

Knowledge Management with K-Infinity

Jan Schümmer, intelligent views
Darmstadt, Germany
j.schuemmer@i-views.de

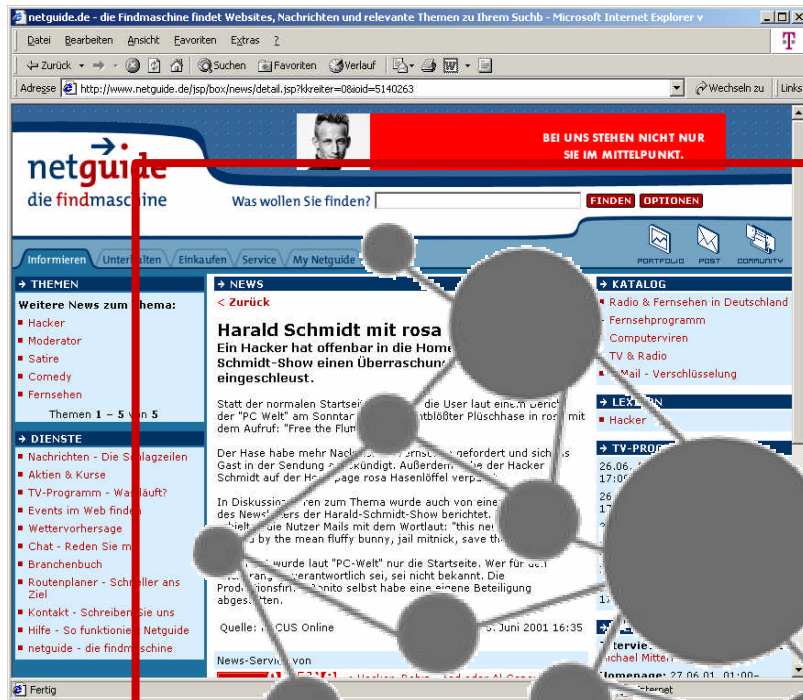
In this talk, I am going to

- show an application - K-Infinity
- give you a short introduction to the COAST framework
- explain how COAST was used to build K-Infinity
- hopefully encourage you to use COAST in your own applications

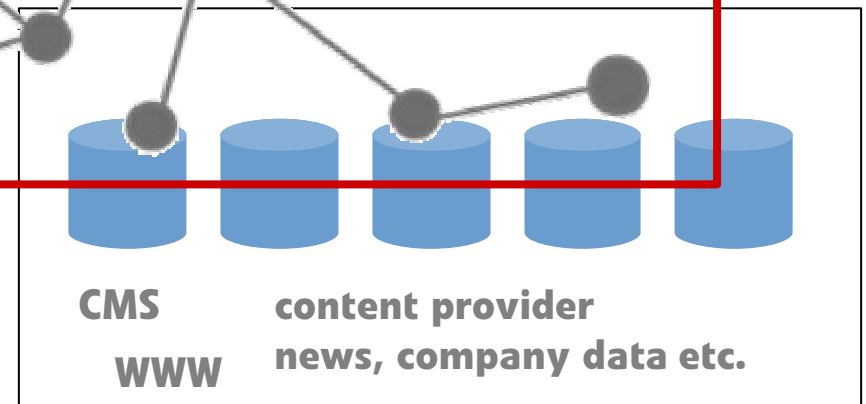
intelligent views

- located in Darmstadt, Germany
- founded in 1997
- spin-off enterprise of GMD – National Research Center for Information Technology, institute IPSI
- about 40 employees

What we are doing



add meaning to your data



Scenario

- You are visiting a classical music news portal to find information about conductors in Bayreuth

Full text query does not help, because

- it only finds strings, not concepts
 - no synonymes
 - no translations
- it produces irrelevant hits
 - homonymes
 - different contexts

A model is needed



K-Infinity allows to

- express knowledge in terms of an object-oriented model
- connect content to arbitrary entities of this model
- use the knowledge net for
 - semantic search
 - presenting related things next to each other
 - ... reasoning

Building blocks

Concept

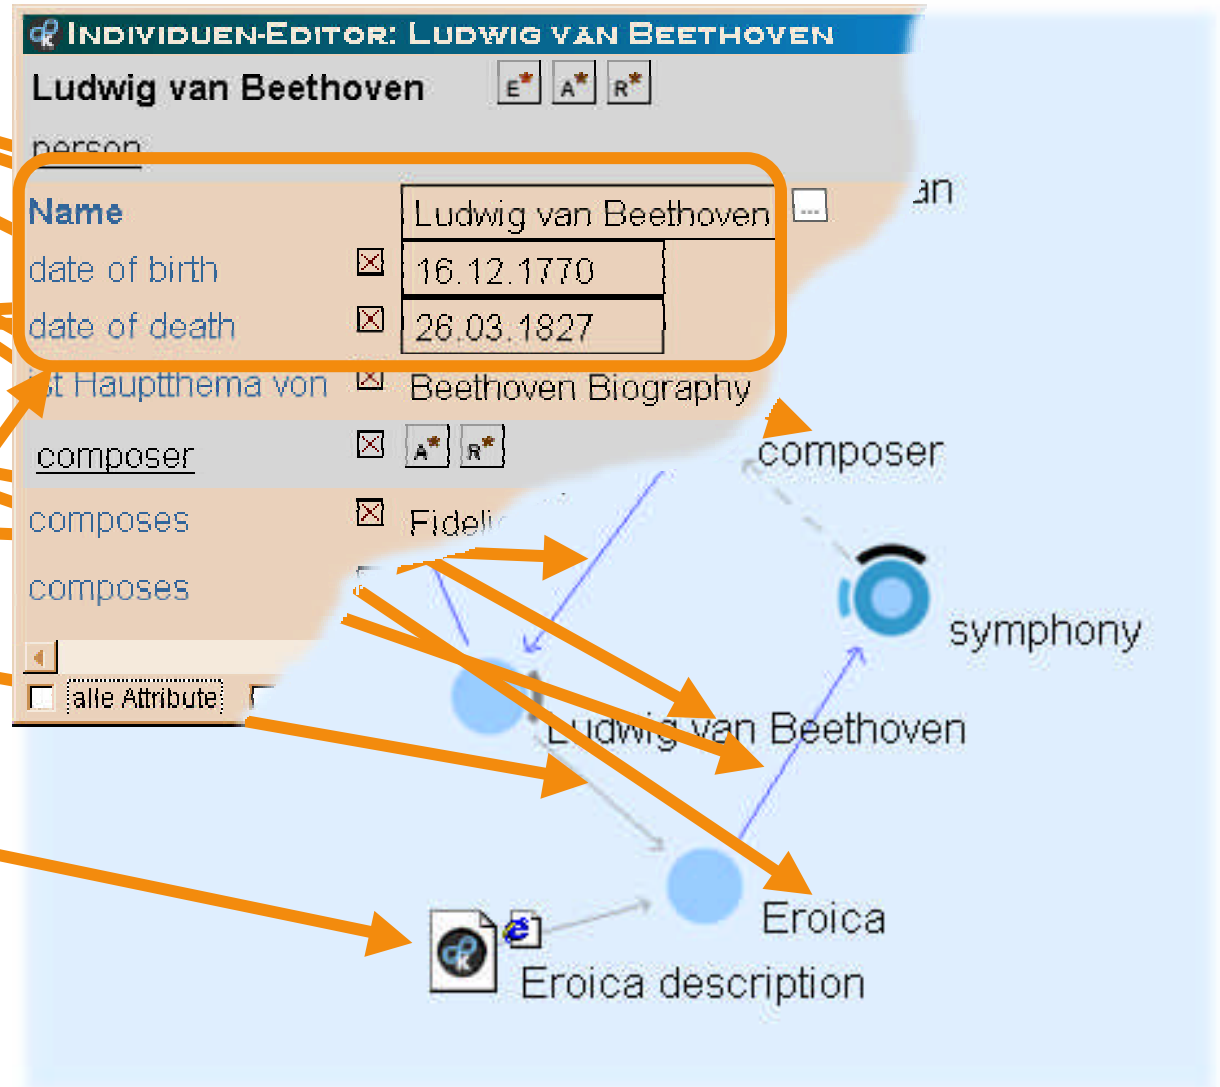
Individual

Relation

- inheritance
- instance of
- extension
- user defined relations

Document

Attribute



- create individual
- create role concept
- define relation for role
- assign role to individual
- create instance of the newly defined relation

- import document
- connect it to the knowledge net
- show matching environment (related documents)

K-Infinity tool suite



K-Infinity

Knowledge Builder

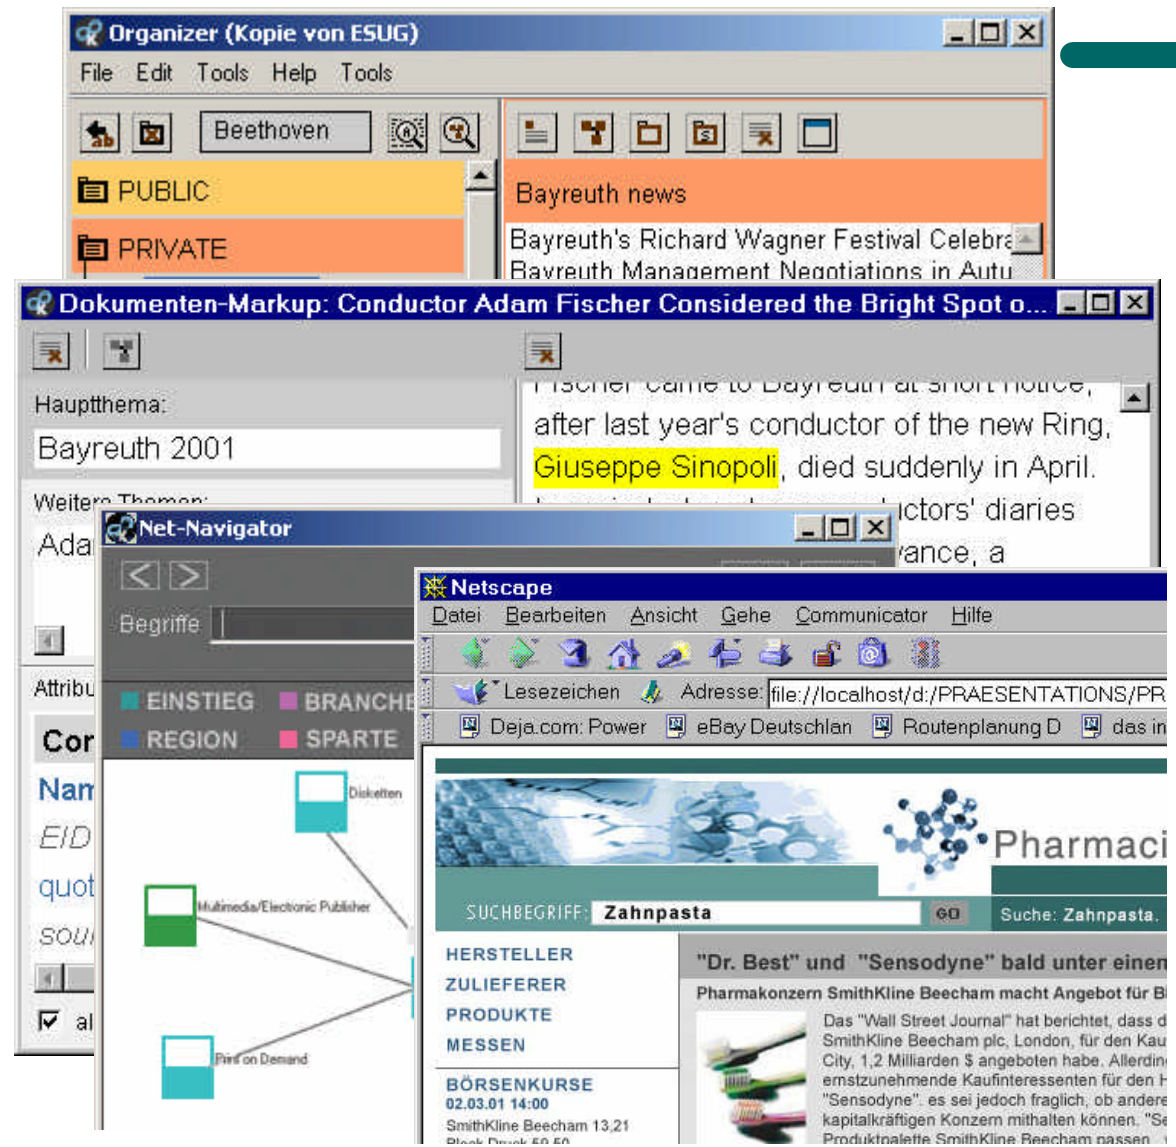
- schema definition

Markup Tool

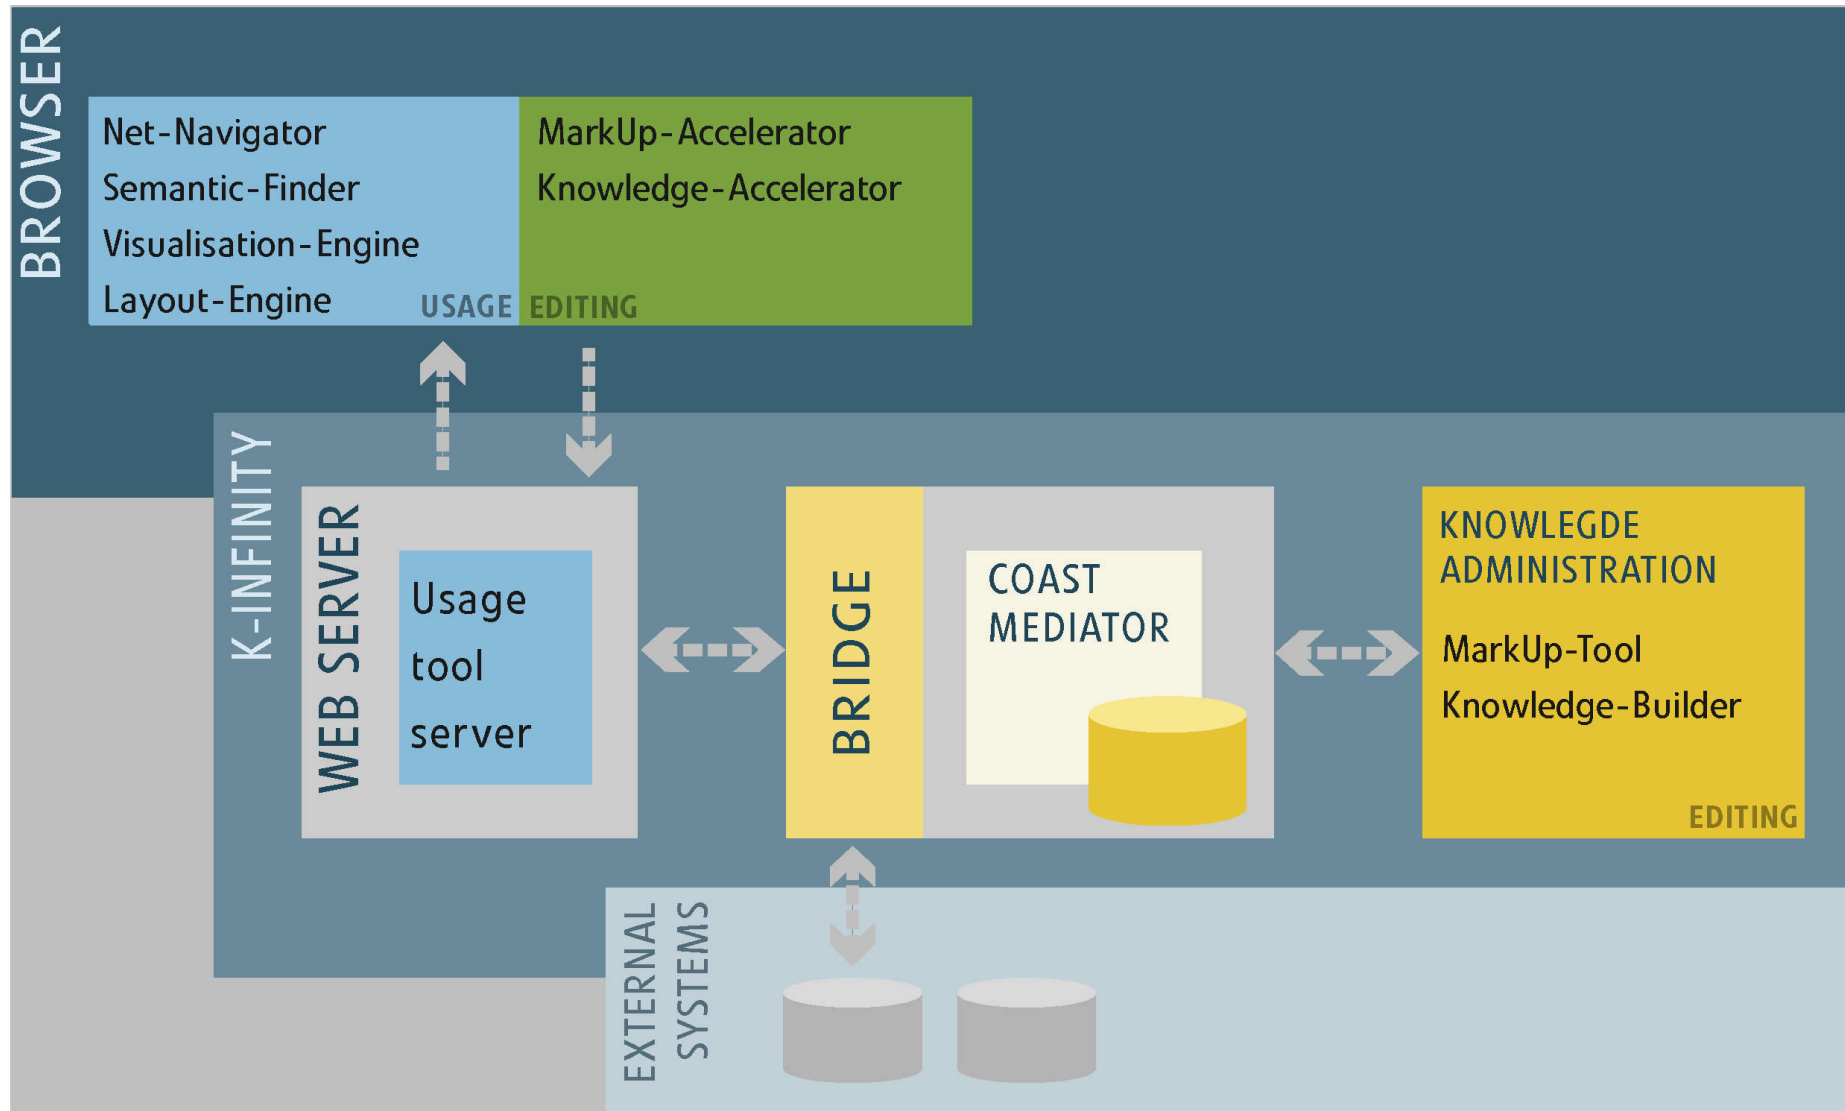
- link documents to the net

Usage tools

- Net Navigator
- Knowledge Accelerator
- Web presentation engines



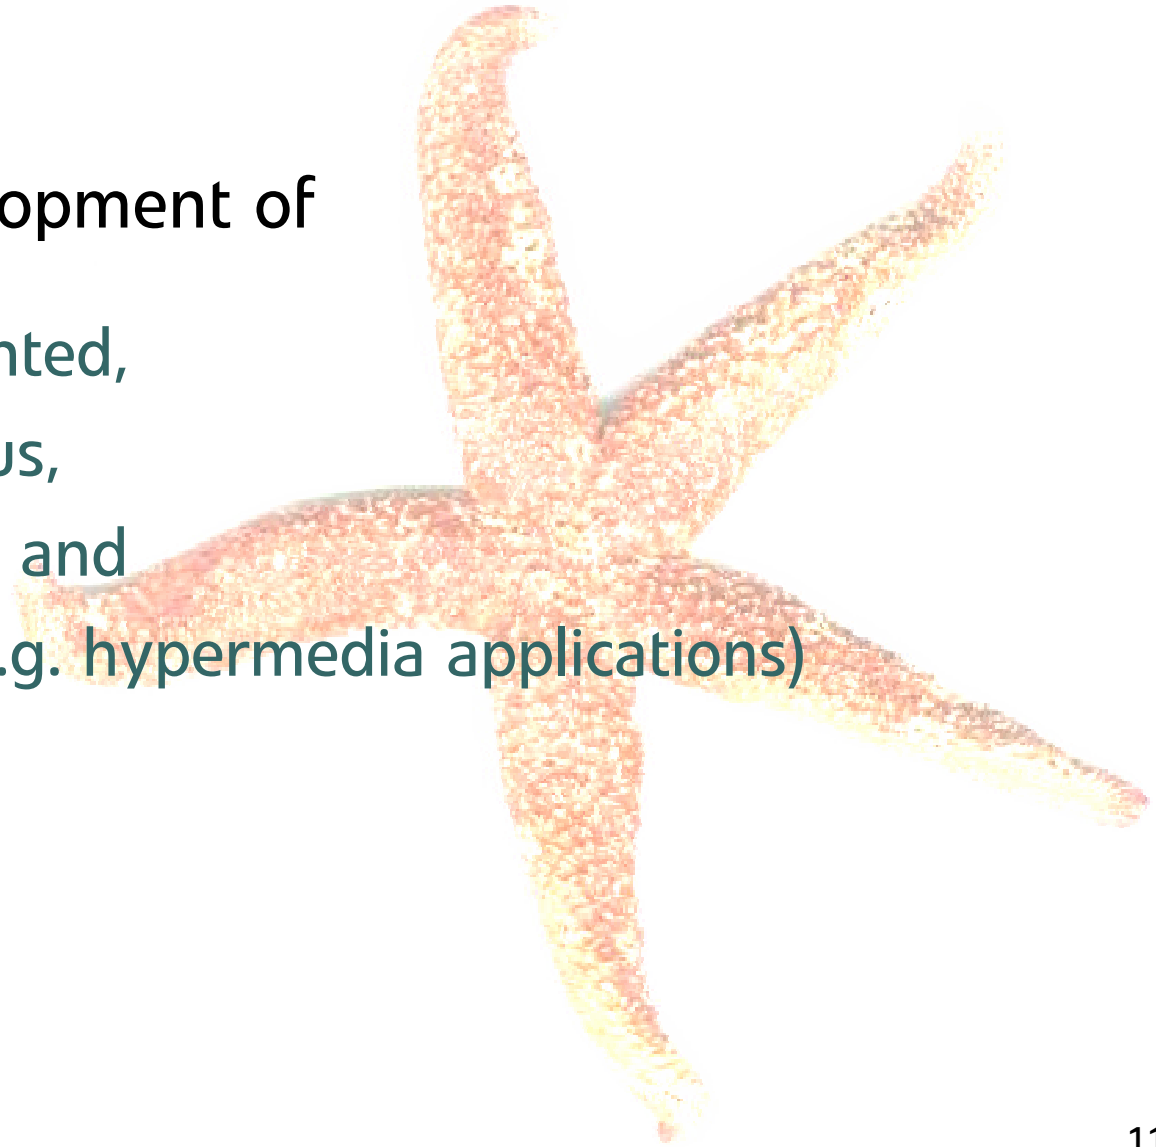
Behind the Scenes



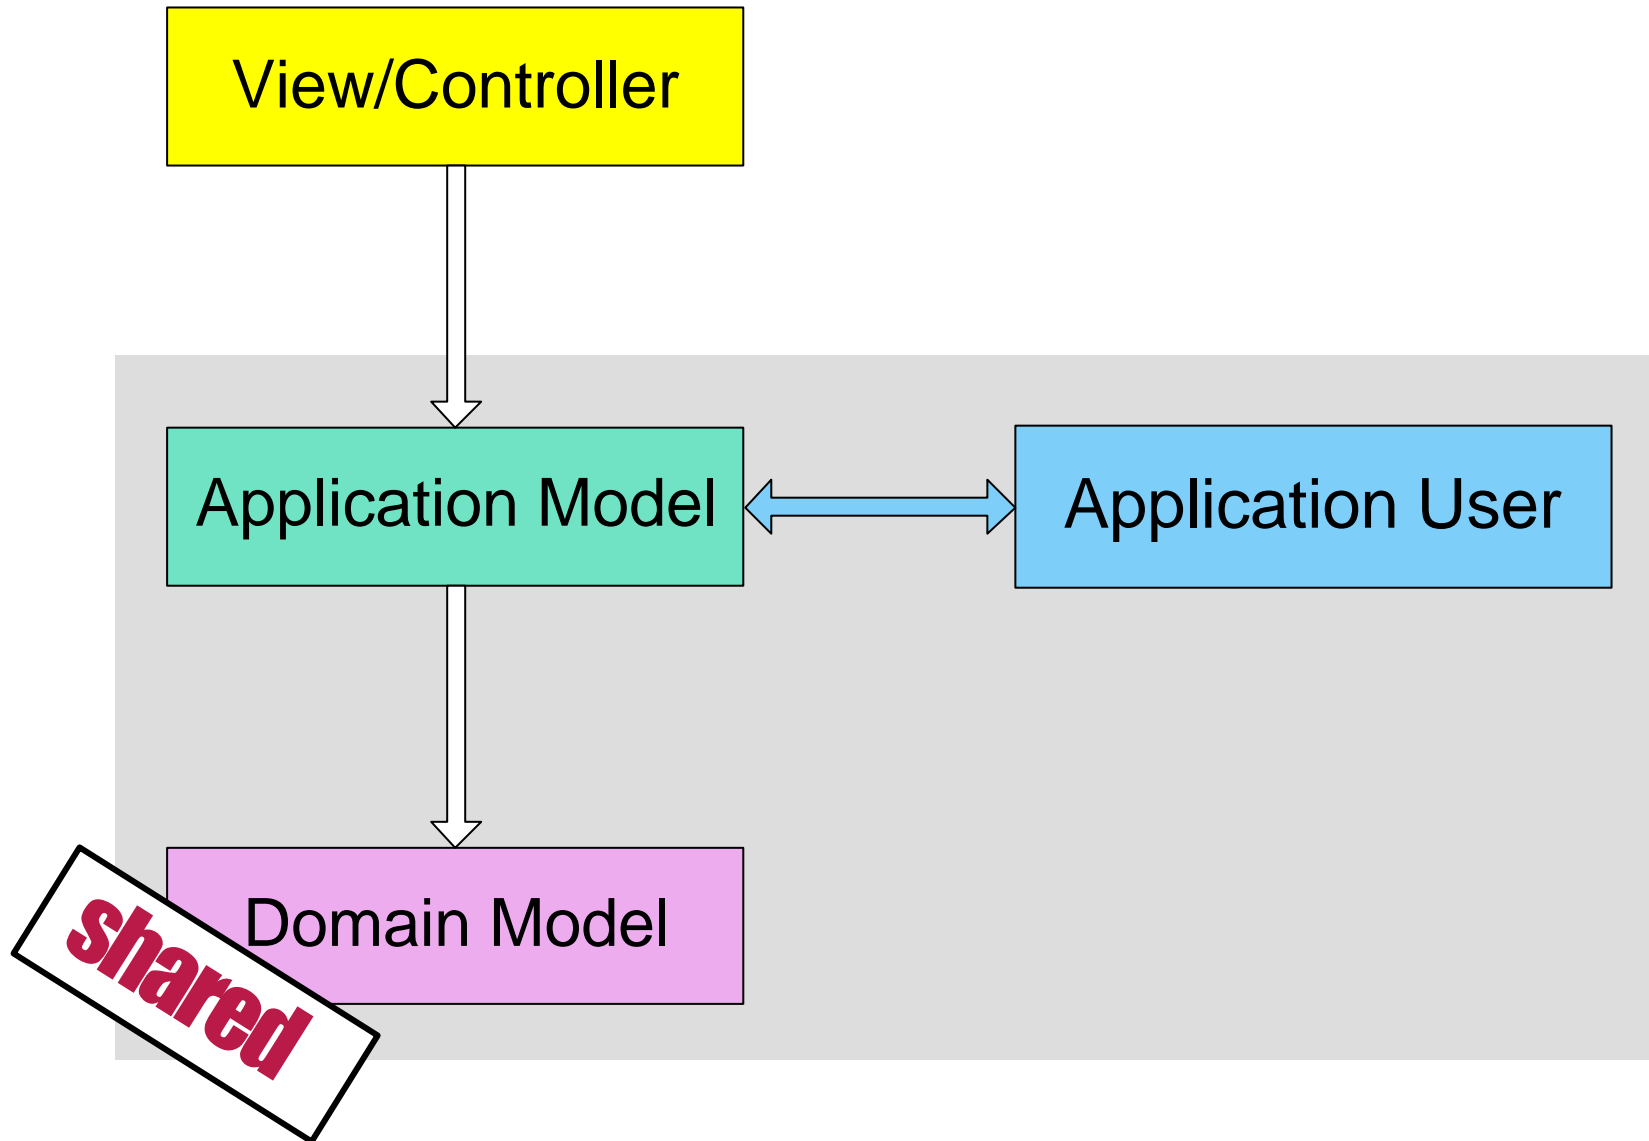
Support the development of

- object oriented,
- synchronous,
- interactive, and
- complex (e.g. hypermedia applications)


groupware.



How COAST applications are structured



The COAST Framework



point out functionality provided by COAST

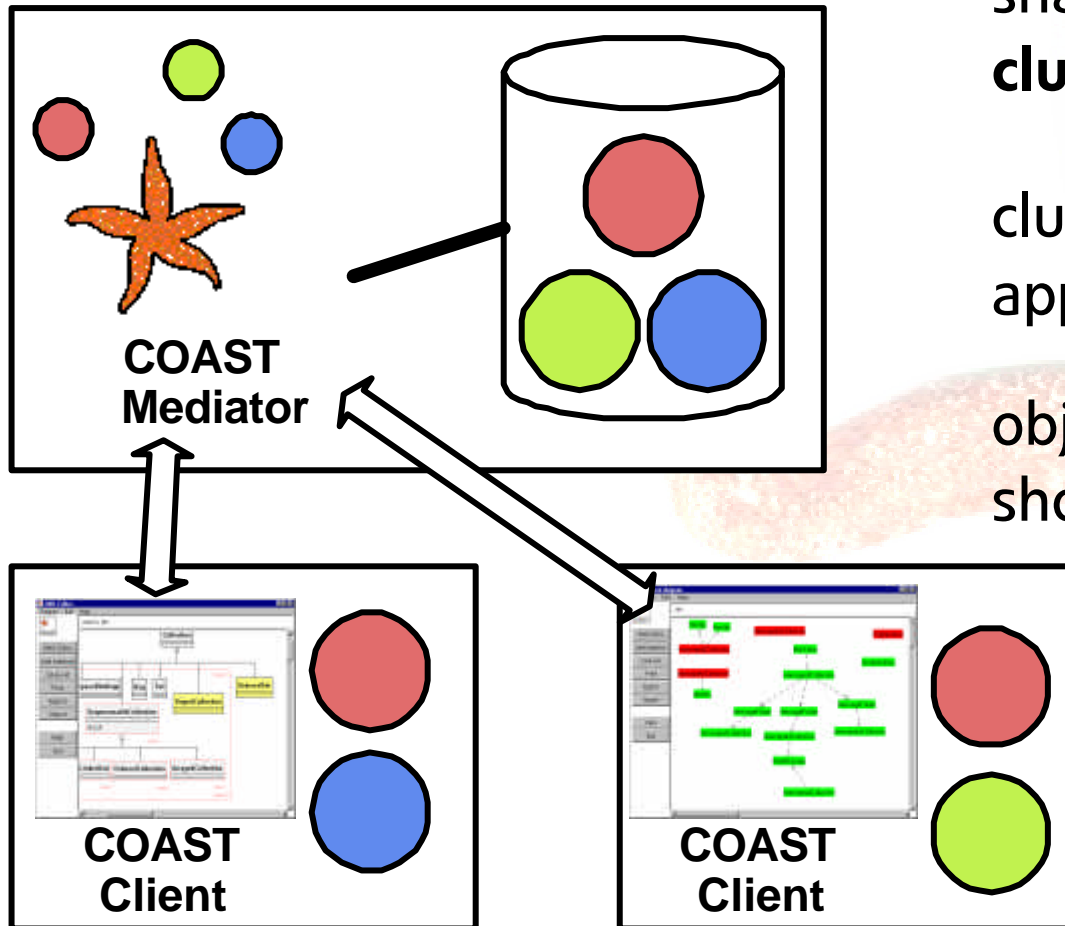
- replication & persistence
- concurrency control
- frame system
- view programming



explain how K-Infinity makes use of the functionality
demonstrate the functionality



Persistency and Replication



shared objects are bundled in **clusters**

clustering is critical for application performance

objects that are used together should be in the same cluster

using the COAST features
clustering strategy

- COAST default behaviour led to bad results
- open issue: re-clustering of COAST volumes

persistency is achieved

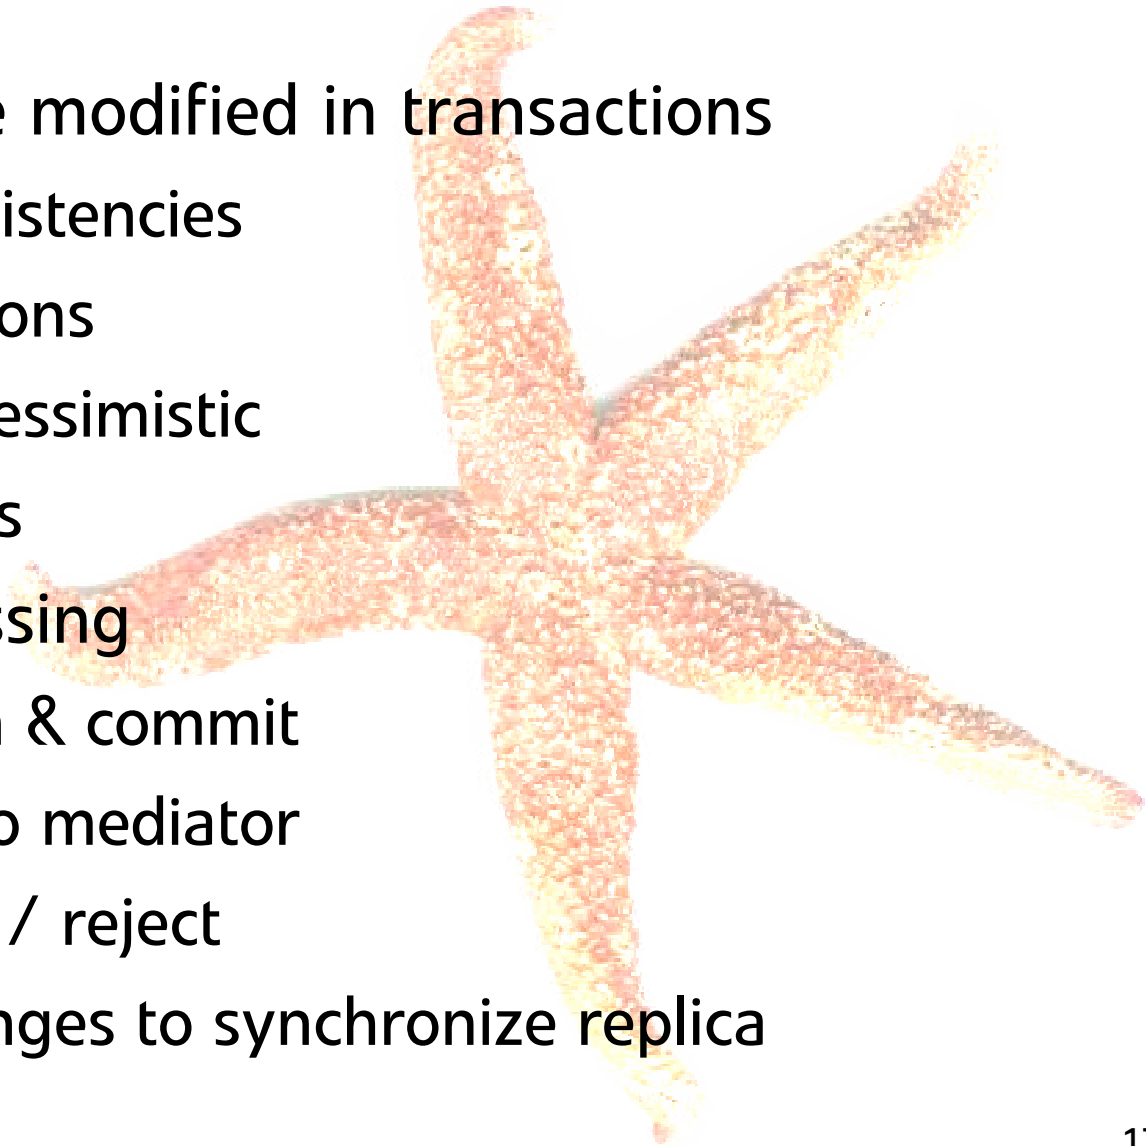
- implicitly by inheritance (COAST standard)
- explicitly by a method call (added for K-Infinity)

shared objects are modified in transactions

- prevent inconsistencies
- short transactions
- optimistic or pessimistic
- ACID properties

transaction processing

- local execution & commit
- send agenda to mediator
- global commit / reject
- broadcast changes to synchronize replica



experiences in K-Infinity

- use optimistic transactions wherever possible
- very little concurrency conflicts
 - separated responsibilities for each user
 - instant visibility of other user's actions

standard approach: locking

- single operations can influence the whole net
- generous locking in extreme leads to single user applications

DEMO: Concurrency Control



K-Infinity

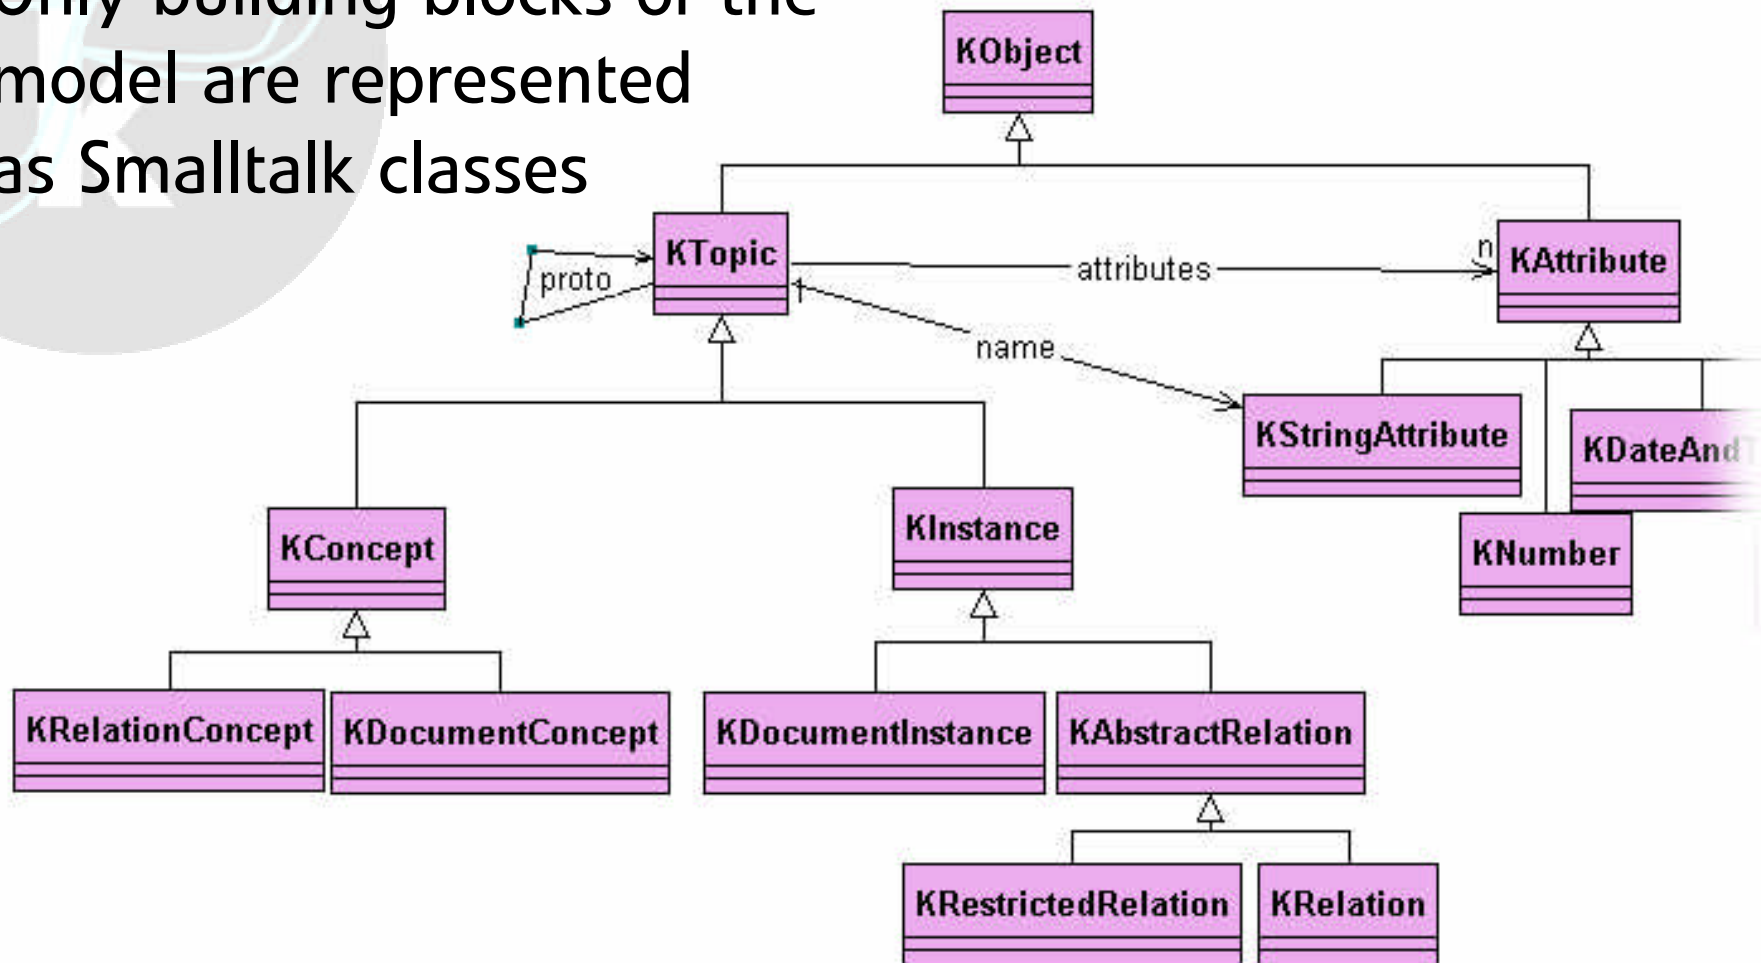
trigger a concurrency conflict

shared objects are modelled as frames and slots
slots are similar to instance variables

- slot properties are declared explicitly
- slots can hold multiple values
- demons
 - inverse relationships
 - type checking
 - constituent slots
 - user defined demons

The K-Infinity domain model

- Prototype system built on top of COAST
- Only building blocks of the model are represented as Smalltalk classes



advantages

- very flexible
- no code changes for adaption to specific knowledge domain

disatvantages

- stuff had to be re-done that is already present in ST (e.g. inheritance – but it is multiple now ;-)
- not as fast as „native“ ST classes would be

automatic view updating

- computed attributes
 - cache computation results
 - computation on demand or on invalidation
 - automatic invalidation
 - dependencies between model and computed attributes are detected by the framework
- views have computed attributes that trigger redisplay
- integrated into transaction scheme


COAST mechanisms work well

- display keeps consistent
- no explicit update notification code

complex visualization

- many dependencies
- wide parts are custom build due to requests of our design departement

hard work where COAST and standard mechanisms meet each other



DEMO: UI behavior in K-Infinity



K-Infinity

change name of a topic & point out the change of all visualizations

change 'allows instances' flag and point out different visualization in the graph editor

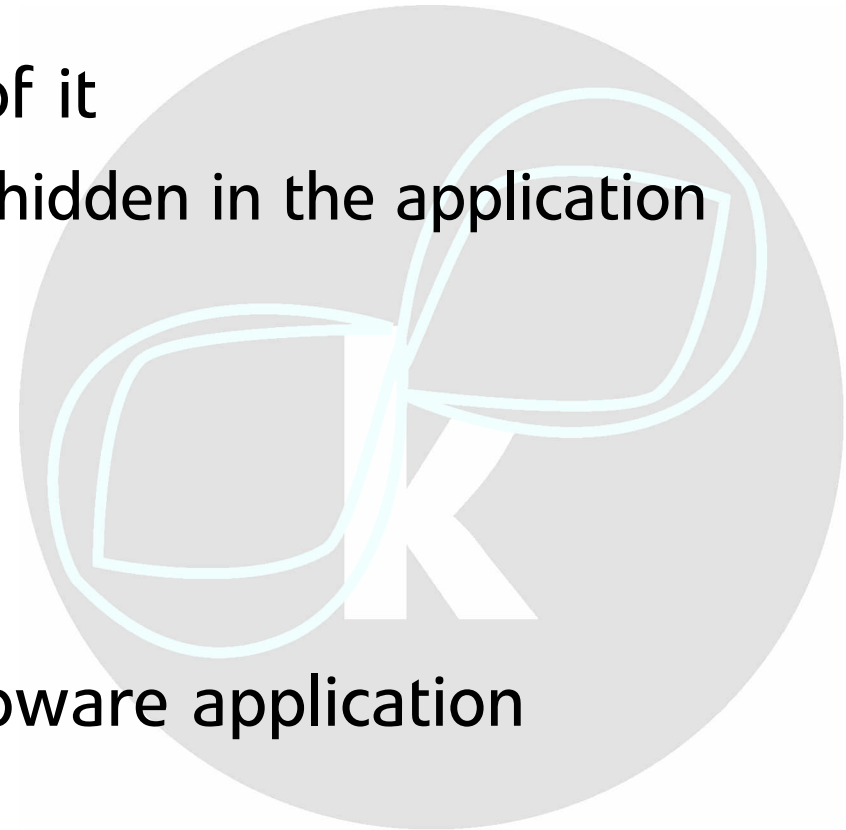
And where is the Groupware?

K-Infinity shows only little of it

- groupware functionality hidden in the application
 - shared workspace
 - user list
 - chat
 - awareness

K-Infinity is used as a groupware application

- distributed users
- demand for co-ordination
- clients are asking for workflows



what did we need to improve

- performance
 - optimizations mainly on the replication layer
- scalability
 - more data
 - factor 15 to TUKAN
 - factor 200 to our prototypes from research
 - more clients
- quality
 - found lots of bugs
 - mediator has to run 24/7

further information

www.i-views.de

www.opencoast.org

