

Stream-Based Reasoning Support for Autonomous Systems¹

Fredrik Heintz and Jonas Kvarnström and Patrick Doherty²

Abstract. For autonomous systems such as unmanned aerial vehicles to successfully perform complex missions, a great deal of embedded reasoning is required at varying levels of abstraction. To support the integration and use of diverse reasoning modules we have developed DyKnow, a stream-based knowledge processing middleware framework. By using streams, DyKnow captures the incremental nature of sensor data and supports the continuous reasoning necessary to react to rapid changes in the environment.

DyKnow has a formal basis and pragmatically deals with many of the architectural issues which arise in autonomous systems. This includes a systematic stream-based method for handling the sense-reasoning gap, caused by the wide difference in abstraction levels between the noisy data generally available from sensors and the symbolic, semantically meaningful information required by many high-level reasoning modules. As concrete examples, stream-based support for anchoring and planning are presented.

1 Introduction

For autonomous systems such as unmanned aerial vehicles (UAVs) to successfully perform complex missions, a great deal of embedded reasoning is required. For this reasoning to be grounded in the environment, it must be firmly based on information gathered through available sensors. However, there is a wide gap in abstraction levels between the noisy numerical data directly generated by most sensors and the crisp symbolic information that many reasoning functionalities assume to be available. We call this the *sense-reasoning gap*.

Bridging this gap is a prerequisite for deliberative reasoning functionalities such as planning, execution monitoring, and diagnosis to be able to reason about the current development of dynamic and incompletely known environments using representations grounded through sensing. For example, when monitoring the execution of a plan, it is necessary to continually collect information from the environment to reason about whether the plan has the intended effects as specified in a symbolic high-level description.

Creating a suitable bridge is a challenging problem. It requires constructing representations of information incrementally extracted from the environment. This information must continuously be processed to generate information at increasing levels of abstraction while maintaining the necessary correlation between the generated information and the environment itself. The construction typically

requires a combination of a wide variety of methods, including standard functionalities such as signal and image processing, state estimation, and information fusion.

These and other forms of reasoning about information and knowledge have traditionally taken place in tightly coupled architectures on single computers. The current trend towards more heterogeneous, loosely coupled, and distributed systems necessitates new methods for connecting sensors, databases, components responsible for fusing and refining information, and components that reason about the system and the environment. This trend makes it less practical to statically predefine exactly how the information processing should be configured. Instead it is necessary to configure the way in which information and knowledge is processed and reasoned about in a context-dependent manner relative to high-level goals while globally optimizing the use of resources and the quality of the results.

To address these issues we have developed the stream-based knowledge processing middleware framework DyKnow [15, 16] which is a central part of our fully deployed unmanned aircraft system architecture [4, 5].

The rest of the paper is structured as follows. Section 2 discusses the requirements on middleware for embedded reasoning and Section 3 describes DyKnow. Sections 4 and 5 give two concrete examples of how DyKnow can facilitate embedded reasoning through support for anchoring and planning. Section 6 discusses DyKnow in relation to the requirements presented in Section 2. Section 7 presents related work. Finally, the conclusions are presented in Section 8.

2 Requirements

A wide range of functionality could conceivably be provided by middleware for embedded reasoning, and no single definition will be suitable for all systems. As a starting point, we present the requirements that have guided our work in distributed UAV architectures. These requirements are not binary in the sense that a system either satisfies them or not. Instead, a system will satisfy each requirement to some degree. Later, we will argue that DyKnow provides a significant degree of support for each of the requirements.

Support integration of existing reasoning functionality. The most fundamental property of middleware is that it supports interoperability. In the case of knowledge processing middleware, the main goal is to facilitate the integration of a wide variety of existing reasoning engines and sensors, bridging the gap between the distinct types of information required by and produced by such entities.

Support distributed sources and processing. Knowledge processing middleware should permit the integration of information from distributed sources and the distribution of processing across multiple computers. For UAVs, sources may include cameras, GPS sensors,

¹ This work is partially supported by grants from the Swedish Foundation for Strategic Research (SSF) Strategic Research Center MOVIII, the Swedish Research Council (VR) Linnaeus Center CADICS, ELLIIT Excellence Center at Linköping-Lund for Information Technology and the Center for Industrial Information Technology CENIT (projects 06.09 and 10.04).

² Dept. of Computer and Information Science, Linköping University, 581 83 Linköping, Sweden. Email: {frehe, jonkv, patdo}@ida.liu.se

and laser range scanners as well as higher level geographical information systems and declarative specifications of objects and their behaviors. Knowledge processing middleware should be sufficiently flexible to allow the integration of these into a coherent processing system while minimizing restrictions on connection topologies and the type of information being processed.

Support quantitative and qualitative processing on many levels of abstraction. In many applications there is a natural information abstraction hierarchy starting with quantitative signals from sensors, through representations of objects with both qualitative and quantitative attributes, to high level events and situations where objects have complex spatial and temporal relations. It should be possible to process information having arbitrary forms at arbitrary levels of abstraction, incrementally transforming it to forms suitable for various types of low-level and high-level reasoning.

Support bottom-up data processing and top-down model-based processing. While each process can be dependent on “lower level” processes for its input, it should also be possible for its output to guide processing in a top-down fashion. For example, if a vehicle is detected a vehicle model could be used to predict possible future locations, thereby directing or constraining processing on lower levels.

Support management of uncertainty. Uncertainty exists not only at the quantitative sensor data level but also in the symbolic identity of objects and in temporal and spatial aspects of events and situations. Therefore, middleware should not be constrained to the use of a single approach to handling uncertainty but should enable the combination and integration of different approaches.

Support flexible configuration and reconfiguration. When an agent’s resources are insufficient, either due to lack of processing power or due to sensory limitations, various forms of trade-offs are required. For example, update frequencies may be lowered, permitted processing delays may be increased, resource-hungry algorithms may be dynamically replaced with more efficient but less accurate ones, or the agent may focus its attention on the most important aspects of its current task. Reconfiguration may also be necessary when the current context or task changes.

Provide a declarative specification of the information being generated and the available information processing functionality. An agent should be able to reason about trade-offs and reconfiguration without outside help, which requires introspective capabilities. Specifically, it must be possible to determine what information is currently being generated as well as the potential effects of a reconfiguration. The declarative specification should provide sufficient detail to allow the agent to make rational trade-off decisions.

3 DyKnow

DyKnow is a fully implemented stream-based knowledge processing middleware framework providing both conceptual and practical support for structuring a knowledge processing system as a set of streams and computations on streams. Input can be provided by a wide range of distributed information sources on many levels of abstraction, while output consists of streams representing for example objects, attributes, relations, and events. As argued in Section 6, it satisfies each of the presented requirements to a significant degree.

Knowledge processing for a physical agent is fundamentally incremental in nature. Each part and functionality in the system, from sensing to deliberation, needs to receive relevant information about the environment with minimal delay and send processed information to interested parties as quickly as possible. Rather than using polling, explicit requests, or similar techniques, we have therefore chosen to




Figure 1. A prototypical knowledge process.

model and implement the required flow of data, information, and knowledge in terms of *streams*, while computations are modeled as active and sustained *knowledge processes* ranging in complexity from simple adaptation of raw sensor data to complex reactive and deliberative processes.

Streams lend themselves easily to a *publish/subscribe* architecture. Information generated by a knowledge process is published using one or more *stream generators*, each of which has a (possibly structured) *label* serving as an identifier within a knowledge processing application. Knowledge processes interested in a particular stream of information can subscribe to it using the label of the associated stream generator, which creates a new stream without the need for explicit knowledge of which process hosts the generator. Information produced by a process is immediately provided to the stream generator, which asynchronously delivers it to all subscribers, leaving the knowledge process free to continue its work. Using an asynchronous publish/subscribe pattern of communication decouples knowledge processes in time, space, and synchronization [9], providing a solid foundation for distributed knowledge processing applications.

Each stream is associated with a declarative *policy*, a set of requirements on its contents. Such requirements may include the fact that elements must arrive ordered by valid time, that each value must constitute a significant change relative to the previous value, that updates should be sent with a specific sample frequency, or that there is a maximum permitted delay. Policies can also give advice on how these requirements should be satisfied, for example by indicating how to handle missing or excessively delayed values.

A knowledge processing application in DyKnow consists of a set of knowledge processes connected by streams satisfying policies. An abstract view of a knowledge process is shown in Fig. 1. Each knowledge process is either an instantiation of a *source* or a *computational unit*. In the first case, it makes external information available through a stream generator, and in the second it refines and processes streams. A formal language called KPL is used to write declarative specifications of DyKnow applications (see [15, 19] for details). It provides a formal semantics for policies and streams. The DyKnow service, which implements the DyKnow framework, sets up the required processing and communication infrastructure for a given set of KPL declarations. Due to the use of CORBA [25] for communication, knowledge processes are location-unaware, supporting distributed architectures running on multiple networked computers.

Fig. 2 provides an overview of how part of the incremental processing required for a traffic surveillance task we have implemented is organized as a set of distinct DyKnow knowledge processes.

At the lowest level, a *helicopter state estimation component* uses data from an *inertial measurement unit* (IMU) and a *global positioning system* (GPS) to determine the current position and attitude of the UAV. A *camera state estimation component* uses this information, together with the current state of the *pan-tilt unit* on which the cameras are mounted, to generate information about the current camera state. The *image processing component* uses the camera state to determine where the camera is currently pointing. Video streams from the *color* and *thermal cameras* can then be analyzed in order to generate *vision*

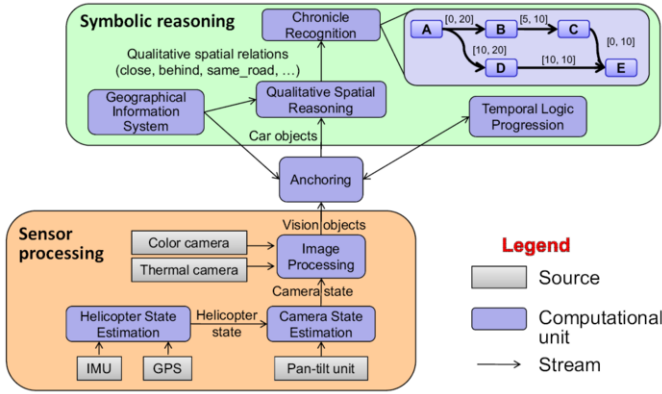


Figure 2. An overview of how the incremental processing required for a traffic surveillance task could be organized.

percepts representing hypotheses about moving and stationary physical entities, including their approximate positions and velocities.

Symbolic formalisms such as chronicle recognition [11] require a consistent assignment of symbols, or identities, to the physical objects being reasoned about and the sensor data received about those objects. Image analysis may provide a partial solution, with vision percepts having symbolic identities that persist over short intervals of time. However, changing visual conditions or objects temporarily being out of view lead to problems that image analysis cannot (and should not) handle. This is the task of the anchoring system to be described in the next section, which uses *progression* of formulas in a metric temporal logic to evaluate potential hypotheses about the observed objects. The anchoring system also assists in object classification and in the extraction of higher level attributes of an object. For example, a *geographic information system* can be used to determine whether an object is currently on a road or in a crossing. Such attributes can in turn be used to derive relations *between* objects, including *qualitative spatial relations* such as *beside*(car_1, car_2) and *close*(car_1, car_2). Concrete events corresponding to changes in such attributes and predicates finally provide sufficient information for the *chronicle recognition system* to determine when higher-level events such as reckless overtakes occur.

4 Support for Anchoring

Many reasoning systems assume perfect knowledge about the identity of objects. For example, a planner assumes that all objects in its planning domain are distinct and unique. An important problem, especially when bridging the sense-reasoning gap, is therefore to detect objects in streams of sensor data and to reason about their identities. The problem of how to create and maintain a consistent correlation between symbolic representations of objects and sensor data that is being continually collected about these objects is called the *anchoring problem* [3], a special case of the symbol grounding problem [14]. A concrete example is to detect and track cars during traffic monitoring using a UAV equipped with color and thermal cameras.

Tracking an object, such as a car, through a series of images is a classical problem. There are many effective solutions for the case where the object is easily distinguishable and can be tracked without interruptions. However, we must also consider the case where an object is temporarily hidden by obstacles (or tunnels in the case of traffic), and where many similar objects may be present in the world. Then, pure image-based tracking is not a complete solution, since it usually only considers the information available in the image itself. A

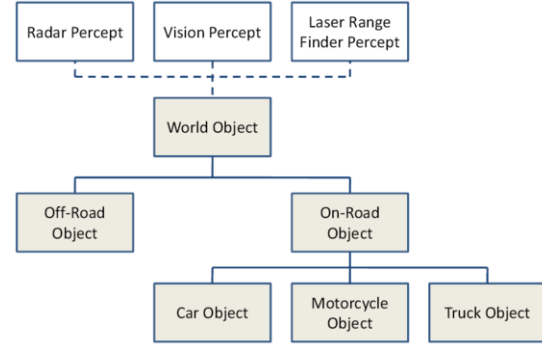


Figure 3. The example percept (white) / object (gray) hierarchy used in the traffic monitoring scenario.

more robust approach would need to actively reason about available knowledge of the world at higher abstraction levels, such as the normative characteristics of specific classes of physical objects. In the case of traffic, this would include the layout of the road network and the typical size, speed, and driving behavior of cars. It has been argued that anchoring can be seen as an extension to classical tracking approaches which handles missing data in a principled manner [10].

Existing approaches to anchoring work under the limiting assumption that each individual piece of sensor data, such as a blob found in a single frame from a video camera, should be anchored to a symbol in a single step. We believe that much can be gained in terms of accuracy as well as speed of recognition by taking advantage of the fact that one generally has access to a timed sequence, or stream, of sensor data related to a particular object.

We have therefore extended DyKnow with a stream-based hierarchical anchoring framework for incrementally anchoring symbols to streams of sensor data [18]. The anchoring process constructs and maintains a set of *object linkage structures* representing the best possible hypotheses at any time. Each hypothesis is continually monitored and refined as more and more information becomes available. Symbols can be associated with an object at any level of classification, permitting symbolic reasoning on different levels of abstraction.

An example hierarchy for traffic monitoring can be seen in Fig. 3. A *world object* represents a physical object detected in the world. Its attributes are based on information from one or more linked percepts and include the absolute coordinates of the object in the physical world. World objects could either be *on-road objects* moving along roads or *off-road objects* not following roads. An on-road object has attributes representing the road segment or crossing the object occupies, enabling more qualitative forms of reasoning, and an improved position estimation which is snapped to the road. Finally, an on-road object could be a *car*, a *motorcycle*, or a *truck*. Each level in the hierarchy adds more abstract and qualitative information while still maintaining a copy of the attributes of the object it was derived from. Thus, an on-road object contains both the original position from the world object and the position projected onto the road network.

Hypotheses about object types and identities must be able to evolve over time. For example, while it might be determined quickly that a world object is an on-road object, more time may be required to determine that it is in fact a car. Also, what initially appeared to be a car might later turn out to be better classified as a truck. To support incremental addition of information and incremental revision of hypotheses, a single physical object is not represented as an indivisible object structure but as an *object linkage structure*.

An object linkage structure consists of a set of *objects* which are *linked* together (note that percepts are also considered to be objects).




Figure 4. An example object linkage structure.

Each object has a type and is associated with a symbol, and represents information about a particular physical object at a given level of abstraction. A symbol is *anchored* if its object is part of an object linkage structure that is grounded in at least one percept. For example, Fig. 4 represents the hypothesis that vision percept vp8, world object wo5, on-road object oo3, and car object co2 all correspond to the same physical object.

Whenever a new object of a given type is generated, it must be determined whether it also belongs to a particular subtype in the hierarchy. For example, a new vision percept originating in image processing may be classified as corresponding to a world object. In this case, it must also be linked to a world object structure, thereby generating an object linkage structure. However, it is essential that sensor data can be anchored even to symbols / objects for which no percepts have arrived for a period of time. Thus, objects and their symbols are not immediately removed when their associated percepts disappear, and any new object at one level might correspond to either a new object or an *existing* object at the next level. To reduce the computational cost, objects that are not likely to be found again are removed. Currently we discard objects which have not been observed or anchored for a certain application-dependent time.

Three conditions are used to determine when to add and remove links between objects belonging to types A and B. These conditions are written in an expressive temporal logic, similar to the well known Metric Temporal Logic [21], and incrementally evaluated by DyKnow using progression over a timed state sequence.³

The unary *establish condition* expresses when an object of type A, which may be a percept or a higher level object, should be linked to a *new* object of type B. When a new object of type A is created, the anchoring system immediately begins evaluating this condition. This is similar to the Find functionality suggested by Coradeschi and Saffiotti [3], which takes a symbolic description of an object and tries to anchor it in sensor data. A world object could for example be linked to a new on-road object if it is on a road for at least five seconds.

The binary *reestablish condition* expresses the condition for an object of type A to be linked to a *known* object of type B, as in the case where a new world object corresponds to an existing on-road object that had temporarily been hidden by a bridge. When a new object of type A is created, the anchoring system immediately begins to evaluate the re-establish condition for every known object of type B that is not linked to an object of type A. If and when one of these conditions becomes true, a link is created between the associated objects. This is similar to the Reacquire functionality [3]. A link from a world object to an existing on-road object could for example be reestablished if the distance is less than thresh within five seconds.

While two objects are linked, a computational unit generates the attributes of the more specific object from those of the less specific object, possibly together with information from other sources.

Finally, since observations are uncertain and classification is imperfect, any link created between two objects is considered a hypothesis and is continually validated through a *maintain condition*. Such conditions can compare the observed behavior of an object with be-

havior that is normative for its type, and possibly with behavior predicted in other ways. For example, one might state that an on-road object should remain continually on the road, maybe with occasional shorter periods being observed off the road due to sensor error.

If a maintain condition is violated, the corresponding link is removed. However, all objects involved remain, enabling reclassification and re-identification at a later time. The state of an object having no incoming links will be predicted based on a general model of how objects of this type normally behave. This is similar to the Track functionality [3].

5 Support for Planning

One approach to solving complex problems is to use a task planner. To integrate task planners into an embedded reasoning system there are a number of issues to consider.

Initial state. For a planner to be able to generate a plan which is relevant in the current situation it must have an accurate and up-to-date domain model. In a static environment it is possible to write a domain model once and for all since the world does not change. In a dynamic environment, such as a disaster area, we do not have the luxury of predefined static domain models. Instead, the UAV must itself generate information about the current state of the environment and encode this in a domain model.

Execution. Executing an action in a plan generally requires sophisticated feedback about the environment on different levels of abstraction. For example, a UAV following a three-dimensional trajectory must continually estimate its position by fusing data from several sensors, such as GPS and IMU. If it loses its GPS signal due to malfunction or jamming, vision-based landing may be needed, which requires processing video streams from cameras in order to estimate altitude and position relative to the landing site.

Monitoring. Classical task planners are built on the fundamental assumption that the only agent causing changes in the environment is the planner itself, or rather, the system or systems that will eventually execute the plan that it generates. Furthermore, they assume that all information provided to the planner as part of the initial state and the operator specifications is accurate. This may in some cases be a reasonable approximation of reality, but it is not always the case. Other agents might manipulate the environment of a system in ways that may prevent the successful execution of a plan. Sometimes actions can fail to have the effects that were modeled in a planning domain specification, regardless of the effort spent modeling all possible contingencies. Consequently, robust performance in a noisy environment requires some form of supervision, where the execution of a plan is constantly monitored in order to detect any discrepancies and recover from potential or actual failures.

We have developed a task planning and execution monitoring framework that handles these issues and integrated it into our unmanned aircraft system [8]. In the spirit of cognitive robotics, this framework makes extensive use of Temporal Action Logic (TAL [7]), a logic for reasoning about action and change. For example, knowledge gathered during plan execution can be used by DyKnow to incrementally create state structures that correspond to partial logical models in TAL, representing the actual development of the system and its environment over time. Desired and expected behaviors can then be described using TAL formulas.

Example 1 Suppose that a UAV supports a maximum continuous power usage of M , but can exceed this by a factor of f for up to τ

³ Progression incrementally evaluates formulas in a state sequence. The result of progressing a formula through the first state in a sequence is a new formula that holds in the remainder of the state sequence iff the original formula holds in the complete state sequence.

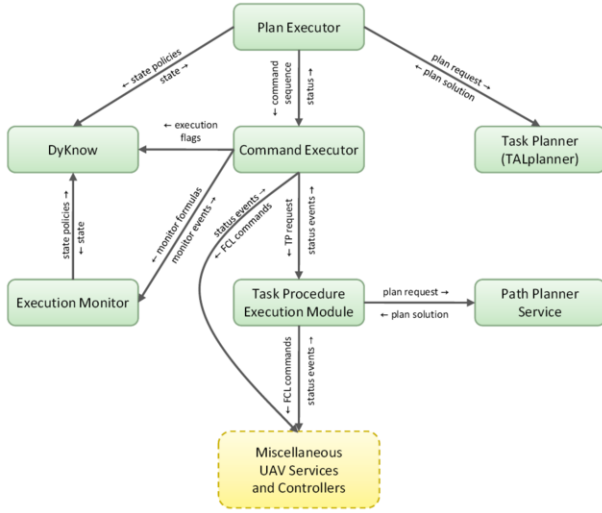


Figure 5. Task planning and execution monitoring overview

units of time, if this is followed by normal power usage for a period of length at least τ' . The following formula can be used to detect violations of this specification:

$$\square \forall uav. (power(uav) > M \rightarrow power(uav) < f \cdot M \text{ } U_{[0, \tau]} \square_{[0, \tau']} power(uav) \leq M) \quad \square$$

The pervasive use of this logic throughout the higher level deliberative layers of the system architecture also provides a solid shared declarative semantics that facilitates knowledge transfer between different modules. Specifically, TAL also forms the semantic basis for our task planner TALplanner [6, 23], which is used to generate mission plans. Basing the execution monitoring framework on the same logic enables us to automatically extract certain necessary conditions that should be monitored during execution.

Fig. 5 shows the relevant part of the UAV system architecture associated with task planning, plan execution, and execution monitoring.

At the top of the center column is the *plan executor* which given a mission request uses DyKnow to acquire essential information about the current contextual state of the world and the UAV's own internal state. Together with a domain specification and a goal specification related to the current mission, this information is fed to TALplanner, which outputs a plan that will achieve the designated goals, under the assumption that all actions succeed and no failures occur. Such a plan can also be automatically annotated with global and/or operator-specific TAL formulas representing conditions to be monitored during execution in order to relax the assumption of success.

The plan executor translates operators in the high-level plan returned by TALplanner into lower level command sequences which are given to the *command executor*. The command executor is responsible for controlling the UAV, either by directly calling the functionality exposed by its lowest level Flight Command Language (FCL) interface or by using Task Procedures, a type of reactive procedures, through the *Task Procedure Execution Module*.

During execution, the command executor adds formulas to be monitored to the *execution monitor*. The monitor configures DyKnow to continuously provide it with current states and uses a progression algorithm [8] to partially evaluate monitor formulas and promptly detect violations. States are generated from potentially asynchronous streams by a synchronization mechanism that uses the formal semantics of the input streams as specified by declarative policies to determine when states should be created.

If a violation is detected, this is immediately signaled by the exe-

cution monitor as an event to the command executor, which can suspend the execution of the current plan, invoke an emergency brake command if required, optionally execute an initial recovery action, and finally signal the new status to the plan executor. The plan executor is then responsible for completing the recovery procedure.

The fully integrated system is implemented on our UAVs. It shows that this approach allows integrating high-level deliberative capability with lower-level reactive and control functionality.

6 Discussion

In the beginning of the paper we introduced a number of requirements for middleware for embedded reasoning. In this section we argue that DyKnow provides a significant degree of support for each of those requirements.

Support integration of existing reasoning functionality. Streams provide a powerful yet very general representation of information varying over time, and any reasoning functionality whose inputs can be modeled as streams can easily be integrated using DyKnow. As two concrete examples, we have shown how progression of temporal logical formulas [8] and chronicle recognition [20] can be integrated using DyKnow.

Support distributed sources and processing. DyKnow satisfies this requirement through the use of the general concepts of streams and knowledge processes. Since the implementation is CORBA-based it provides good support for distributed applications. DyKnow explicitly represents both the time when information is valid and when it is available. Therefore it has excellent support for integrating information over time even with varying delays. DyKnow also provides a very useful stream synchronization mechanism that uses the formal semantics of the declarative policies to determine how to synchronize a set of asynchronous streams and derive a stream of states. This functionality is for example used to create synchronized state sequences over which temporal logical formulas can be evaluated.

Support processing on many levels of abstraction. General support is provided in DyKnow through streams, where information can be sent at any abstraction level from raw sampled sensor data and upwards. Knowledge processes also provide general support for arbitrary forms of processing. At the same time, DyKnow is explicitly designed to be extensible to provide support for information structures and knowledge processing that is more specific than arbitrary streams. DyKnow directly supports specific forms of high-level information structures, such as object linkage structures, and specific forms of knowledge processing, including formula progression and chronicle recognition. This provides initial support for knowledge processing at higher levels than plain streams of data. In [17] we argue that this provides an appropriate framework for supporting all the functional abstraction levels in the JDL Data Fusion Model [27].

Support quantitative and qualitative processing. Streams provide support for arbitrarily complex data structures, from real values to images to object structures to qualitative relations. The structured content of samples also allows quantitative and qualitative information to be part of the same sample. DyKnow also has explicit support for combining qualitative and quantitative processing in the form of chronicle recognition, progression of metric temporal logical formulas, and object linkage structures. Both chronicles and temporal logical formulas support expressing conditions combining quantitative time and qualitative features.

Support bottom-up data processing and top-down model-based processing. Streams are directed but can be connected freely, giving the application programmer the possibility to do both top-down and

bottom-up processing. Though this article has mostly used bottom-up processing, chronicle recognition is a typical example of top-down model-based processing where the recognition engine may control the data being produced depending on the general event pattern it is attempting to detect.

Support management of uncertainty. In principle, DyKnow supports any approach to representing and managing uncertainty that can be handled by processes connected by streams. It is for example easy to add a probability or certainty factor to each sample in a stream. This information can then be used by knowledge processes subscribing to this stream. Additionally, DyKnow has explicit support for uncertainty in object identities and in the temporal uncertainty of complex events that can be expressed both in quantitative and qualitative terms. The use of a metric temporal logic also provides several ways to express temporal uncertainty.

Support flexible configuration and reconfiguration. Flexible configuration is provided by the declarative specification language KPL, which allows an application designer to describe the processes in a knowledge processing application and how they are connected with streams satisfying specific policies. The implementation uses the specification to instantiate and connect the required processes.

Provide a declarative specification of the information being generated and the available information processing functionality. This requirement is satisfied through the formal language KPL for declarative specifications of DyKnow knowledge processing applications. The specification explicitly declares the properties of the streams by policies and how they connect the different knowledge processes.

7 Related Work

Stream processing has a long history [26] and today data stream management systems such as XStream [12], Aurora [2], and Borealis [2] provide continuous query languages supporting filters, maps, and joins on streams. Similarly, systems for complex event processing, such as Cayuga [1] and SASE [13], provide specification languages for defining complex events in terms of more primitive events occurring in event streams. However, these systems do not provide the appropriate abstractions for stream-based reasoning by themselves.

There are also many development frameworks for robotic systems [22, 24] which focus on integration aspects, but do not explicitly provide support for stream-based reasoning. What these frameworks lack are formal ways to specify streams. For example, in DyKnow it is possible to formally specify the start and end time of a stream, the sampling period, and how to approximate missing values. DyKnow streams are also active, producing samples according to the semantics of a policy while other frameworks use streams only to transport given input to multiple subscribers. Another important aspect that is often not supported is an explicit representation of time. In DyKnow each sample is tagged with a valid and an available time which makes it possible to reason about when a value is true, when it was actually available to the process, and how much it was delayed by previous processing. Using policies and these time-stamps synchronized states can be extracted from asynchronous streams with a precise meaning.

8 Conclusion

As autonomous systems become more sophisticated and are expected to handle increasingly complex and challenging tasks and missions, there is a growing need to integrate a variety of reasoning functionalities. In this paper we have given a high-level overview of the stream-based middleware DyKnow and how it can support embedded rea-

soning. We discussed general requirements on this type of middleware and showed how DyKnow can support stream-based reasoning such as anchoring and planning in an autonomous system. We have thereby taken the first steps towards developing powerful, robust, and general support for stream-based reasoning, which we believe will continue to increase in importance since it captures the incremental nature of the information available and the continuous reasoning necessary to react to rapid changes in the environment.

REFERENCES

- [1] L. Brenna, A. Demers, J. Gehrke, M. Hong, J. Ossher, B. Panda, M. Riedewald, M. Thatte, and W. White, 'Cayuga: a high-performance event processing engine', in *Proc. ICMD*, (2007).
- [2] U. Çetintemel, D. Abadi, Y. Ahmad, H. Balakrishnan, M. Balazinska, M. Cherniack, J. Hwang, W. Lindner, S. Madden, A. Maskey, A. Rasin, E. Ryvkina, M. Stonebraker, N. Tatbul, Y. Xing, and S. Zdonik, 'The Aurora and Borealis Stream Processing Engines', in *Data Stream Management: Processing High-Speed Data Streams*, (2007).
- [3] S. Coradeschi and A. Saffiotti, 'An introduction to the anchoring problem', *Robotics and Autonomous Systems*, **43**(2–3), 85–96, (2003).
- [4] P. Doherty, 'Advanced research with autonomous unmanned aerial vehicles', in *Proc. KR*, (2004).
- [5] P. Doherty, P. Haslum, F. Heintz, T. Merz, P. Nyblom, T. Persson, and B. Wingman, 'A distributed architecture for autonomous unmanned aerial vehicle experimentation', in *Proc. DARS*, (2004).
- [6] P. Doherty and J. Kvarnström, 'TALplanner: A temporal logic-based planner', *AI Magazine*, **22**(3), 95–102, (2001).
- [7] P. Doherty and J. Kvarnström, 'Temporal action logics', in *Handbook of Knowledge Representation*, Elsevier, (2008).
- [8] P. Doherty, J. Kvarnström, and F. Heintz, 'A temporal logic-based planning and execution monitoring framework for unmanned aircraft systems', *J. of Auton. Agents and Multi-Agent Systems*, **19**(3), (2009).
- [9] P. Eugster, P. Felber, R. Guerraoui, and A.-M. Kermarrec, 'The many faces of publish/subscribe', *ACM Comput. Surv.*, **35**(2), (2003).
- [10] J. Fritsch, M. Kleinhagenbrock, S. Lang, T. Plötz, G. A. Fink, and G. Sagerer, 'Multi-modal anchoring for human-robot interaction', *Robotics and Autonomous Systems*, **43**(2–3), 133–147, (2003).
- [11] M. Ghallab, 'On chronicles: Representation, on-line recognition and learning', in *Proc. KR*, (1996).
- [12] L. Girod, Y. Mei, R. Newton, S. Rost, A. Thiagarajan, H. Balakrishnan, and S. Madden, 'Xstream: a signal-oriented data stream management system', in *Proc. ICDE*, (2008).
- [13] D. Gyllstrom, E. Wu, H. Chae, Y. Diao, P. Stahlberg, and G. Anderson, 'Sase: Complex event processing over streams', in *Proc. CIDR*, (2007).
- [14] S. Harnad, 'The symbol-grounding problem', *Physica*, **D**(42), (1990).
- [15] F. Heintz, *DyKnow: A Stream-Based Knowledge Processing Middleware Framework*, Ph.D. dissertation, Linköpings universitet, 2009.
- [16] F. Heintz and P. Doherty, 'DyKnow: An approach to middleware for knowledge processing', *J. of Intelligent and Fuzzy Syst.*, **15**(1), (2004).
- [17] F. Heintz and P. Doherty, 'A knowledge processing middleware framework and its relation to the JDL data fusion model', *J. of Intelligent and Fuzzy Syst.*, **17**(4), (2006).
- [18] F. Heintz, J. Kvarnström, and P. Doherty, 'A stream-based hierarchical anchoring framework', in *Proc. IROS*, (2009).
- [19] F. Heintz, J. Kvarnström, and P. Doherty, 'Bridging the sense-reasoning gap: DyKnow – stream-based middleware for knowledge processing', *Journal of Advanced Engineering Informatics*, **24**(1), 14–26, (2010).
- [20] F. Heintz, P. Rudol, and P. Doherty, 'From images to traffic behavior – a UAV tracking and monitoring application', in *Proc. Fusion*, (2007).
- [21] R. Koymans, 'Specifying real-time properties with metric temporal logic', *Real-Time Systems*, **2**(4), 255–299, (1990).
- [22] J. Kramer and M. Scheutz, 'Development environments for autonomous mobile robots: A survey', *Autonomous Robots*, **22**(2), (2007).
- [23] J. Kvarnström, *TALplanner and Other Extensions to Temporal Action Logic*, Ph.D. dissertation, Linköpings universitet, 2005.
- [24] N. Mohamed, J. Al-Jaroodi, and I. Jawhar, 'Middleware for robotics: A survey', in *Proc. RAM*, (2008).
- [25] Object Management Group. The CORBA specification v 3.1, jan 2008.
- [26] R. Stephens, 'A survey of stream processing', *Acta Informatica*, **34**(7), (1997).
- [27] F. White, 'A model for data fusion', in *Proc. Sensor Fusion*, (1988).