature addresses the challenges related to  
892 sensor and agent heterogeneity, relying on standards for representing agent negotiation protocols  
893 and sensor metadata. The simulation results developed in this paper provide indication that this  
894 approach can have deep impact in digital rehabilitation scenarios, allowing the self-configuration of  
895 decentralized sensor solutions. Complementary to these results, the evaluation performed on the RDF  
896 stream processing agent implementation, provides evidence of the feasibility of allowing agent-based  
897 interactions among sensing devices. Finally, the simulator employed in this study might be a strategic  
898 tool for future system design and setup, before including the human in the loop.

899

900 *Limitations* - The simulation scenarios tested in this study have relied on synthetic data generated by  
901 the execution environment. In real-world applications, getting such information might require further  
902 tasks such as a more complex signal processing and more complex agent interactions (depending  
903 on the kinematic chain of a given motor exercise), which might entail more complex semantic  
904 representations of the exchanged information. Moreover, many real sensors are still unable to “run”  
905 RT-agents (due to the lack of a proper RT-MAS framework for embedded systems). Operating in  
906 safety-critical conditions, the development of 3rd-party hardware and software might require a longer  
907 developing time, thus slowing down the adoption of the proposed solution. **Another important aspect  
908 to consider refers to the fact that in this work we do not handle uncertainty of both streaming and  
909 static knowledge. Agents beliefs may have different levels of uncertainty, for which techniques such as  
910 fuzzy multi-criteria decision-making [59]. Moreover, RSP agents may require to adopt strategies for  
911 scheduling streaming task under uncertainty conditions [60], or rely on discrepancy measures in case  
912 of disagreements [61]**

913

914 In terms of impact, this approach may constitute a first step towards a more decentralized  
915 understanding of how IoT devices can be used for supporting eHealth applications. Particularly in  
916 digital physiotherapy, it would be important to explore the challenges of deploying stream processing  
917 agents in clinical environments. Furthermore, it will be crucial to study how real-time constraints  
918 might be included as extensions of RDF validation languages such as SHACL [62].

919 **Author Contributions:** Conceptualization, J.P.C. and D.C.; methodology, J.P.C.; software, D.C.; validation, J.P.C.  
920 and D.C.; formal analysis, J.P.C.; investigation, J.P.C. and D.C.; resources, D.C.; writing—original draft preparation,  
921 J.P.C. and D.C.; writing—review and editing, D.C. and J.P.C.

922 **Acknowledgments:** Thanks to Giuseppe Albanese for the support with the simulation environment.

923 **Conflicts of Interest:** The authors declare no conflict of interest.

924 **References**

- 925 1. Acree, L.S.; Longfors, J.; Fjeldstad, A.S.; Fjeldstad, C.; Schank, B.; Nickel, K.J.; Montgomery, P.S.; Gardner,  
926 A.W. Physical activity is related to quality of life in older adults. *Health and quality of life outcomes* **2006**,  
927 *4*, 37.
- 928 2. Elavsky, S.; McAuley, E.; Motl, R.W.; Konopack, J.F.; Marquez, D.X.; Hu, L.; Jerome, G.J.; Diener, E. Physical  
929 activity enhances long-term quality of life in older adults: Efficacy, esteem, and affective influences. *Annals*  
930 *of Behavioral Medicine* **2005**, *30*, 138–145.
- 931 3. Faber, M.J.; Bosscher, R.J.; Paw, M.J.C.A.; van Wieringen, P.C. Effects of exercise programs on falls and  
932 mobility in frail and pre-frail older adults: a multicenter randomized controlled trial. *Archives of physical*  
933 *medicine and rehabilitation* **2006**, *87*, 885–896.
- 934 4. Daley, M.J.; Spinks, W.L. Exercise, mobility and aging. *Sports medicine* **2000**, *29*, 1–12.
- 935 5. Vissers, M.M.; Bussmann, J.B.; Verhaar, J.A.; Arends, L.R.; Furlan, A.D.; Reijman, M. Recovery of physical  
936 functioning after total hip arthroplasty: systematic review and meta-analysis of the literature. *Physical*  
937 *therapy* **2011**, *91*, 615–629.
- 938 6. Speck, R.M.; Courneya, K.S.; Mâsse, L.C.; Duval, S.; Schmitz, K.H. An update of controlled physical  
939 activity trials in cancer survivors: a systematic review and meta-analysis. *Journal of Cancer Survivorship*  
940 **2010**, *4*, 87–100.
- 941 7. Hugues, A.; Di Marco, J.; Janiaud, P.; Xue, Y.; Pires, J.; Khademi, H.; Cucherat, M.; Bonan, I.; Gueyffier, F.;  
942 Rode, G. Efficiency of physical therapy on postural imbalance after stroke: study protocol for a systematic  
943 review and meta-analysis. *BMJ open* **2017**, *7*, e013348.
- 944 8. Reisdorf, B.C.; Rikard, R. Digital rehabilitation: a model of reentry into the digital age. *American Behavioral*  
945 *Scientist* **2018**, *62*, 1273–1290.
- 946 9. Tommasini, R.; Calvaresi, D.; Calbimonte, J.P. Stream Reasoning Agents: Blue Sky Ideas Track. Proceedings  
947 of the 18th International Conference on Autonomous Agents and MultiAgent Systems. International  
948 Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 1664–1680.
- 949 10. Della Valle, E.; Ceri, S.; Barbieri, D.F.; Braga, D.; Campi, A. A first step towards stream reasoning. Future  
950 Internet Symposium. Springer, 2008, pp. 72–81.
- 951 11. Dell’Aglio, D.; Le Phuoc, D.; Le-Tuan, A.; Ali, M.I.; Calbimonte, J.P. On a Web of Data Streams. Proc. ISWC  
952 Workshop on Decentralizing the Semantic Web DeSemWeb 2017, 2017.
- 953 12. Calvaresi, D.; Marinoni, M.; Dragoni, A.F.; Hilfiker, R.; Schumacher, M. Real-time multi-agent systems for  
954 telerehabilitation scenarios. *Artificial intelligence in medicine* **2019**, *96*, 217–231.
- 955 13. Calvaresi, D.; Marinoni, M.; Sturm, A.; Schumacher, M.; Buttazzo, G. The challenge of real-time multi-agent  
956 systems for enabling IoT and CPS. Proceedings of the international conference on web intelligence. ACM,  
957 2017, pp. 356–364.
- 958 14. Mikołajewska, E.; Mikołajewski, D. Neurological telerehabilitation—current and potential future  
959 applications. *J. Health Sci* **2011**, *1*, 7–14.
- 960 15. Borel, S.; Schneider, P.; Newman, C. Video analysis software increases the interrater reliability of video gait  
961 assessments in children with cerebral palsy. *Gait & posture* **2011**, *33*, 727–729.
- 962 16. Munro, A.; Herrington, L.; Carolan, M. Reliability of 2-dimensional video assessment of frontal-plane  
963 dynamic knee valgus during common athletic screening tasks. *Journal of sport rehabilitation* **2012**, *21*, 7–11.
- 964 17. Buonocunto, P.; Giantomassi, A.; Marinoni, M.; Calvaresi, D.; Buttazzo, G. A limb tracking platform for  
965 tele-rehabilitation. *ACM Transactions on Cyber-Physical Systems* **2018**, *2*, 30.
- 966 18. Matarić, M.J.; Eriksson, J.; Feil-Seifer, D.J.; Winstein, C.J. Socially assistive robotics for post-stroke  
967 rehabilitation. *Journal of NeuroEngineering and Rehabilitation* **2007**, *4*, 5.
- 968 19. Burgar, C.G.; Lum, P.S.; Shor, P.C.; Van der Loos, H.M. Development of robots for rehabilitation therapy:  
969 The Palo Alto VA/Stanford experience. *Journal of rehabilitation research and development* **2000**, *37*, 663–674.
- 970 20. Zhang, W.; Gao, J.; Shi, B.; Cui, H.; Zhu, H. Health monitoring of rehabilitated concrete bridges using  
971 distributed optical fiber sensing. *Computer-Aided Civil and Infrastructure Engineering* **2006**, *21*, 411–424.
- 972 21. Ferreira, C.; Guimarães, V.; Santos, A.; Sousa, I. Gamification of stroke rehabilitation exercises using a  
973 smartphone. Proceedings of the 8th International Conference on Pervasive Computing Technologies for  
974 Healthcare. ICST (Institute for Computer Sciences, Social-Informatics and . . . , 2014, pp. 282–285.

- 975 22. Alimanova, M.; Borambayeva, S.; Kozhamzharova, D.; Kurmangaiyeva, N.; Ospanova, D.;  
976 Tyulepberdinova, G.; Gaziz, G.; Kassenkhan, A. Gamification of hand rehabilitation process using virtual  
977 reality tools: Using leap motion for hand rehabilitation. 2017 First IEEE International Conference on  
978 Robotic Computing (IRC). IEEE, 2017, pp. 336–339.
- 979 23. Fraile, J.A.; Bajo, J.; Corchado, J.M.; Abraham, A. Applying wearable solutions in dependent environments.  
980 *IEEE Transactions on Information Technology in Biomedicine* **2010**, *14*, 1459–1467.
- 981 24. Bergmann, J.; McGregor, A. Body-worn sensor design: what do patients and clinicians want? *Annals of*  
982 *biomedical engineering* **2011**, *39*, 2299–2312.
- 983 25. Chen, K.H.; Chen, P.C.; Liu, K.C.; Chan, C.T. Wearable sensor-based rehabilitation exercise assessment for  
984 knee osteoarthritis. *Sensors* **2015**, *15*, 4193–4211.
- 985 26. Lee, S.I.; Adans-Dester, C.P.; Grimaldi, M.; Dowling, A.V.; Horak, P.C.; Black-Schaffer, R.M.; Bonato, P.;  
986 Gwin, J.T. Enabling stroke rehabilitation in home and community settings: a wearable sensor-based  
987 approach for upper-limb motor training. *IEEE journal of translational engineering in health and medicine* **2018**,  
988 *6*, 1–11.
- 989 27. Sherrill, D.M.; Moy, M.L.; Reilly, J.J.; Bonato, P. Using hierarchical clustering methods to classify motor  
990 activities of COPD patients from wearable sensor data. *Journal of NeuroEngineering and Rehabilitation* **2005**,  
991 *2*, 16.
- 992 28. Rodriguez, A.C.; Roda, C.; González, P.; Navarro, E. Contextualizing Tasks in Tele-Rehabilitation Systems  
993 for Older People. International Workshop on Ambient Assisted Living. Springer, 2015, pp. 29–41.
- 994 29. Felisberto, F.; Costa, N.; Fdez-Riverola, F.; Pereira, A. Unobstructive Body Area Networks (BAN) for  
995 efficient movement monitoring. *Sensors* **2012**, *12*, 12473–12488.
- 996 30. Mutingi, M.; Mbohwa, C. Developing Multi-agent Systems for mHealth Drug Delivery. In *Mobile Health*;  
997 Springer, 2015; pp. 671–683.
- 998 31. Barbieri, D.F.; Braga, D.; Ceri, S.; Della Valle, E.; Grossniklaus, M. C-sparql: a continuous query language  
999 for rdf data streams. *Intl. J. Semantic Computing* **2010**, *4*, 3–25.
- 1000 32. Calbimonte, J.P.; Jeung, H.; Corcho, O.; Aberer, K. Enabling query technologies for the semantic sensor  
1001 web. *Int. J. Semantic Web Inf. Syst.* **2012**, *8*, 43–63.
- 1002 33. Le-Phuoc, D.; Dao-Tran, M.; Parreira, J.X.; Hauswirth, M. A native and adaptive approach for unified  
1003 processing of linked streams and linked data. In *ISWC*; Springer, 2011; pp. 370–388.
- 1004 34. Komazec, S.; Cerri, D.; Fensel, D. Sparkwave: continuous schema-enhanced pattern matching over RDF  
1005 data streams. In *Proc. 4th ACM International Conference on Distributed Event-Based Systems DEBS*; ACM,  
1006 2012; pp. 58–68.
- 1007 35. Anicic, D.; Fodor, P.; Rudolph, S.; Stojanovic, N. EP-SPARQL: a unified language for event processing and  
1008 stream reasoning. *WWW*, 2011, pp. 635–644.
- 1009 36. Barbieri, D.F.; Della Valle, E. A Proposal for Publishing Data Streams as Linked Data - A Position Paper.  
1010 LDOW, 2010.
- 1011 37. Sequeda, J.F.; Corcho, O. Linked stream data: A position paper. SSN. CEUR-WS. org, 2009, pp. 148–157.
- 1012 38. Balduini, M.; Valle, E.D.; Tommasini, R. SLD Revolution: A Cheaper, Faster yet more Accurate Streaming  
1013 Linked Data Framework. *RSP*, 2017, pp. 1–15.
- 1014 39. Mauri, A.; Calbimonte, J.P.; Dell'Aglio, D.; Balduini, M.; Brambilla, M.; Valle, E.D.; Aberer, K. TripleWave:  
1015 Spreading RDF Streams on the Web. *ISWC*, 2016, pp. 140–149.
- 1016 40. Berners-Lee, T.; Hendler, J.; Lassila, O. The semantic web. *Scientific american* **2001**, *284*.
- 1017 41. Ciorrea, A.; Mayer, S.; Gandon, F.; Boissier, O.; Ricci, A.; Zimmermann, A. A Decade in Hindsight:  
1018 The Missing Bridge Between Multi-Agent Systems and the World Wide Web. Proceedings of the 18th  
1019 International Conference on Autonomous Agents and MultiAgent Systems. International Foundation for  
1020 Autonomous Agents and Multiagent Systems, 2019, pp. 1659–1663.
- 1021 42. Vetrice, G.; Deaconescu, T. Actuating systems of elbow rehabilitation devices. *Annals of the Academy of*  
1022 *Romanian Scientists Series on Engineering Sciences* **2016**, *8*.
- 1023 43. Calvaresi, D.; Cesarini, D.; Sernani, P.; Marinoni, M.; Dragoni, A.F.; Sturm, A. Exploring the ambient  
1024 assisted living domain: a systematic review. *Journal of Ambient Intelligence and Humanized Computing* **2017**,  
1025 *8*, 239–257.



- 1026 44. Dell’Aglío, D.; Della Valle, E.; Calbimonte, J.P.; Corcho, O. RSP-QL semantics: a unifying query model  
1027 to explain heterogeneity of RDF stream processing systems. *International Journal on Semantic Web and*  
1028 *Information Systems (IJSWIS)* **2014**, *10*, 17–44.
- 1029 45. Keskiärrkkä, R.; Blomqvist, E.; Lind, L.; Hartig, O. RSP-QL\*: Enabling Statement-Level Annotations in  
1030 RDF Streams. *International Conference on Semantic Systems*. Springer, 2019, pp. 140–155.
- 1031 46. Buttazzo, G.C. *Hard real-time computing systems: predictable scheduling algorithms and applications*; Vol. 24,  
1032 Springer Science & Business Media, 2011.
- 1033 47. Greenwood, D.; Lyell, M.; Mallya, A.; Suguri, H. The IEEE FIPA approach to integrating software agents  
1034 and web services. *Proceedings of the 6th international joint conference on Autonomous agents and*  
1035 *multiagent systems*. ACM, 2007, p. 276.
- 1036 48. Stonebraker, M.; Çetintemel, U.; Zdonik, S.B. The 8 requirements of real-time stream processing. *SIGMOD*  
1037 *Record* **2005**, *34*, 42–47.
- 1038 49. Calvaresi, D.; Marinoni, M.; Lustrissimini, L.; Appoggetti, K.; Sernani, P.; Dragoni, A.F.; Schumacher,  
1039 M.; Buttazzo, G. Local scheduling in multi-agent systems: getting ready for safety-critical scenarios. In  
1040 *Multi-Agent Systems and Agreement Technologies*; Springer, 2017; pp. 96–111.
- 1041 50. Tommasini, R.; Sedira, Y.A.; Dell’Aglío, D.; Balduini, M.; Ali, M.I.; Le Phuoc, D.; Della Valle, E.; Calbimonte,  
1042 J.P. VoCaLS: Vocabulary and Catalog of Linked Streams. *International Semantic Web Conference*. Springer,  
1043 2018, pp. 256–272.
- 1044 51. Albanese, G.; Calvaresi, D.; Sernani, P.; Dubosson, F.; Dragoni, A.F.; Schumacher, M. MAXIM-GPRT:  
1045 A Simulator of Local Schedulers, Negotiations, and Communication for Multi-Agent Systems in  
1046 General-Purpose and Real-Time Scenarios. *International Conference on Practical Applications of Agents*  
1047 *and Multi-Agent Systems*. Springer, 2018, pp. 291–295.
- 1048 52. Pardo-Castellote, G.; Hamilton, M.; Thiebaut, S.S. Real-time publish-subscribe system, 2011. US Patent  
1049 7,882,253.
- 1050 53. Bellavista, P.; Corradi, A.; Foschini, L.; Pernafini, A. Data Distribution Service (DDS): A performance  
1051 comparison of OpenSplice and RTI implementations. *2013 IEEE symposium on computers and*  
1052 *communications (ISCC)*. IEEE, 2013, pp. 000377–000383.
- 1053 54. Calvaresi, D.; Appoggetti, K.; Lustrissimini, L.; Marinoni, M.; Sernani, P.; Dragoni, A.F.; Schumacher, M.  
1054 *Multi-Agent Systems’ Negotiation Protocols for Cyber-Physical Systems: Results from a Systematic*  
1055 *Literature Review*. *Proceedings of the 10th International Conference on Agents and Artificial*  
1056 *Intelligence, ICAART 2018, Volume 1, Funchal, Madeira, Portugal, January 16-18, 2018*, 2018, pp. 224–235.  
1057 doi:10.5220/0006594802240235.
- 1058 55. Smith, R.G. The contract net protocol: High-level communication and control in a distributed problem  
1059 solver. *IEEE Transactions on computers* **1980**, pp. 1104–1113.
- 1060 56. Zhang, Y.; Duc, P.M.; Corcho, O.; Calbimonte, J.P. SRBench: a streaming RDF/SPARQL benchmark. In *The*  
1061 *Semantic Web–ISWC 2012*; Springer, 2012; pp. 641–657.
- 1062 57. Patni, H.; Henson, C.; Sheth, A. Linked sensor data. *2010 International Symposium on Collaborative*  
1063 *Technologies and Systems*. IEEE, 2010, pp. 362–370.
- 1064 58. Della Valle, E.; Ceri, S.; Van Harmelen, F.; Fensel, D. It’s a streaming world! Reasoning upon rapidly  
1065 changing information. *IEEE Intelligent Syst.* **2009**, pp. 83–89.
- 1066 59. Xiao, F. EFMCDM: Evidential fuzzy multicriteria decision making based on belief entropy. *IEEE Transactions*  
1067 *on Fuzzy Systems* **2019**.
- 1068 60. Xiao, F.; Zhang, Z.; Abawajy, J. Workflow scheduling in distributed systems under fuzzy environment.  
1069 *Journal of Intelligent & Fuzzy Systems* **2019**, *37*, 5323–5333.
- 1070 61. Xiao, F. A new divergence measure for belief functions in D–S evidence theory for multisensor data fusion.  
1071 *Information Sciences* **2020**, *514*, 462–483.
- 1072 62. Knublauch, H.; Kontokostas, D. Shapes constraint language (SHACL). *W3C Recommendation* **2017**, *20*.