Examining the Role of Context in Data Interoperability

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Jim Smith, Patrick Place, Marc Novakouski, David Carney
October 26, 2011

© 2011 Carnegie Mellon University

Overview

Background

Research Approach

Discussion

Conclusions

Next Steps

© 2011 Carnegie Mellon University

Background

Motivation for research

- Interoperability continues to be a problem for the DoD
  - Emphasis has been on “filling the gaps,” or imposing commonality (i.e., lexicons, data models, etc.)
- The evolution towards looser coupling between systems of systems constituents (i.e., moving from “directed” to “acknowledged” and “collaborative” systems of systems) has highlighted the limitations with current approaches

The challenges are many

- Most data-centric systems cannot be easily made to interoperate: they weren't designed to work together, or can do so only over a limited range
- This lack of interoperability causes the warfighter to employ workarounds (e.g., “fat fingering” data between incompatible systems)
- The systems of systems that result are brittle – any change (to one or more of the constituent systems, or their operational employment) may “break” interoperability

What is needed is a way to reduce time and cost to achieve data interoperability in a system of systems context, while allowing for greater independence of the constituent systems

- *Without* requiring common data models, ontologies, lexicons, etc.

© 2011 Carnegie Mellon University

Background: Guiding Scenario - 1

Take a really simple example (drawn from real life):

- Two systems (A and B) process and displayed track information
  - System A is on an airplane; system B is on a ship
  - Both systems are “interoperable,” and have all the usual and customary documentation
- A mission thread (T) requires A and B to exchange air track information to create a common tactical picture
- Using a standard messaging format over a common tactical data link, they exchange track data, and integrate those tracks held by the other system into a common view
  - Interoperability is achieved in a specific instantiation of mission thread T

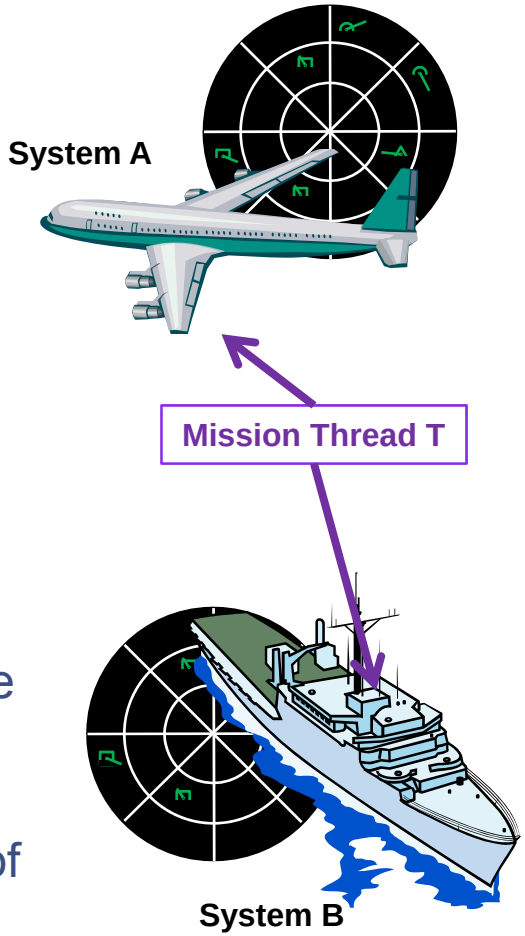

© 2011 Carnegie Mellon University

Background: Guiding Scenario – 2a

What happens when we change the mission thread context to include surface tracks (designated by T£)?

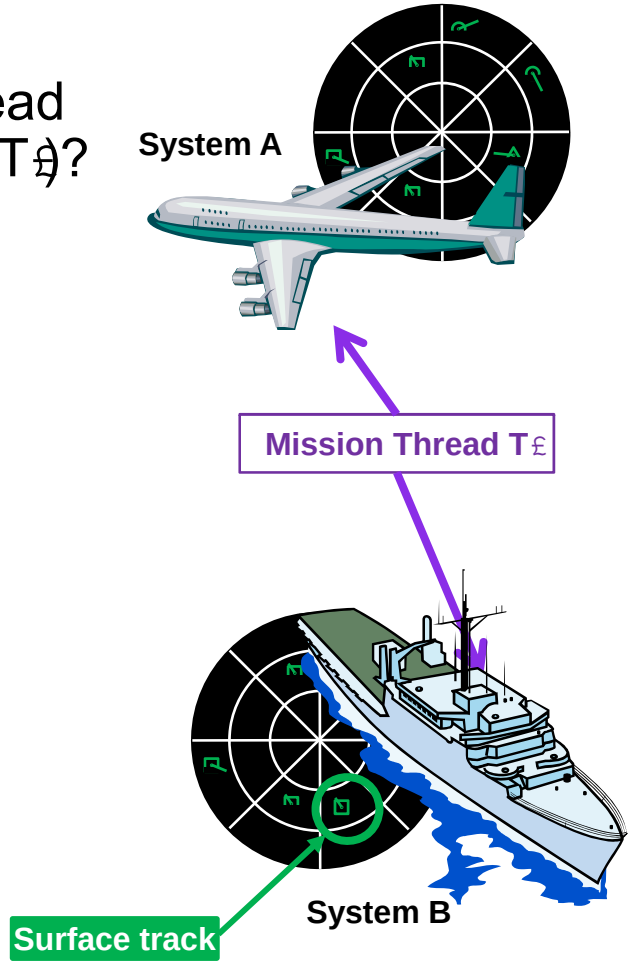

© 2011 Carnegie Mellon University

Background: Guiding Scenario – 2b

What happens when we change the mission thread context to include surface tracks (designated by T£)?

- Though not apparent from the available documentations, system A didn't correctly handle surface tracks
  - There was an unstated contextual constraint on system A that all tracks were assumed to be air tracks. As a consequence, system A's data models, ontology, interfaces, etc. didn't account other types of tracks

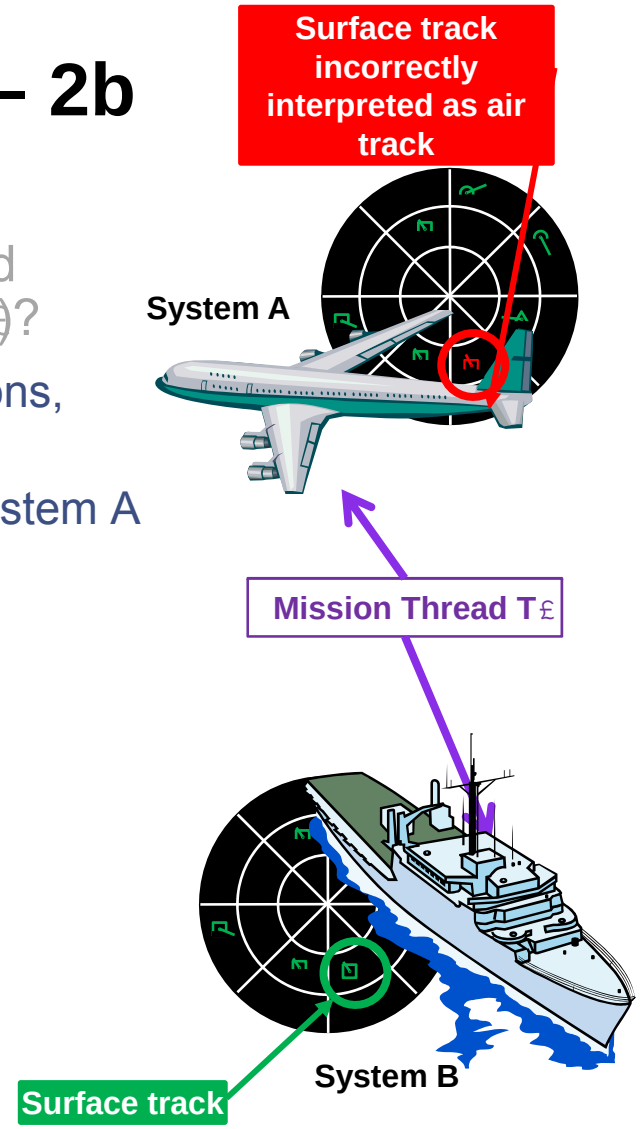

© 2011 Carnegie Mellon University

Background: Guiding Scenario – 2c

What happens when we change the mission thread context to include surface tracks (designated by T£)?

- Though not apparent from the available documentations, system A didn't correctly handle surface tracks
  - There was an unstated contextual constraint on system A that all tracks were assumed to be air tracks. As a consequence, system A's data models, ontology, interfaces, etc. didn't account other types of tracks
  - Ironically, for compatibility with the messaging format, system A **did** force the appropriate message field to "air track" when passing track information (resulting in system B seeing "phantom" air tracks duplicating surface tracks)
    - Discovering this required reading detailed system specifications, technical requirements, and (ultimately) source code for system A

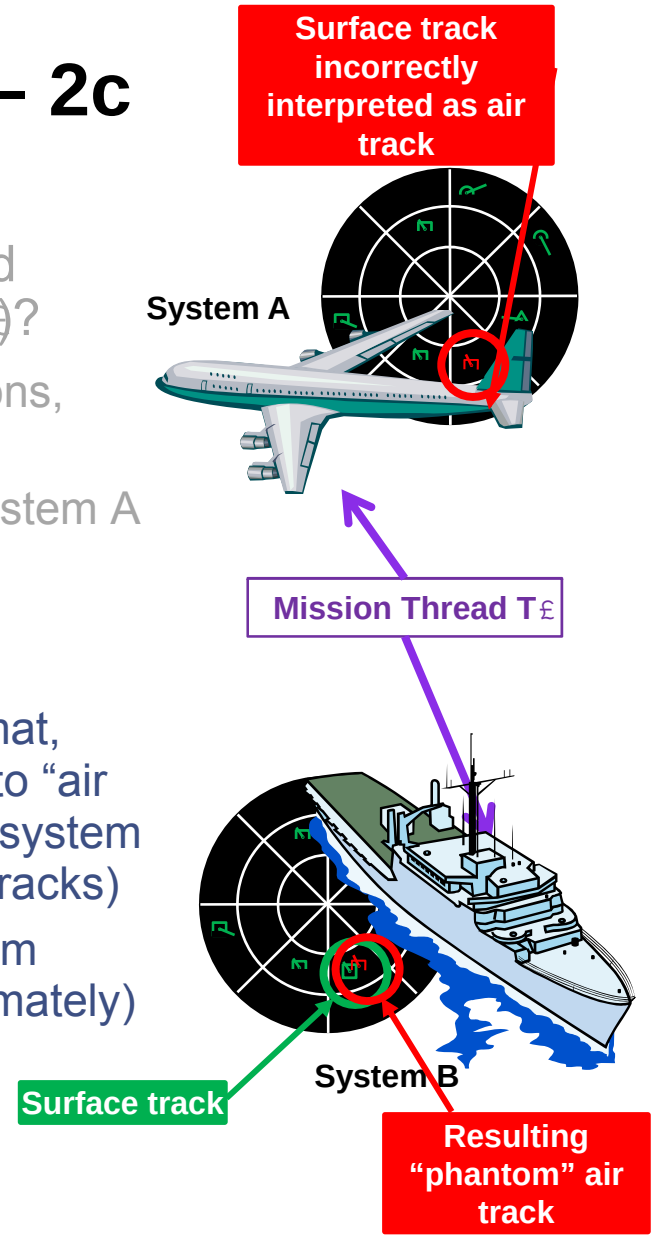

© 2011 Carnegie Mellon University

Background: Guiding Scenario – 2d

What happens when we change the mission thread context to include surface tracks (designated by T£)?

- Though not apparent from the available documentations, system A didn't correctly handle surface tracks
  - There was an unstated contextual constraint on system A that all tracks were assumed to be air tracks. As a consequence, system A's data models, ontology, interfaces, etc. didn't account other types of tracks
  - Ironically, for compatibility with the messaging format, system A **did** force the appropriate message field to "air track" when passing track information (resulting in system B seeing "phantom" air tracks duplicating surface tracks)
    - Discovering this required reading detailed system specifications, technical requirements, and (ultimately) source code for system A

Had the mission thread context remained air tracks only, then systems A and B would have remained interoperable *in the original context* (T)

This unrecognized contextual limitation resulted in a failure of interoperability – no common tactical picture

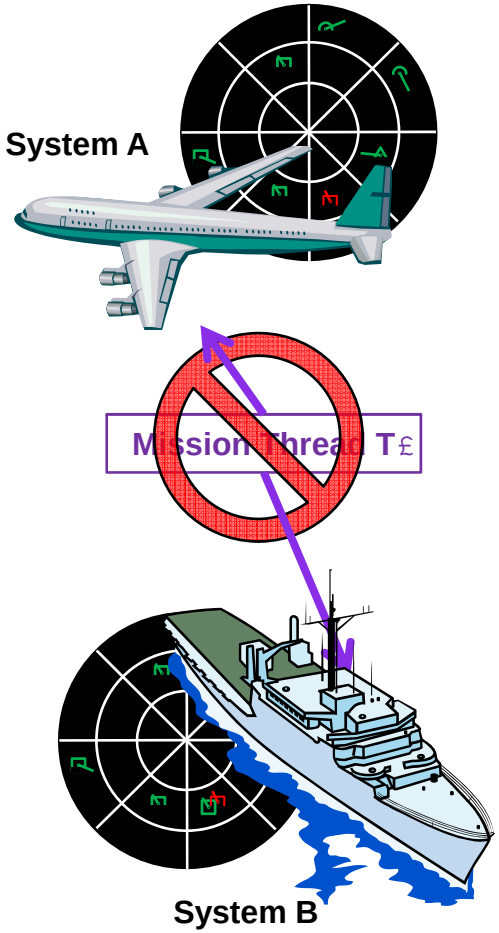

© 2011 Carnegie Mellon University

Research Approach

Question: Can we represent context in a mathematically rigorous manner so that it can be used to improve understanding and reasoning about interoperability?

Approach:

- Characterize the relationship between context and interoperability
- Identify the aspects of context that are necessary to understand interoperability
  - Develop a schema to represent context
- Develop a mathematical framework for expressing and reasoning about context and interoperability
- Test this approach with a *very* simple case
  - Does it work???

© 2011 Carnegie Mellon University

Context - 1

As illustrated by the guiding scenario, a seemingly insignificant change in the system of systems (SoS) operational environment, patterns of use, or in the constituent systems can have a major impact on interoperability

- In SoSs, *all* of these are continually changing

Conventional representations of interoperability (e.g., DoDAF views OV-6, SV-6) focus on the relationships between the constituent systems: filling the gaps in data formats, information exchange requirements, etc.

- The SoS context emerges as a by-product of these views
- As a result, the consequences of a change in the SoS context may not be immediately apparent

Similarly, conventional representations of context focus on the "situation of an entity," and don't encompass many of the factors that are important to understanding interoperability, including:

- Constraints on permissible behaviors and operational environment
- Pragmatic aspects of interoperable data exchanges
- Assumptions made by the developers, users, etc.

© 2011 Carnegie Mellon University

Context – 2a

What if we make context a “first class” object, as opposed to something that is inferred?

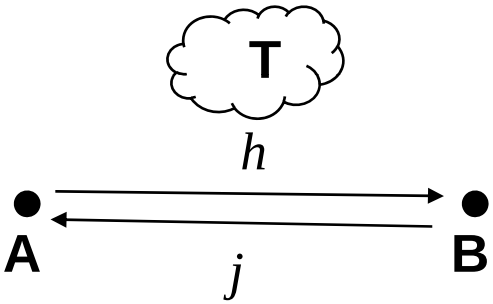

© 2011 Carnegie Mellon University

Context – 2b

What if we make context a “first class” object, as opposed to something that is inferred?

Doing so allows us to shift the focus from “filling the gaps” between systems to understanding how systems relate to the SoS operational thread context

- The relationships between the systems - “the gaps” - become a byproduct of the relationships between the systems and the SoS context, rather than the other way around

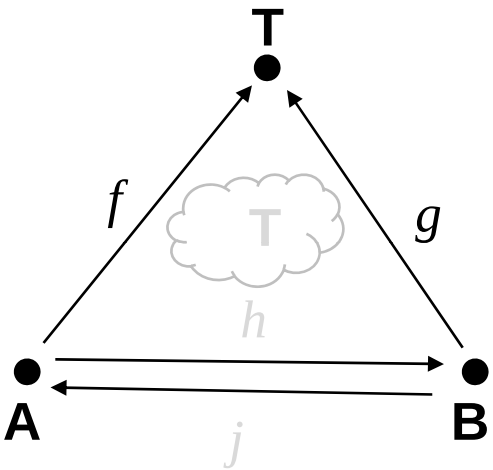

© 2011 Carnegie Mellon University

Context - 3

The context, T, has an ontology, data models, functions/behaviors, and constraints - just like any system

- For convenience, we can represent these with a "contextual schema," consisting of a "signature" and constraints:
  - This signature is analogous to the signature of an algebra or a logic system, and consists of
    - O : Ontology for the domain of discourse
    - D: Data models for the information exchanges
    - F: Functional behaviors over the ontology and data models
  - The constraints define any restrictions on the allowable behaviors for the
    - C: Contextual constraints
- Combining the signature and constraints, we obtain

$$Sch_x = (Sig_x, C_x) = (O_x, D_x, F_x, C_x)$$

© 2011 Carnegie Mellon University

Context - 4

We can also see that the SoS context is separate and distinct from the context of any of the constituent systems:

System of Systems	System
O_T: Ontology for *SoS* domain of discourse	O_{Sys}: Ontology for *system* domain of discourse
D_T: Data models for the *SoS information exchanges*	D_{Sys}: *System* data models
F_T: Set of functional behaviors that implement an SoS *operational thread*	F_{Sys}: *System* functional behaviors
C_T: *SoS operational thread* contextual constraints	C_{Sys}: *System* contextual constraints

Note: $C_T \neq \hat{u}\ (C_{Sys_i})$ (likewise for O,D, and F)

] $Sch_T \neq \hat{u}\ (Sch_{Sys_i})$

© 2011 Carnegie Mellon University

Interoperability - 1

Interoperability can be expressed as mappings between *logical systems*:

- Interoperability represents truth-preserving mappings from the contextual schema of one logical system to another or, as previously shown, between the systems and the SoS

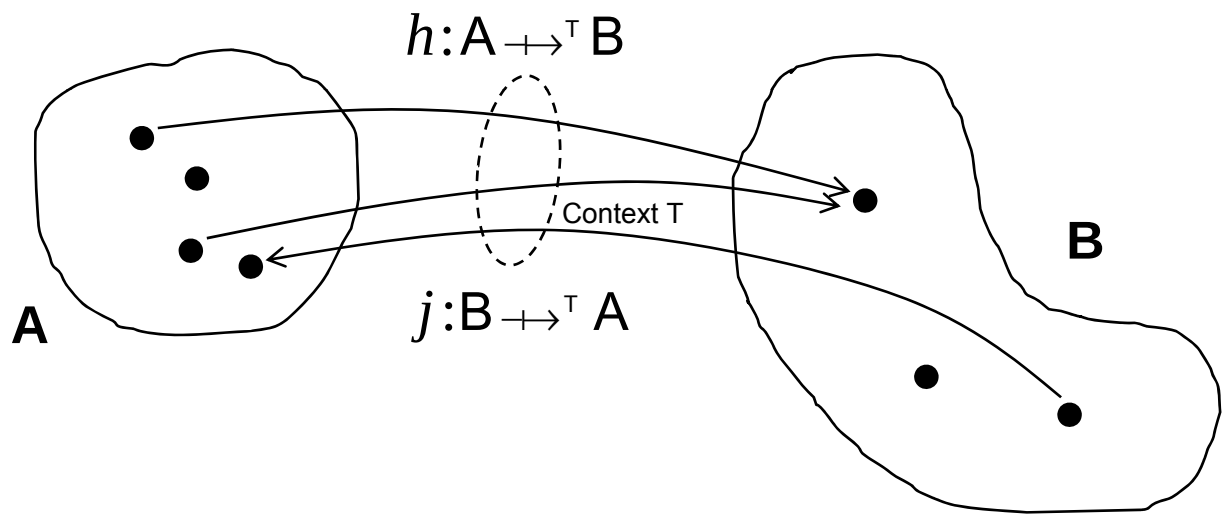

Note: these mappings are *partial* functions—interoperability doesn't require that everything be mapped from one logical system to another, *just those aspects that are relevant to a specified context*

© 2011 Carnegie Mellon University

Category Theory "Lite" - 1

By expressing interoperability in this way, we are able to leverage a branch of mathematics, known as *category theory*, to simplify how we define and reason about interoperability

Category theory provides a rigorous approach to understanding the relationships between objects and their morphisms (transformations)

- Objects can be sets, algebras, or even contextual schemas
- Morphisms are mappings between objects that preserve their structure and meet certain requirements (i.e., identity, associativity)

Several uses in software engineering

- Specification transformation and module composition
- Database schema transformation

Provides a powerful tool to "prove" interoperability in a given—or in a new or novel—context

- Allows us to understand what is happening in interoperable information exchanges
  - Both semantic and pragmatic aspects of interoperability
- Helps us identify how and why interoperability has failed

© 2011 Carnegie Mellon University

Category Theory "Lite" - 2

University of Toronto — Department of Computer Science

Definition of a Category

- **A category consists of:**
  - → a class of *objects*
  - → a class of *morphisms* ("arrows")
  - → for each morphism, *f*, one object as the *domain* of *f* and one object as the *codomain* of *f*.
  - → for each object, *A*, an *identity morphism* which has domain *A* and codomain *A*. ("ID_A")
  - → for each pair of morphisms f:A→B and g:B→C, (i.e. cod(f)=dom(g)), a *composite morphism*, g ∘ f: A→C

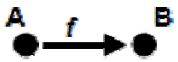

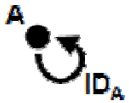

- **With these rules:**
  - → *Identity composition:* For each morphism *f*:A→B,
    $f \circ ID_A = f$ and $ID_B \circ f = f$
  - → *Associativity:* For each set of morphisms f:A→B, g:B→C, h:C→D,
    $(h \circ g) \circ f = h \circ (g \circ f)$

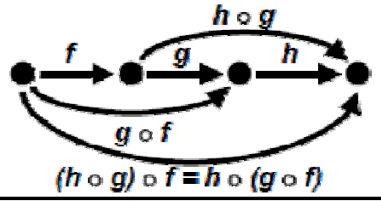

© Steve Easterbrook, 1999

5

From "Category Theory for Software Engineers," Steve Easterbrook, University of Toronto, 1999.

© 2011 Carnegie Mellon University

Interoperability – 2a

Sch_A

O_A

D_A

F_A

C_A

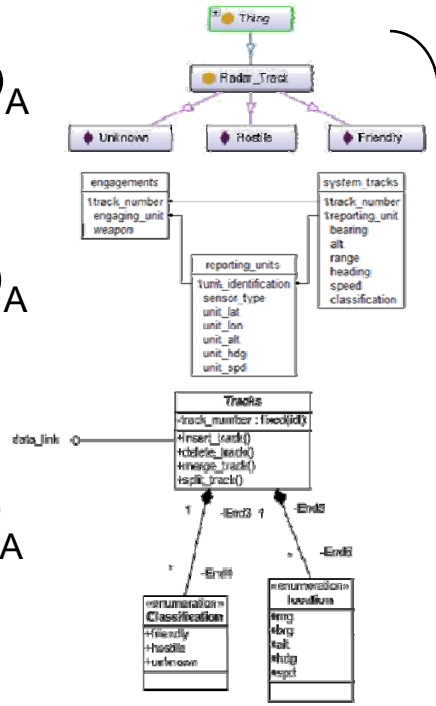

$\bigcirc track_type = air$

$\bigcirc track_alt = \{0..120000\}$

$\bigcirc track_spd = \{0..2500\}$

$\bigcirc track_rng = \{0..500\}$

The morphism $f : A \text{ fi } T$ represents the mapping from the ontology, data models, and functional behaviors of Sch_A to their images in Sch_T, while maintaining consistency with contextual constraints

$$A \bullet \xrightarrow{f} \bullet T$$

Sch_T

O_T

D_T

F_T

C_T

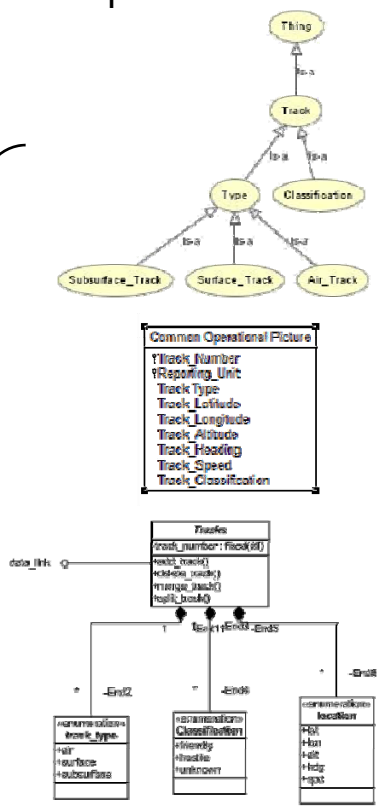

$\bigcirc track_type = \{air, surface, subsurface\}$

$\forall t \in air : \bigcirc track_alt > 0$

$\forall t \in surface \cup subsurface : track_alt = 0$

© 2011 Carnegie Mellon University

Interoperability – 2b

Sch_A $\quad A \xrightarrow{f} T \quad$ Sch_T

O_A

D_A

F_A

C_A

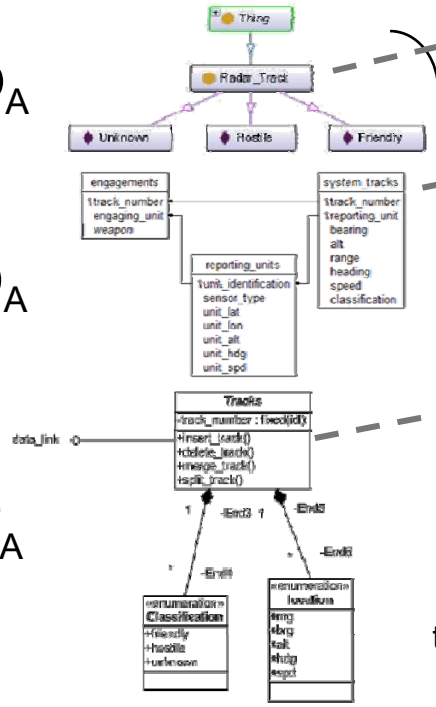

$\bigcirc track_type = air$

$\bigcirc track_alt = \{0..120000\}$

$\bigcirc track_spd = \{0..2500\}$

$\bigcirc track_rng = \{0..500\}$

Radar_Track fi Track

system_tracks fi Common Operational Picture

add_track() fi insert_track()

delete_track() fi delete_track()

...

track_type=air ⌀track_type={air, surface, subsurface} fi

track_type=air

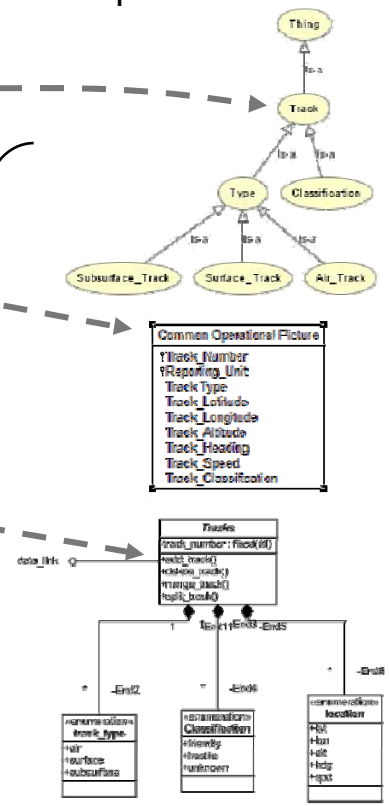

$\bigcirc track_type = \{air, surface, subsurface\}$

$\forall t \in air : \bigcirc track_alt > 0$

$\forall t \in surface \cup subsurface : track_alt = 0$

O_T

D_T

F_T

C_T

© 2011 Carnegie Mellon University

Interoperability – 3a

A concrete example can help illustrate this…

We have objects (systems A and B; SoS operational thread T) in the category of contextual schemas:

A= |a_1, a_2~ (set of air tracks held by A)

B= |b_1 ,b_2, b_3~ (set of air (b_1, b_2) and surface (b_3) tracks held by B)

T= |t_1, t_2~ (set of tracks resulting from the mapping of A and B to the context of T)

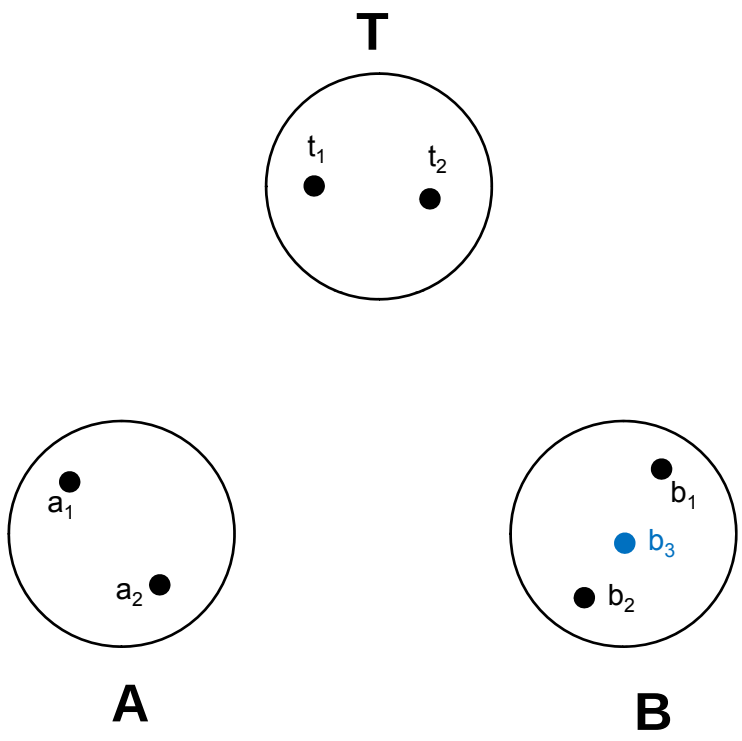

© 2011 Carnegie Mellon University

Interoperability – 3b

A concrete example can help illustrate this…

We have objects (systems A and B; SoS operational thread T) in the category of contextual schemas:

A= |a_1, a_2~ (set of air tracks held by A)

B= |b_1 ,b_2, b_3~ (set of air (b_1, b_2) and surface (b_3) tracks held by B)

T= |t_1, t_2~ (set of tracks resulting from the mapping of A and B to the context of T)

and morphisms (mappings between them):

h : A fi B = |$a_1$fi b_1~ (track reported by A to B)

j : B fi A = |$b_2$fi a_2~ (track reported by B to A)

f : A fi T = |$a_1$fi t_1, $a_2$fi t_2~ (The mapping from A to the context T)

g : B fi T = |$b_1$fi t_1, $b_2$fi t_2~ (The mapping from B to the context T)

Note: Track b_3 is not reported by B to A because it is not part of context T

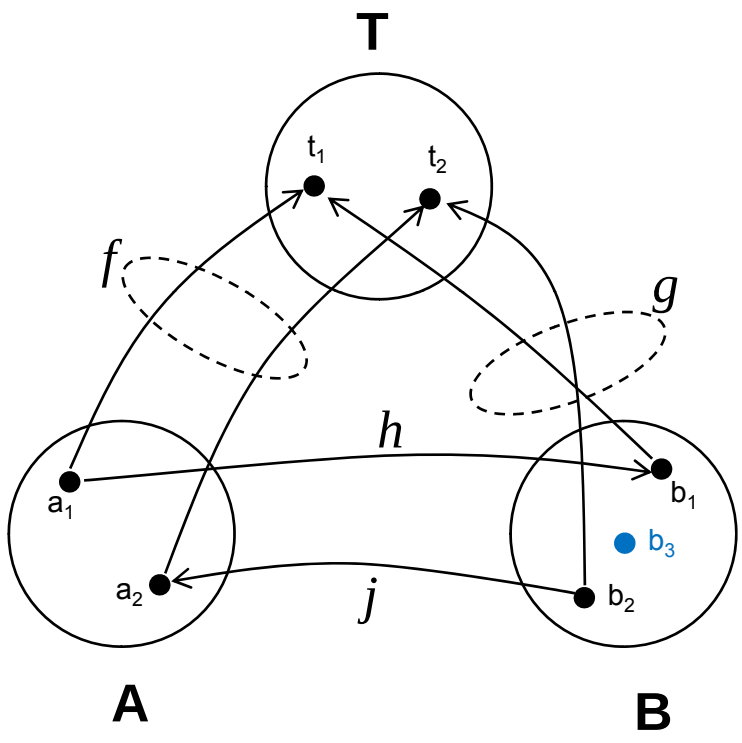

© 2011 Carnegie Mellon University

Interoperability – 3c

A concrete example can help illustrate this…

We have objects (systems A and B; SoS operational thread T) in the category of contextual schemas:

A= |a_1, a_2~ (set of air tracks held by A)

B= |b_1 ,b_2, b_3~ (set of air (b_1, b_2) and surface (b_3) tracks held by B)

T= |t_1, t_2~ (set of tracks resulting from the mapping of A and B to the context of T)

and morphisms (mappings between them):

h : A fi B = |$a_1$fi b_1~ (track reported by A to B)

j : B fi A = |$b_2$fi a_2~ (track reported by B to A)

f : A fi T = |$a_1$fi t_1, $a_2$fi t_2~ (The mapping from A to the context T)

g : B fi T = |$b_1$fi t_1, $b_2$fi t_2~ (The mapping from B to the context T)

Note: Track b_3 is not reported by B to A because it is not part of context T

By inspection, we see that

A fi T = A fi B fi T (or $f = g \,''\, h$)

and

B fi T = B fi A fi T (or $g = f \,''\, j$)

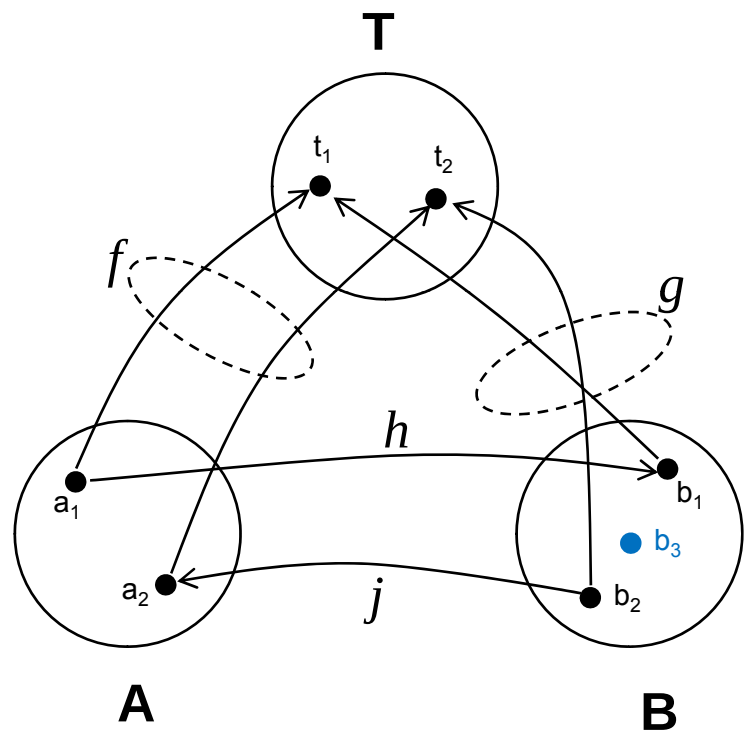

© 2011 Carnegie Mellon University

Category Theory "Lite" - 3

From our guiding scenario, A and B represent two systems, and T represents the SoS operational context within which they must interoperate

- Node P (the "pullback* of A and B over T") represents the interoperability between A and B in context T
  - A and B exchange track data in context T so that both systems hold the same tracks

When we changed the SoS context to include surface tracks, interoperability failed

- Since A didn't properly handle surface tracks, A and B had different tactical pictures
  - In this case, P would highlight the fact that the "common" picture was actually not common

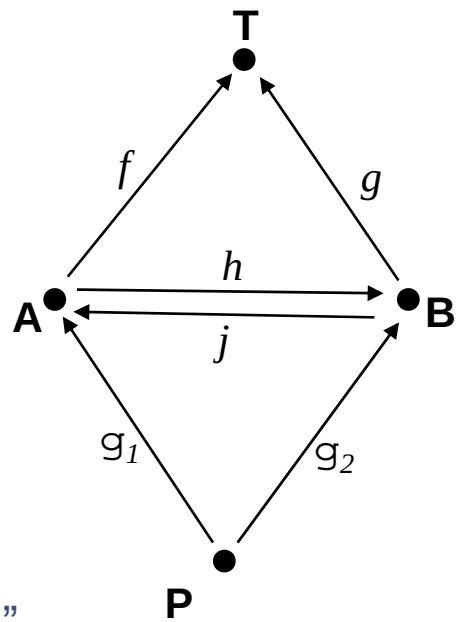

* The pullback is a special case of the more general "limit"

© 2011 Carnegie Mellon University

Interoperability – 4a

Continuing our example…

Construct the pullback, P

- In context T, P = |ćt$_1$,ća$_1$, b$_1$®®, ćt$_2$,ća$_2$, b$_2$®®~
- The pullback of A and B over T shows the tracks jointly held by A and B in context T: the common tactical picture

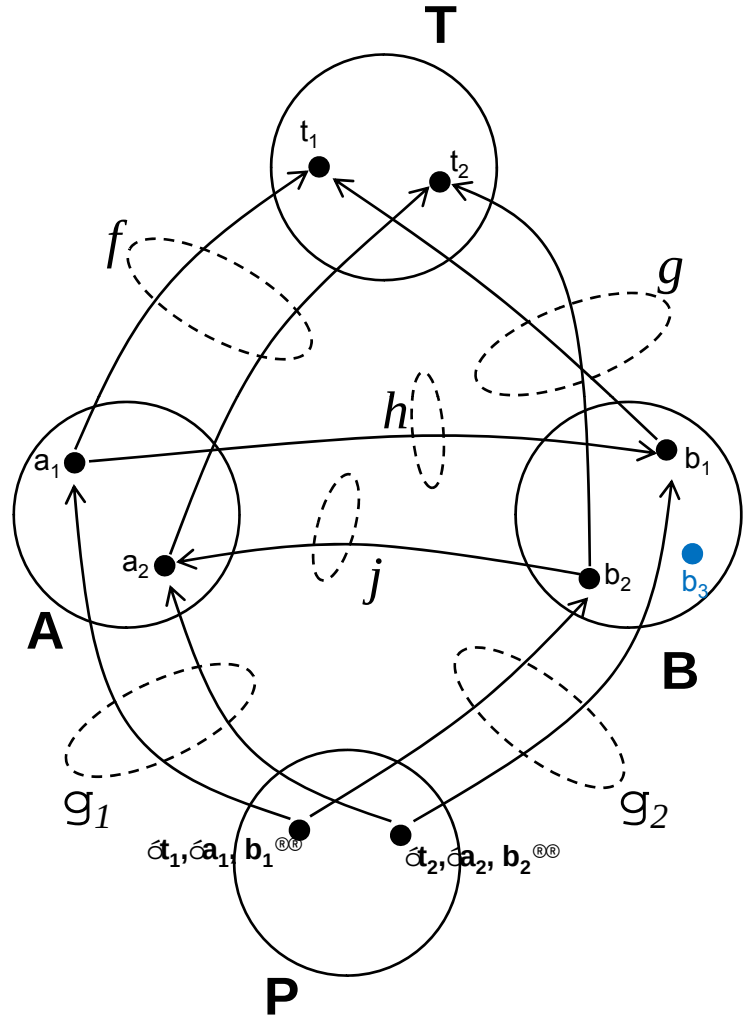

© 2011 Carnegie Mellon University

Interoperability – 4b

Continuing our example…

Construct the pullback, P

- In context T, P = |ćt$_1$,ća$_1$, b$_1$®®, ćt$_2$,ća$_2$, b$_2$®®~
- The pullback of A and B over T shows the tracks jointly held by A and B in context T: the common tactical picture

Now, change the mission thread context to include surface tracks (T£), and construct a new pullback (P£):

- b$_3$ is a surface track held by B which, unlike the previous case, is *now* part of the context (T£)
- A incorrectly interprets it as an air track, a$_3^*$
- A subsequently passes this invalid track back to B as a "phantom" air track, b$_3^*$

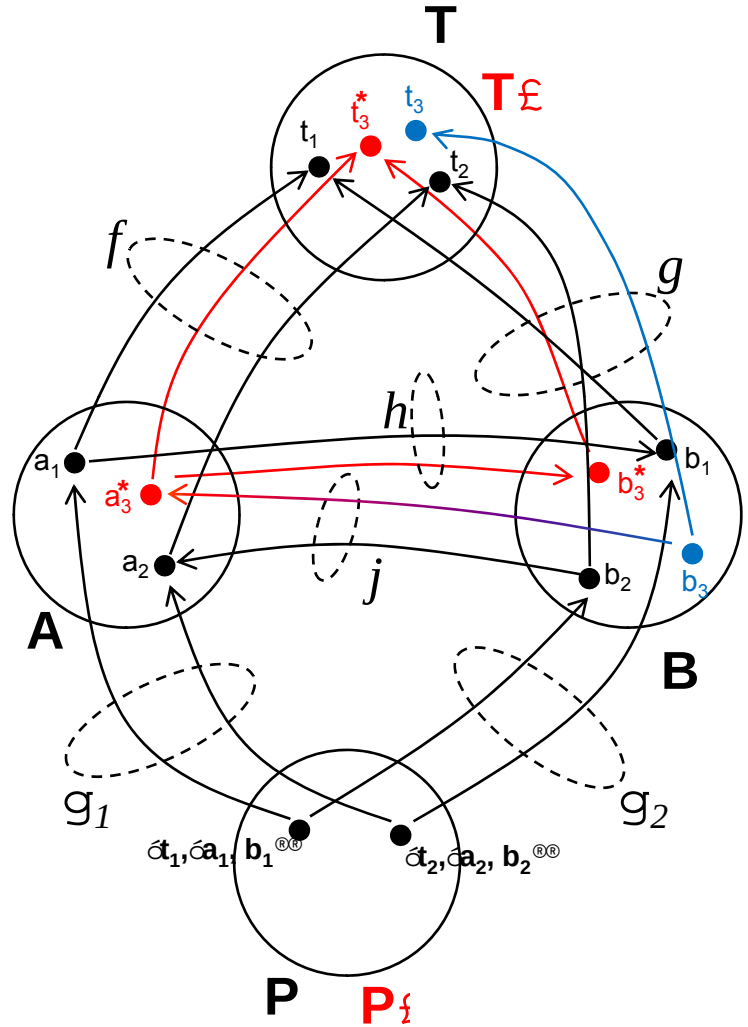

© 2011 Carnegie Mellon University

Interoperability – 4c

Continuing our example…

Construct the pullback, P

- In context T, P = |ćt$_1$,ća$_1$, b$_1$®®, ćt$_2$,ća$_2$, b$_2$®®~
- The pullback of A and B over T shows the tracks jointly held by A and B in context T: the common tactical picture

Now, change the mission thread context to include surface tracks (T£), and construct a new pullback (P£):

- b$_3$ is a surface track held by B which, unlike the previous case, is *now* part of the context (T£)
- A incorrectly interprets it as an air track, a$_3^*$
- A subsequently passes this invalid track back to B as a "phantom" air track, b$_3^*$

Thus, in context T£ the revised pullback P£ = |ćt$_1$,ća$_1$, b$_1$®®, ćt$_3^*$,ća$_3^*$, b$_3^*$®®, ćt$_2$,ća$_2$, b$_2$®®~

- Note that new pullback reflects the "phantom" track (t$_3^*$) resulting from the incorrect interpretation of surface track b$_3$ by system A – no longer have a common tactical picture

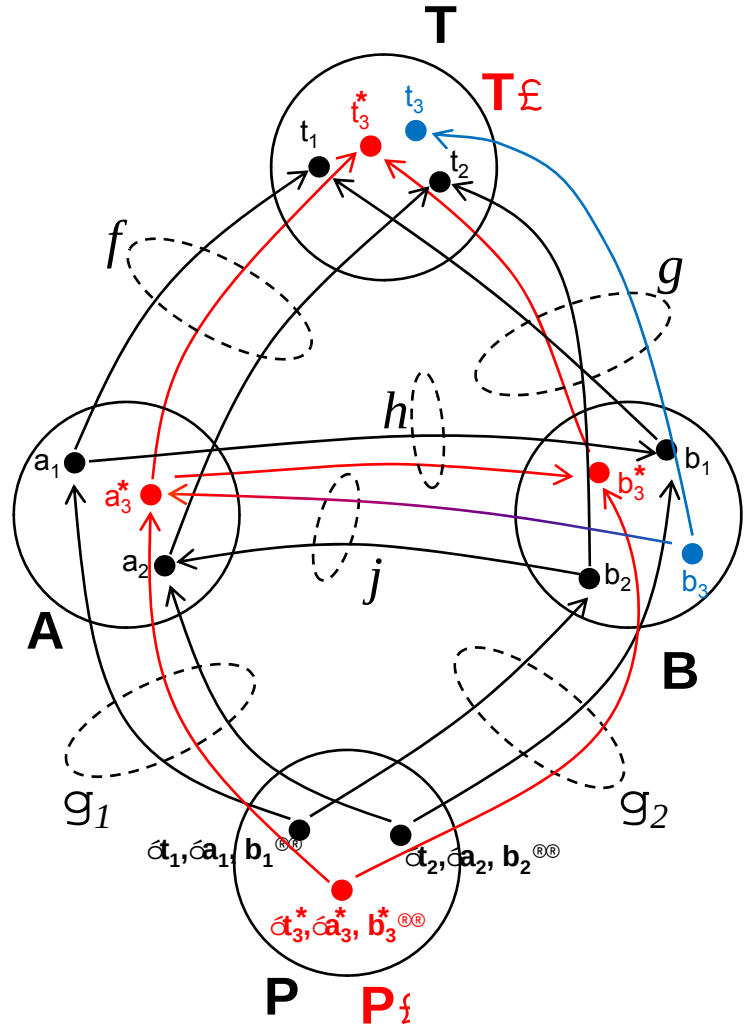

© 2011 Carnegie Mellon University

Reasoning Framework - 1

We can represent the contextual schemas using an ontology language (e.g., RDF, OWL), XML schema and applicable DoDAF architecture views:

- OV-3
- OV-6
- SV-6

} various aspects of the data exchange semantics

- SV-10: system constraints/permissible behaviors
- SvcV-6 and SvcV-10: same as SV-6 and SV-10 for services

The Rei ontology* provides a language for expressing constraints as policies, including conflict resolution (e.g., “A should do x, unless it is overridden by a higher-priority demand”)

```
<deontic:Permission rdf:ID="Perm_StudentPrinting"> <deontic:actor rdf:resource="#PersonVar"/> <deontic:action
rdf:resource="#ObjVar"/> <deontic:constraint rdf:resource="#IsStudentAndBWPrinter"/> </deontic:Permission>
```

Deontic logic then provides a framework to reason about constraints

- Allows us to “prove” that interoperability between systems can exist in a given context, and identify its bounds

* http://www.cs.umbc.edu/~lkagal1/rei/

© 2011 Carnegie Mellon University

Reasoning Framework - 2

From our guiding scenario, we define the following database schemas:

System A

```
schema Tracks {
        interface Track {
                int tracknum();
                int lat();
                int lon();
                int alt();
                int hdg();
                int spd();
                Set<Track> tracks();
                }
        Collection<Track> dbTracks;
        Constraints:
        Tracks X,Y;
        tracknum(X) != tracknum(Y);

}
```

System B

```
schema Tracks {
        interface Track {
                int track_number();
                int track_type();
                int track_position();
                int alt();
                int hdg();
                int speed();
                Set<Track> tracks();
                }
        Collection<Track> dbTracks;
        Constraints:
        Tracks X,Y;
        tracknum(X) != tracknum(Y);
        track_type(X) = {surface, air, subsurface}
}
```

Context T

```
schema Tracks {
        interface Track {
                int track_no();
                int track_type();
                int track_lat();
                int track_lon();
                int track_alt();
                int track_head();
                int track_spd();
                Set<Track> tracks();
                }
        Collection<Track> dbTracks;
        Constraints:
        Tracks X,Y;
        tracknum(X) != tracknum(Y);
        track_type(X) = air

}
```


© 2011 Carnegie Mellon University

Reasoning Framework - 3

In the original context, the pullback of A and B over T is represented by

Tracks =

<track_no, <tracknum, track_number>>
<air, <air, track_type>>
<track_lat, <lat, position.lat>>
<track_lon, <lon, position.lon>>
<track_head, <hdg, head>>

<track_spd, <spd, speed>>

The track type constraints are given by

(i) $\bigcirc_A air$ (system A must handle "air" tracks)

(ii) $\bigcirc_B \left(air \wedge surface \wedge subsurface \right)$ (system B must handle tracks of all three types)

(iii) $\Box_T air$ (it is necessary that tracks be of type "air")

This gives us :

(iv) $\Box_T air \Rightarrow \mathrm{F}_A \neg air \wedge \mathrm{F}_B \neg air$

By inspection, we can see that the constraints on A and B are consistent in this context:

(v) $\bigcirc_A air \Rightarrow \mathrm{F}_A \neg air$

(vi) $\bigcirc_B air \Rightarrow \mathrm{F}_B \neg air$

© 2011 Carnegie Mellon University

Reasoning Framework - 4

When we changed the mission thread context to include other track types (i.e., surface and subsurface), we have

```
Context T'
schema Tracks {
        interface Track {
                int track_no();
                int track_type();
                int track_lat();
                int track_lon();
                int track_alt();
                int track_head();
                int track_spd();
                Set<Track> tracks();
                }
        Collection<Track> dbTracks;
        Constraints:
        Tracks X,Y;
        tracknum(X) != tracknum(Y);
        track_type(X) = {surface, air, subsurface}

}
```

The pullback of A and B over T' is represented by

```
Tracks =
        <track_no, <tracknum, track_number>>
        <track_type, <air, track_type>>
        <track_lat, <lat, position.lat>>
        <track_lon, <lon, position.lon>>
        <track_head, <hdg, head>>
        <track_spd, <spd, speed>>
```


© 2011 Carnegie Mellon University

Reasoning Framework - 5

This gives us a interoperability limitation

(i) $\bigcirc_A air$ (system A tracks must handle "air" tracks)

(ii) $\bigcirc_B(air \wedge surface \wedge subsurface)$ (system B must handle tracks of all three types)

(iii) $\Box_{T'}(air \wedge surface)$ (it is necessary that tracks include "air" and "surface")

As before, we restate the universal necessity as :

(iv) $\Box_{T'}(air \wedge surface) \Rightarrow \bigcirc_A(air \wedge surface) \wedge \bigcirc_B(air \wedge surface)$

However, since system A only supports air tracks, this violates one of its contextual constraints

(v) $\bigcirc_A air \boxplus \exists \mathrm{F}_A \neg air \Rightarrow \mathrm{F}_A\ surface$

(vi) $\bigcirc_B(air \wedge surface \wedge subsurface) \Rightarrow \mathrm{P}_B(air \vee surface)$

So, even though neither system was changed, a change in mission thread context resulted in loss of interoperability

© 2011 Carnegie Mellon University

Conclusions

Treating context as a “first class” object (like a system) allows us to re-think interoperability

- Focus on the relationships between the constituent systems and the SoS context, rather than “filling the gaps” between the systems
- Making context explicit makes it much easier to see how a change in context—either for the SoS, or a constituent system—can affect interoperability
  - SoS and system context can be represented by their contextual schemas

Category theory and deontic logic can be used to reason about interoperability in an SoS context

© 2011 Carnegie Mellon University

Next Steps

We've proven it *can* work—for a very simple problem

- Still need to see if this approach can work for more realistic problems
  - Does it scale up?
  - Is the cost greater than the benefits?
  - Is there some way to do this without requiring program managers to have a Ph.D. in applied mathematics???
- Looking for potential candidates to pilot this approach

A category-theoretic approach can also be employed in reasoning about interoperability in a new SoS context, or with a novel composition of systems

- Instead of a pullback/limit, use a pushout/co-limit

© 2011 Carnegie Mellon University

Contact Information

Jim Smith

Senior Member of the Technical Staff

Telephone: +1 703-908-8221

Email: jds@sei.cmu.edu

Web: www.sei.cmu.edu/staff/jds

Web

www.sei.cmu.edu

www.sei.cmu.edu/contact.cfm

U.S. Mail

Software Engineering Institute

Customer Relations

4500 Fifth Avenue

Pittsburgh, PA 15213-2612

USA

Customer Relations

Email: info@sei.cmu.edu

Telephone: +1 412-268-5800

SEI Phone: +1 412-268-5800

SEI Fax: +1 412-268-6257

© 2011 Carnegie Mellon University

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.