Perspective Filters as a Means for Interoperability Among Information-Centric Decision-Support Systems

By
Kym J. Pohl
Collaborative Agent Design Research Center (CADRC)
California Polytechnic University
San Luis Obispo, California
kpohl@cadrc.calpoly.edu

If the key to symbolic reasoning is representation then it certainly follows that the foundation of expert-system-based, decision-support systems is the rich manner in which the entities, concepts, and notions relevant to the domain space(s) are represented [4, 10]. This requirement can be accommodated through the development and employment of one or more ontologies. An ontology in this sense can be defined as a relationship-rich, typically object-based representation of the entities, concepts, and notions relevant to the domain(s) of operation. The problem arises when two or more of these systems, each operating over a potentially extensive ontology attempt to collaborate with each other. While collaboration within each of these systems may be based on very high-level descriptions of entities, concepts and notions, it will undoubtedly be subject to various application-specific biases.

For example, in a tactical command and control system an entity such as an M1A1 tank may be viewed, and therefore represented as a tactical asset. In this case the bias would be toward tactical utility. However, in a logistics system the same M1A1 tank would most appropriately be viewed as a potential supply item with emphasis on logistical inventory and supply. In both cases, however, the subject is still the exact same M1A1 tank with basic characteristics. The difference resides in the manner in which the tank is being viewed by each of these systems. Another term for this bias-based filter is *perspective*. Perspective is not only a natural component of the way in which we perceive the world but moreover should be viewed as a highly beneficial and desirable characteristic. Perspective is the ingredient in an ontology-based decision-support system that allows for the representation of domain-specific notions and bias. For example, if a decision-support system is to assist in the formulation of logistical supply missions then it is more appropriate, and beneficial for an entity such as a howitzer to be primarily viewed as a supply item instead of a tactical asset. If viewed as a supply item the description of a howitzer could provide great detail in terms of the items shipping weight, shipping dimensions, tie-down points, etc. In the context of a tactical command and control system such information is essentially irrelevant and certainly not of primary focus. What would be relevant in such a tactical system would be characteristics such as projectile range, effective casualty radius, advancement velocity, etc. Again, it may be the exact same howitzer that is being discussed between the two disparate systems. However, it is being discussed within two different contexts exhibiting two distinctly different perspectives. While collaboration within or across systems supported by the exact same perspective-based representation performs well, the problem arises when collaboration needs to occur between systems or system components where the perspectives are in fact not the same and potentially drastically dissimilar. In this unto common case, the extent to which systems can collaborate on events and information is essentially limited to low-level data-passing with receivers having little or no understanding of content and implication. Simply stated, the problem at the heart of interoperability between symbolic reasoning-based systems resides in the means by which information-centric systems exhibiting wholly, or even partially disparate perspectives, can interoperate at a meaningful and useful level.

The solution to this dilemma can take primarily two different directions. The first of these paths focuses on the development of a *universal* ontology. Such an ontology would represent a single, all inclusive view of the world. Each system would utilize this representation as the core informational basis for operation. Since each system would have knowledge of this common representation of the entities, notions, and concepts, interoperability at the information level would be clear and concise requiring no context-diminishing translation. However, as straightforward as this may appear there are two major flaws with this approach. First, in practicality it is highly unlikely that such a universal description could actually be successfully developed. Considering the amount of forethought and vision this task would require, such an undertaking would be of monumental scale as well as being plagued with misrepresentation. Inevitably, certain notions or concepts would be inappropriately represented in a particular domain in an effort to model them adequately in another.

The second flaw with the universal ontology approach is less obvious but perhaps even more destructive. Considering the number of domains across which such an ontology would need to encompass the resulting ontology would most likely be comprised mainly of generalities. These generalities would typically only partially represent the manner in which any one particular system wished to *see the world*. In other words, due to the number of perspectives a universal ontology would attempt to represent, the resulting ontology would ironically end up being just the opposite, a perspective-absent description falling far short of system needs and expectations. While perspective was the cause of the original interoperability problem it is still a highly valuable characteristic that should not only be preserved but should be wholeheartedly embraced and promoted. As mentioned earlier, perspective is a valuable and useful means of conveying domain-specific notions and bias, which are crucial to information-centric decision-support systems. To omit its presence is to significantly reduce the usefulness of an ontology and therefore the effectiveness of the utilizing decision-support system(s). This coupled with the highly unlikely potential for developing such a comprehensive, inter-domain description of the world renders the universal ontology approach both unrealistic and wholly ineffective.

The second, more promising solution to interoperability between decision-support systems introduces the notion of a *perspective filter*. Based on the façade design pattern [1, 2, 3] perspective filters allow core entities, concepts and notion accessible to interoperating systems to be viewed in a more appropriate form relative to each collaborator's perspective. In brief, the façade pattern allows for a certain description to be viewed, and consequently interacted with in a more appropriate manner. Similar to a pair of infrared *night vision* goggles, overlaying a filter may enhance or refine otherwise limited information. In the case of ontology-based collaboration this filter essentially superimposes a more perspective-oriented, ontological layer over the initial representation. The filter may not only add or modify the terminology and constraints of the core descriptions but may also extend and enhance it through the incorporation of additional characteristics.

These characteristics may take the form of additional attributes and relationships as well as refining constraints. For example, Figure 1 illustrates the use of a logistically oriented perspective filter over a core description of conveyances. Note first that while the core conveyance ontology appears to represent only a limited amount of bias the effectiveness of perspective filters certainly does not require such a general core description. If the core ontology were heavily biased toward a foreign set of perspectives it would simply mean that the perspective filters would need to be more extensive and incorporate additional constraints, extensions, etc. However, for clarity of illustration a limited, rather general core ontology was selected.

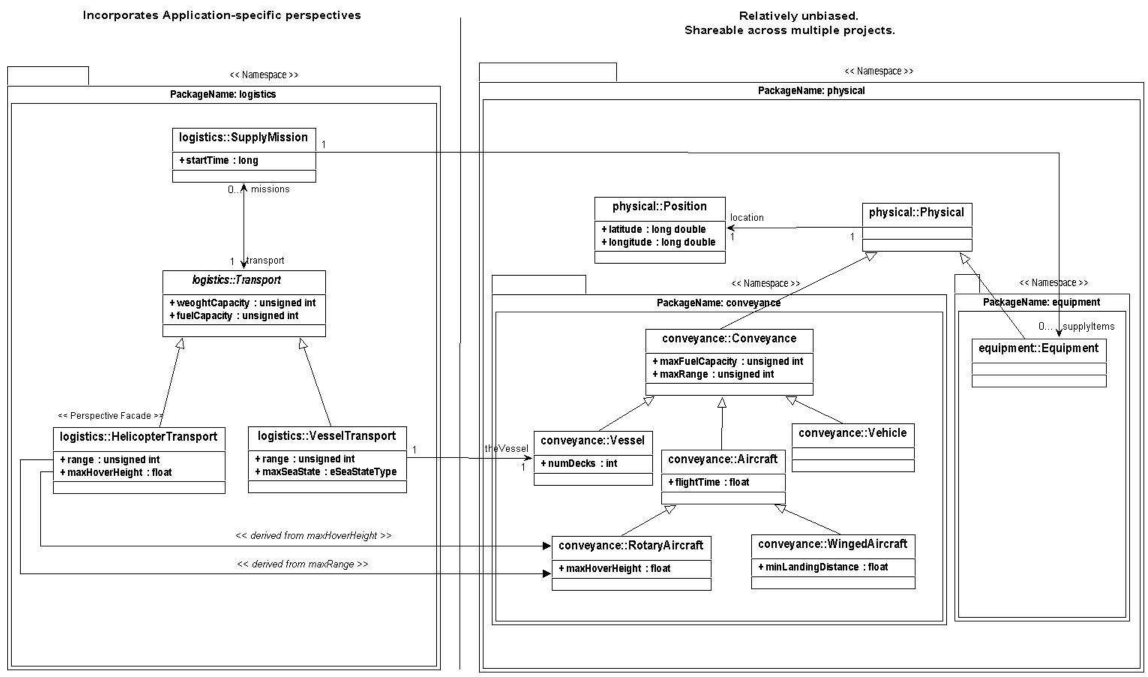

Figure 1 – Partially Derived Logistics Ontology

Core to the logistics perspective presented in Figure 1 is the notion of a transport. However, although the logistics system may have a notion of all of the types of conveyances (i.e., vessels, vehicles, and aircraft) represented in the core ontology it, in the context of this example, may only consider vessels and rotary aircraft as potential transports. In this situation it would be valuable to represent this refined constraint in the ontology forming the representational heart of the logistics system while still employing the core conveyance ontology. As Figure 1 illustrates, representing such refinement can be accomplished by explicitly introducing a constrained notion of a transport in the application-specific filter ontology. An abstract *Transport* is defined to have two specific derivations (*VesselTransport* and *HelicopterTransport*). At this point it is immediately apparent that a vehicle is not a transport candidate. In the context of the example logistics system transports can only be *VesselTransports* or *HelicopterTransports*. The task now becomes linking these two system specific notions to the core conveyance ontology. Relating these two transport types to their conveyance ontology counterparts can be achieved in two different ways. For illustration purposes, the definition of *VesselTransport* adopts the first method while *HelicopterTransport* employs the second. The first method defines an explicit relationship between the *VesselTransport* and the core description of a vessel outlined in the conveyance ontology. Utilizing this approach, obtaining the core information relative to the corresponding *Vessel* from a *VesselTransport* requires both knowledge of their relationship in addition to another level of indirection. For reasons of performance and logical integrity, both of these requirements may not be desirable.

The second method, illustrated in Figure 1 using *HelicopterTransport*, avoids both shortcomings inherent in the first approach. In this case, *HelicopterTransport* exists as a façade, or filter, which transparently links at the attribute level into the core *RotaryAircraft* description. That is, each attribute of *RotaryAircraft* desired to be exposed to users of *HelicopterTransport* is explicitly declared in the façade. For example, since the maximum range of travel is relevant to the definition of a *HelicopterTransport* the *maxRange* attribute of *RotaryAircraft* (inherited from *Conveyance*) is subsequently exposed in the HelicopterTransport façade description. By virtue of being declared in a façade any access to such an attribute would be transparently mapped into the corresponding attribute(s) on which it is based. In the case of the *range* attribute of *HelicopterTransport*, access would transparently be directed to the inherited *maxRange* attribute of *RotaryAircraft*. Notice also the use of alternative terminology over that used in the core ontology (i.e., *range* vs. *maxRange*). It should also be noted that the derivative nature of a façade attribute is not limited to mapping into another attribute. Rather, the value of a façade attribute may also be derived through calculation, perhaps based on the values of multiple attributes residing in potentially several different core objects. In either case, the fact that the value of the façade attribute is derived is completely transparent to the façade user.

Another perspective-oriented enhancement to the core ontology illustrated in Figure 1 is the notion of a *SupplyMission*. Being a fundamental concept in the example logistics system a supply mission essentially relates supply items in the form of equipment to the transports by which they will be delivered. Once again, the definition of a logistics-specific notion (i.e., supply items) is derived from a notion defined in the core ontology (i.e., equipment). In this case, an explicit relationship is declared linking *SupplyMission* to zero or more *Equipment* items. Since, from the perspective of the logistics system *Equipment* scheduled for delivery are viewed as items that are to be supplied, the term *supplyItems* is used as the referencing nomenclature. Such an enhancement demonstrates the ability to integrate new concepts (i.e., supply missions) with existing core notions.

In the context of interoperability among information-centric, decision-support systems significant benefits could be obtained from essentially drawing relevant concepts and notions into a system's local set of perspective-rich, filter ontologies. As the above example illustrates, key components of these perspective-oriented ontologies could be derived from a set of core, relatively unbiased common notions forming the basis for informational collaboration among systems. There are several benefits to adopting this approach. Collaboration among information-centric, decision-support systems would take place in terms of various core ontologies (i.e., *Conveyance*) with each collaborator viewing these core entities, concepts and notions according to its own perspective. Figure 2 briefly extends the logistics example presented in Figure 1 showing collaboration between the original logistics system and a tactical command and control system. Collaboration between these two example systems is in terms of the common, core ontologies on which they share their derivations. A conveyance is still a conveyance whether it is viewed in the context of logistics or tactical command and control. To represent domain-specific notions (e.g., transport, supply item, tactical asset, etc.) each collaborating system would apply the appropriate filter. Although discussing a conveyance from partially disparate perspectives both systems can collaborate about core entities, concepts, and notions.

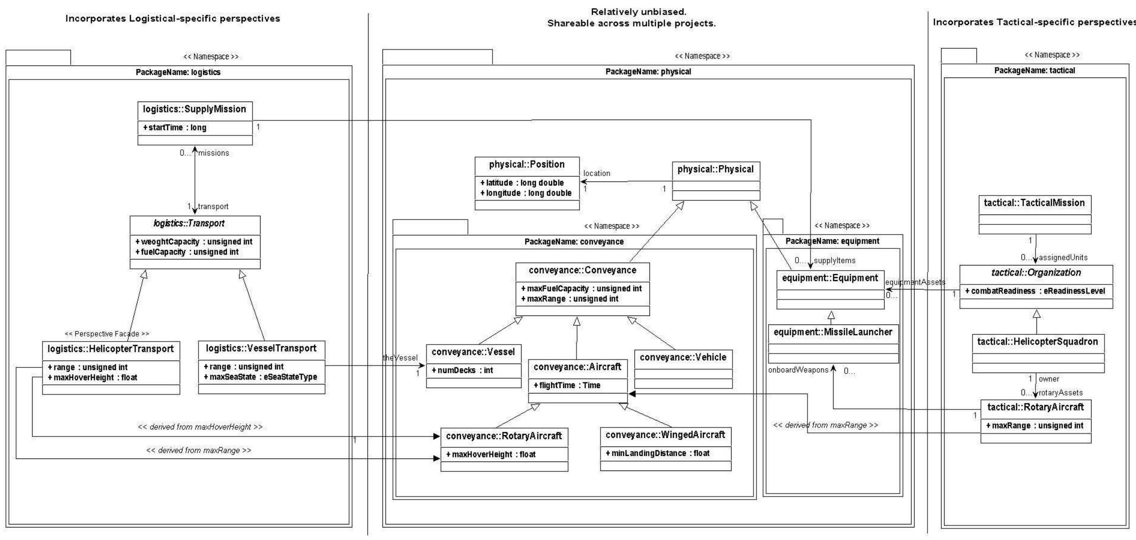

Figure 2 – Two disparate domains linked into the same core ontology

Another advantage of supplementing core, non-system-specific ontologies with perspective rich filters is the preservation of both time and effort during the development of such information-centric systems. Core ontologies could be archived in a sort of ontology library forming a useful reference assisting in the development of new system ontologies. Models created for new decision-support systems could make use of this ontology library as a strong basis for deriving system-specific filters. In addition, such a process would promote the use of common core descriptions increasing the potential for interoperability even further.

Interoperability between disparate decision-support systems is crucial to the operational effectiveness of information-centric, decision-support systems. As the emergence of such systems increases the need to support inter-system collaboration at the information level becomes increasingly critical. By constraining valuable, perspective-based biases to local, system-specific filter ontologies coupled with the use of core, relatively unbiased ontologies, interoperability between disparate information-centric decision-support systems becomes both feasible and effective.

References

[1] Buschmann F, D Schmidt, H Rohnert, and M Stal; "*Pattern-Oriented Software Architecture: A System of Patterns*"; Vol. 2, John Wiley and Sons, New York, New York, 1996.

[2] Buschmann F, R Meunier, H Rohnert, P Sommerlad, and M Stal; "*Pattern-Oriented Software Architecture: A System of Patterns*"; Vol. 1, John Wiley and Sons, New York, New York, 1996.

[3] Fowler M; "*Analysis Patterns: Reusable Object Models*"; Addison-Wesley, Reading, Massachusetts, 1997.

[4] Gomez-Perez, A., N. Juristo and J. Pazos (1995); *'Evaluation and Assessment of Knowledge Sharing Technology'*; in Mars N. (ed.) Towards Very Large Knowledge Bases - Knowledge Building and Knowledge Sharing 1995, IOS Press, Amsterdam (pp.289-296).

[5] Hay D.; "*Data Model Patterns: Conventions of Thought*"; Dorset House, New York, New York, 1996.

[6] Hayes-Roth F, D Waterman and D Lenat (eds.); "*Building Expert Systems*"; Addison-Wesley, Reading, Massachusetts, 1983.

[7] Orfali R, D Harkey and J Edwards; "*The Essential Distributed Objects Survival Guide*"; John Wiley and Sons, Inc., New York, New York, 1996.

[8] Pohl J, A Chapman, K Pohl, J Primrose and A Wozniak; "*Decision-Support Systems: Notions, Prototypes, and In-Use Applications*"; Technical Report, CADRU-11-97, CAD Research Center, Design Institute, College of Architecture and Environmental Design, Cal Poly, San Luis Obispo, CA, January, 1997.

[9] Pohl J; "*Human-Computer Partnership in Decision-Support Systems: Some Design Guidelines*"; in Pohl J (ed.) Advances in Collaborative Design and Decision-Support Systems, focus symposium: International Conference on Systems Research, Informatics and Cybernetics, Baden-Baden, Germany, August 18-22, 1997 (pp. 71-82).

[10] Tansley D. and C. Hayball (1993); *'Knowledge-Based Systems Analysis and Design: A KAD Developers Handbook'*, Prentice Hall.