Introduction to the ToolBus Coordination Architecture

Paul Klint







Introduction to the ToolBus Coordination Architecture

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The problem: component interconnection

- Systems become heterogeneous because we want to couple existing and new software components
 - different implementation languages
 - different implementation platforms
 - different user-interfaces
- Systems become distributed in local area networks
- Needed: interoperability of heterogeneous systems

Component interconnection: reasons

- Reusing existing components decreases construction costs of new systems
- Decomposing large, monolithic systems into smaller, cooperating components increases
 - modularity
 - flexibility



Component interconnection: issues

- Data integration: exchange of data between components
- Control integration: flow of control between components
- User-interface integration: how do the userinterfaces of components cooperate?



Data integration

- Data representations differ per
 - machine: word size, byte order, floating point representation, ...
 - language implementation: size of integers, emulation of IEEE floating point standard, ...
- How can we exchange data between components:
 - integers, reals, record => linear encoding
 - pointers => impossible in general



Data integration

- Assume a common representation *R*
- For each component C_i (with data domain D_i) there exist conversion functions

$$-f_i: D_i \to R \text{ and } f_i^{-1}: R \to D_i$$

- Convert a value d_i from C_i to C_i by $f_i^{-1}(f_i(d_i))$
- Examples: IDDL, ASN-1, XML, ...
- ToolBus uses ATerms as common representation

Control integration

- Broadcasting: each component can notify other components of state changes
- Remote procedure calls: components can call each other as procedures
- General message passing: the most general approach
- In the ToolBus Tscripts are used to model the interactions between components



User-interface integration

• The ToolBus does not address user-interface integration as separate issue but can be used to achieve it



Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

Brief history of the ToolBus

- In 1992 the first implementation of the ASF+SDF Meta-Environment was completed:
 - 200 KLOC Lisp code
 - Monolithic
 - Hard to maintain
- ... all traits of a legacy system



Time line

- 1992: Unsuccessful decomposition experiments
- 1994: First generation: ToolBus
- 1995: Second generation: Discrete time ToolBus
- 2001: Meta-Environment based on ToolBus
- 2002/7: Extensions, new functions and structure
- 2007: Third generation: Java-based ToolBus



1992



Old





Introduction to the ToolBus Coordination Architecture

1993

- Difficult synchronization and communication problems problems start to appear
- PSF specification of communication; simulation reveals several deadlocks
- Problems with this specification:
 - complex (> 20 pages) and ad hoc
 - difficult to extend
 - cannot be used to directly coordinate the components



1993/1994

- Idea of a "ToolBus" as general communication structure appeared
- First design and implementation
- Several experiments
 - Feature interaction in telephone switches (RUU/PTT)
 - Traffic control (Nederland Haarlem/UvA/CWI/RUU)
 - Management of complex bus stations (idem)
 - Definition of user interfaces (UvA)



1994/1995

- Fall 1994: redesign based on this experience
- Spring 1995: design and implementation of Discrete Time ToolBus completed
- First experiments to prototype parts of the Meta-Environment started



More recently ...

- In 2001 a new implementation of the Meta-Environment based on the ToolBus was completed
- In 2007 we have completed a new generation ToolBus (Java-based) that is used by the Meta-Environment
- The ToolBus can be seen as a Service-oriented Architecture (SOA) *avant la lettre* ...



Structuring and Composition of Software

- Structured programming
- Functions, procedures & libraries
- Object-oriented programming & Modules
- Unix pipes
- DCOM
- Coordination languages & SOA



Introduction to the ToolBus Coordination Architecture

Modules

System

Procedures

Statements

Service-oriented Architecture (SOA)

- Loose coupling
- Service contract
- Autonomy
- Abstraction
- Reusability
- Composability
- Statelessness

The Meta-Environment

• Discoverability

- Message exchange patterns
- Coordination
- Atomic transactions

ToolBus requirements

- Flexible interconnection architecture for software components
- Good control over communication
- Relatively simple descriptions
- Uniform data exchange format
- Multi-lingual: C, Java, Perl, ASF+SDF, ...
- Potential for verification
- Use existing concurrency theory: Process Algebra

Process Algebra

- A theoretical framework to describe process behaviour
- Consists of
 - Constants: deadlock (δ), silent step (τ)
 - Atomic actions: a, b, c, ...
 - Processes x, y, z, ... built with the operators:
 - sequential compositions: .
 - non-deterministic choice: +
 - parallell composition: ||

Basic Process Algebra (BPA) The basic axioms for choice (+) and sequential composition (.):

A1.
$$x + y = y + x$$

A2. $(x + y) + z = x + (y + z)$
A3. $x + x = x$
A4. $(x + y) \cdot z = x \cdot z + y \cdot z$
A5. $(x \cdot y) = x \cdot y \cdot z$

Axioms for deadlock: A6. $x + \delta = x$ A7. $\delta \cdot x = \delta$

The Meta-Environment

Introduction to the ToolBus Coordination Architecture

Merge (||)

Use the auxiliary operator left merge (\parallel):

M1.
$$x || y = x ||_y + y ||_x$$

M2. $a ||_x = a \cdot x$
M3. $a \cdot x ||_y = a \cdot (x ||y)$
M4. $(x + y) ||_z = x ||_z + y ||_z$

Examples:

$$a \parallel b = a \parallel _b + b \parallel _a = a .b + b .a$$

 $a .b \parallel c = a .b \parallel _c + c \parallel _a .b =$
 $a . (b \parallel c) + c.a.b = a . (b.c + c.b) + c.a.b$

The Meta-Environment

Introduction to the ToolBus Coordination Architecture

Process Algebra versus ToolBus

- Process Algebra can be used to describe all (possibly infinite) behaviours of a collection of parallel processes
- This behaviour has the form of a process tree still containing all possible choices
- Properties of the parallel processes can be verified by verifying this behaviour description, e.g.
 - absence of deadlock

Process Algebra versus ToolBus

- Atomic actions may be enabled/disabled as a result of conditions or time constraints
- The ToolBus executes a process expression but randomly selects one of the enabled arguments of a choice operator
- The steps taking by the ToolBus are thus just one possible series of steps that is contained in the complete behaviour of the process expression:
 - a || b executes as a . b or as b .a (and not both!)



Coordination, Representation & Computation

- Coordination: the way in which program and system parts interact (procedure calls, RMI, ...)
- Representation: language and machine neutral data exchanged between components
- Computation: program code that carries out a specialized task

A rigorous separation of coordination from computation is the key to flexible and reusable systems



Architectural Layers



Cooperating Components



Introduction to the ToolBus Coordination Architecture

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

Why not using XML as terms?

- Has been tried in various language processing projects
- XML is too verbose to represent parse trees of large (> 100 KLOC) programs
- XML does not provided sharing
- For discussion see: M.G.J. van Brand and P. Klint, ATerms for manipulation and exchange of structured data: It's all about sharing, *Information and Software Technology*, **49**(1), 2007, 55-64.



Generic Representation Annotated Terms (ATerms)

• Applicative, prefix terms

- Maximal subterm sharing $(\Rightarrow DAG)$
 - cheap equality test, efficient rewriting
 - automatic generational garbage collection
- Annotations (text coordinates, dataflow info, ...)
- Very concise, binary, sharing preserving encoding
- Language & machine independent exchange format





A term is ...

• a Boolean, integer, real or string

- true, 37, 3.14e-12, "rose"

- a value occurrence of a variable
 - X, InitialAmount, Highest-bid
- a result occurrence of a variable
 - X?, InitialAmount?



A term is ...

- a single identifier
 - f, pair, zero
- a function application
 - pair("rose", address("Street", 12345))
- a list
 - [a, b, c], [a, 1.25, "last"], [[a, 1], [b, 2]]
- a placeholder

- <int>, add(<int>,<int>)

The Meta-Environment

Introduction to the ToolBus Coordination Architecture
Matching of terms

- Term matching is used to
 - determine which actions can communicate
 - to transfer data between sender and receiver
- Intuition:
 - terms match if the are structurally identical
 - value occurrence: use variable's value
 - result occurrence: assign matched subterm to variable (only if overall match succeeds!)



Example of term matching



Types

- The ToolBus uses its own type system
 - static checks & automatic generation of interface code
- bool, int, real, str
- list: list with arbitrary elements
- list(*Type*): list with *Type* elements
 list(int)
- *term*: arbitrary term

Types

- *Id*: all terms with function symbol *Id* (allows partial type declarations)
 - f accepts f, f(1), f("abc",3), ...

•
$$Id(T_1, ..., T_n)$$

- f(int, str) accepts f(3,"abc") but not f(3)
- $[T_1, ..., T_n]$: list of elements with given types

- [int, str] accepts [1,"abc"] but not [1,2,3]

Types

- All variables have types
- Types are checked statically when possible
- Types play a role during matching:
 - I is int variable, S is str variable, T is term variable
 - match f(13) and f(I?)
- succeeds
- match f(13) and f(S?)
- match f(13) and f(T?)

fails

succeeds



Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

The ToolBus architecture





The ToolBus architecture

- Processes inside the ToolBus can communicate with each other
- Tools can not communicate with each other
- Tools can communicate using a fixed protocol:









Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

ToolBus scripts: processes

- The ToolBus: a parallel composition of processes
- Private variables per process
- $P_1 + P_2$ $P_1 \cdot P_2$ $P_1 || P_2 P_1 * P_2$
- :=, if then else
- All data are terms that can be matched
- A limited set of built-in operations on terms
- No other support for datatypes

ToolBus scripts: processes

- Send, receive message (handshaking)
- Send/receive notes (broadcasting)
- Subscription to notes
- Dynamic process creation
- Absolute/relative delay, timeout



ToolBus scripts: tools

- Execute/terminate tools
- Connect/disconnect tools
- Communication between process and tool is synchronous
- Process can send evaluation request to tool (which returns a value later on)
- Tool can generate events to be handled by the ToolBus



Hello World





Hello World: string generated by tool



Simple clock with user-interface





Simple clock with user-interface

process CLOCK is process-expression-1 tool clock is tool-definition-1

process UI is process-expression-2 tool ui is tool-definition-2

toolbus(CLOCK, UI)





User-interface



Tscripts: in more detail

- Process communication: messages & notes
- Composite processes
- Expressions & built-in functions
- Time primitives
- Tools



Process communication: messages

- Messages used for synchronous, two-party communication between processes
- snd-msg and rec-msg synchronize sender/receiver
- Communication is possible if the arguments match
- There is two-way data transfer between sender and receiver (using result variables)



Process communication: notes

- Notes used for asynchronous, broadcasting communication between processes
- Each process must subscribe to the notes it wants to receive
- Each process has a private note queue on which snd-note, rec-note and no-note operate



Process communication: notes

- subscribe to notes of a given form
 - subscribe(compute(<str>,<int>))
- unsubscribe from certain notes
- snd-note to all subscribers
 - snd-note(compute(E,V))
- rec-note: receive a note of a given form
- no-note received of given form

Composite process expressions

- One of the atomic processes mentioned above
- delta (deadlock), tau (silent step)
- $P_1 + P_2$: choice (non-deterministic)
- P₁. P₂: sequential composition
- $P_1 || P_2$: parallel composition
- $P_1 * P_2$: repetition

Composite process expressions

- $P(T_1, T_2, ...)$: a named process (with optional parameters) will be replaced by its definition
- create(P(T₁, T₂, ...), Pid?): dynamic process
 creation
- V := Expr: evaluate Expr and assign result to V
- if Expr then P_1 else P_2 fi
- if Expr then P_1 fi = if Expr then P_1 else delta fi

Expressions

- An expression is evaluated in the current environment of the process in which it occurs
- Constants evaluate to themselves: a
- Variables evaluate to their current values
- Lists evaluate to a list of their evaluated elements
- Some function symbols have a built-in meaning



Built-in functions

- Booleans: not, and, or
- Integers: add, sub, mul, mod, less, less-equal, greater, greater-equal
- Lists: first, next, get, put, join, member, subset, diff, inter, size
- Miscellaneous: equal, not-equal, process-id, process-name, current-time, quote



Time primitives

- A (relative or absolute) delay or time out may be associated with each atomic process
- Relative time: delay(Expr) or timeout(Expr)
- Absolute time: abs-delay(y, mon, d, h, min, s) or abs-timeout(y, mon, d, h, min, s)
- Example:
 - printf("expired") delay(10)
 - printf("Renew account") abs-timeout(2008,4,1,12,0,0)

Process definitions

- Process definition: process *Pname Formals* is *P*
- *Formals* are optional and contain a list of formal parameter names

- process MakeWave(N: int) is ...

- All variables (including formals) must be declared and have a type
- let *VarDecls* in *P* endlet introduces variables:

```
- let E : str, V : int in ... endlet
```

Tools

- Tools have to be executed or connected before they can be used
- Requires a tool definition: tool ui is { ... }
- Introduces a new type, e.g. ui

The Meta-Environment

- Execute a tool: execute(ui, Uid?)
- Receive connection request: rec-connect(ui, Uid?)
- Tool identification is assigned to Uid (of type uid)

Tools

- snd-terminate: terminate an executing tool
 - snd-terminate(Tid)
- rec-disconnect: receive disconnection request from tool
 - rec-disconnect(Uid)
- shutdown: terminate the whole ToolBus
 - shutdown("Auction ends")



Tools

• snd-eval, rec-value: request tool to evaluate a term, and receive the resulting value from tool

- initiative: ToolBus

- snd-do: request tool to perform some action, there is no reply
 - initiative: ToolBus
- rec-event, snd-ack-event: receive event from tool, acknowledge it after appropriate processing
 - initiative: tool

The Meta-Environment

Tscripts

- A Tscripts consists of
 - a list of process and tool definitions
 - a single ToolBus configuration
- A ToolBus configuration describes the initial set of active processes in the ToolBus:
 - toolbus($Pname_1, ..., Pname_n$)
 - Eache *Pname* is optionally followed by parameters



Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment
Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples: calculator; auction; waves
- Implementation issues
- Conclusions

The Meta-Environment

Example: calculator





Example: calculator

- CALC: the calculation process
- BATCH: reads expressions from file, calculates their value, writes result back to file
- UI: the user-interface
- LOG maintains a log of all calculations
- CLOCK provides current time



Process CALC





Process BATCH





User-interface



- When the user presses <u>Calc</u>, a dialog window aetExpr appears to enter an expression plus(3,4)
- The result is diplayed in a separate window
- Pressing showLog display all calculations so far
- Pressing <u>showTime</u> displays the current time
- Pressing <u>Quit</u> ends the application



Give expression:

Cancel

Value is: 7

ok.

User-interface: process UI





User-interface: CALC-BUTTON



Introduction to the ToolBus Coordination Architecture

User-interface: LOG-BUTTON

```
process LOG-BUTTON(Tid : ui) is
  let N : int, L : term
  in
     rec-event(Tid, N?, button(showLog)).
     snd-msg(showLog).
     rec-msg(showLog, L?).
     snd-do(Tid, display-log(L)).
     snd-ack-event(Tid, N)
  endlet
```



User-interface: TIME-BUTTON

```
process TIME-BUTTON(Tid : ui) is
  let N : int, T : str
  in rec-event(Tid, N?, button(showTime)).
     snd-msg(showTime).
     rec-msg(showTime, T?).
     snd-do(Tid, display-time(T)).
     snd-ack-event(Tid, N)
  endlet
```

process QUIT-BUTTON(Tid : ui) is
 rec-event(Tid, button(quit)) .
 shutdown("End of calc demo")

Process LOG



Process LOG1



Process CLOCK

```
process CLOCK is
  let Tid : clock, T : str
  in
       execute(clock, Tid?).
            rec-msg(showTime).
        (
            snd-eval(Tid, readTime) .
             rec-value(Tid, time(T?)).
            snd-msg(showTime, T)
         ) * delta
  endlet
```



ToolBus Configuration

toolbus (CALC, BATCH, UI, LOG, CLOCK)

Creates the processes for the calculator application

Start calculator application: toolbus calc.tb



Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples: calculator; auction; waves
- Implementation issues
- Conclusions

The Meta-Environment

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples: calculator; auction; waves
- Implementation issues
- Conclusions

The Meta-Environment





Introduction to the ToolBus Coordination Architecture

- How are bids synchronized?
- How to inform bidders about higher bids?
- How to decide when the bidding is over and the item is sold?
- Bidders may come and go during the auction







- The Auction process
 - executes master tool: user-interface of auction master
 - connection/disconnection of new bidders
 - introduces new items for sale (at the initiative of the auction master)
 - controls the bidding provess via OneSale
- A Bidder process is created for each new bidder



Process Auction



tool master is { command = "wish-adapter -script master.tcl" }

Introduction to the ToolBus Coordination Architecture

Process ConnectBidder





process OneSale(Mid : master) is						
	let Descr : str,	%% Description of current item for sale				
	InAmount : int,	%% Initial amount for item				
	Amount : int,	%% Current amount				
	HighestBid : int,	%% Highest bid so far				
	Final : bool,	%% Did we already issue a final call for bids?				
	Sold : bool,	%% Is the item sold?				
	Bid : bidder	%% New bidder tool connected during sale				
	in rec-event(Mid, new-item(Descr?, InAmount?)) .					
	HighestBid := InAmount .					
	<pre>snd-note(new-item(Descr, InAmount)) .</pre>					
	Final := false . Sold :	= false. Where the action is				
	•					
) * if Sold then snd-ack-event(Mid, new-item(Descr, InAmount)) fi						
endlet						
The Meta-Environment Introduction to the ToolBus Coordination Architecture 96						

(if not(Sold) then ... fi + if not(or(Final, Sold)) then ... fi + if and(Final, not(Sold)) then ... fi + ConnectBidder(Mid, Bid?) ...) * if Sold then ... fi





(if not(Sold) then fi	Not yet sold, not asked for final bids			
+ if not(or(Final, Sold)) then		Wait	10 sec, then ask for final bids	
snd-note(any-higher-bid) delay(sec(10) snd-do(Mid, any-higher-bid(10)) .)).	Inform auction master		
Final := true •	Yes,	Yes, now we have asked for final b		
if and(Final, not(Sold)) then •	Not yet sold, but asked for final bids			
snd-note(sold(HighestBid)) delay(sec(1 Sold := true fi Yes, item is now s	10)) .• sold		Wait 10 sec, then inform all bidders that item is sold	
+ ConnectBidder(Mid, Bid?) . • snd-msg(Bid, new-item(Descr, HighestBi	id)).	Bidder is connected during sale		
Final := false) * if Sold then fi	[Inform new bidder about progress		
			Restart, final bids (if any)	



Process Bidder

```
process Bidder(Bid : bidder) is
let Descr : str, Amount : int, Acceptance : term
in
    subscribe(new-item(<str>, <int>)) . subscribe(update-bid(<int>))
    subscribe(sold(<int>)) . subscribe(any-higher-bid) .
    ( ...
    )
    * delta
endlet
```



Get info about item for sale after connection

Process Bidder



Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples: calculator; auction; waves
- Implementation issues
- Conclusions

The Meta-Environment

Road map

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples: calculator; auction; waves
- Implementation issues
- Conclusions

The Meta-Environment

One-dimensional wave equation

Simulate a string attached at the two end points:







Introduction to the ToolBus Coordination Architecture

One-dimensional wave equation

Amplitude at point *i* at $t+\Delta t$ is given by:

$$y_i(t+\Delta t) = F(y_i(t), y_i(t-\Delta t), y_{i-1}(t), y_{i+1}(t))$$

and

$$F(z_1, z_2, z_3, z_4) = 2z_1 - z_2 + (c \Delta t / \Delta x)^2 (z_3 - 2z_1 + z_4)$$

 Δx : the (small) interval between sampling points

c: constant representing the propagation velocity of the wave



Example: wave equation





One-dimensional wave equation

- Auxiliary process F computes function F
- Process P models a sampling point
- Process Pend models the end points
- Process MakeWave constructs *N* connected instances of P and two end points
- Tool display visualizes the simulation



Process F

Compute $F(z_1, z_2, z_3, z_4) = 2z_1 - z_2 + (c \Delta t / \Delta x)^2 (z_2 - 2z_1 + z_4)$





Introduction to the ToolBus Coordination Architecture
Process P

```
process P(Tid : display, L : int, I : int, R : int, Dstart : real, Estart : real) is
 let AL : real, AR : real, D : real, D1 : real, E : real
 in
                                                    L: left, I: this point, R: right
   D := Dstart . E := Estart . •
   (( rec-msg(L, I, AL?)
                                                   D, E: amplitutes of this point
    || rec-msg(R, I, AR?)
    || snd-msg(I, L, E)
                                              Receive amplitudes of neighbours
    || snd-msq(I, R, E)
                                              Send our amplitude to neighbours
    || snd-do(Tid, update(I, E))
                                               Update our amplitude on display
    D1 := E .
    F(E, D, AL, AR, E?) ...
                                             Compute new versions of D and E
    D := D1
   ) * delta
 endlet
```







Process MakeWave



The Meta-Environment

Introduction to the ToolBus Coordination Architecture

Tool definition and ToolBus configuration

tool display is { command = "wish-adapter -script ui-wave.tcl"}

toolbus(MakeWave(8))



- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples: calculator; auction; waves
- Implementation issues
- Conclusions

The Meta-Environment

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

Requirements ATerms

- Open: independent of hw/sw platform
- Simple: a small API
- Efficient: fast reading and writing
- Concise: small memory usage
- Language-independent
- Annotations: applications can transparantly store additional information in data structure



ATerm Types

- INT
- REAL
- APPL
- LIST
- PLACEHOLDER
- BLOB (Binary Large OBject)
- ANNOTATION



Examples

- 1 3.14 -0.7E34
- f(a,b) "test!"(1, 2.1, "hello")
- [] [1, 2, "abc"]
- <int> f(<int>, <real>)
- BLOBs
 - used to encode images, binary files, ...
 - have no textual representation



The ATerm Implementation

- C and Java API
- Only applicative operations
 - No destructive operations on ATerms
- Maximal subterm sharing
- Automatic garbage collection
- Binary encoding (BAF: Binary ATerm Format)



The ATerm C API

- Level 1: 41 functions
- Level 2: 80 functions (superset of Level 1)
- All function start with AT
- Defines types ATerm and ATbool
- Make and Match
- Read and Write
- Annotate



Intermezzo: Patterns

- A pattern is an ATerm with placeholders: incr(<int>)
- A string pattern is a pattern represented as string: "incr(<int>)"
- A string pattern resembles the format string in printf/scanf in C
- Placeholders correspond to typed arguments of ATmake/ATmatch



Make and Match

• ATerm ATmake(String p, ATerm a1, ...)

- parse **p** and fill placeholders with **a1**, **a2**, ...

- ATerm ATmatch(ATerm t, String p, ATerm *a1, ...)
 - match † against p; assign subterms at placeholders to a1, a2,...
- ATbool ATisEqual(ATerm t1, ATerm t2)
- int ATgetType(ATerm t)

Read and Write

- ATerm ATreadFromString(String s)
- ATerm ATreadFromTextFile(File f)
- ATerm ATreadFromBinaryFile(File f)
- ATbool ATwriteToTextFile(ATerm t, File f)
- ATbool ATwriteToBinaryFile(ATerm t, File f)
- char *ATwriteToString(ATerm t)



Annotate

- ATerm ATsetAnnotation(ATerm t, ATerm I, ATerm a)
 - add annotation [I, a] to copy of \dagger
- ATerm ATgetAnnotation(ATerm t, ATerm I)
- ATerm ATremoveAnnotation(ATerm t, ATerm I)



Other Functions in the Level 1 API

- Variations on the preceeding functions
- ATprintf
- handlers (warnings and errors)
- protect/unprotect



Structure of an ATerm-based Application



Introduction to the ToolBus Coordination Architecture

The Level 2 API

- Detailed operations for efficient ATerm manipulation
- Dictionaries
- Tables
- Indexed sets



The Java API

- Two versions:
 - Native (uses the C version via JNI, not implemented)
 - Pure (a pure Java reimplementation)
- Interface ATermFactory encapsulates the whole API
- Separate interfaces for each kind of ATerm (AFun, ATermList, etc.)



Class Structure



Using ATermFactory

```
import aterm.*
factory = new PureFactory();
ATerm t1 = factory.makeInt(3)
ATerm t2 = factory.readFromFile("test.trm");
ATerm t3 = factory.makeAFun("f1", 1, false);
ATerm t4 = factory.make("f(<int>)", 3);
ATerm t5 = factory.parse("f(1, [a, b])");
```

- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

ToolBus design/implementation method

- Specification of ToolBus using ASF+SDF
- Execution of small test cases
 - Tool behaviour is defined very abstractly
- Hand translation of ASF+SDF specification to C
 - Literal translation of Tscripts
 - Implementation of tools is more concrete (see later)
- Very few bugs in ToolBus implementation
 - Some bugs turned out to be bugs in the specification!



The ToolBus implementation





The ToolBus Interpreter

- Syntax analysis of Tscript (lex/yacc)
- Typechecking of Tscript
- Create the initial ToolBus configuration
- Start execution
- Delays and timeouts
- Garbage collection of terms



The ToolBus Interpreter

- Execute tools as separate Unix process
- On creation: send expected input signature to tool
 - Permits detection of Tscript/tool mismatches
- During execution of tool: check terms received from tool against their output signature
 - Permits detection of misbehaving tools
- Enforce ToolBus protocol for each tool



Main Interpreter Loop

- Wait for
 - an event coming from one of the tools
 - expiration of a timer
- Compute effect of event/timer on ToolBus state
- Perform any enabled atomic actions
- Repeat as long as possible
- Go back to waiting state



ToolBus Interpreter

- Interpreter maintains a lists of processes
- Each process is compiled into a finite automaton with an action associated with each transition
 - From the enabled actions one is selected randomly and executed
 - The process goes to corresponding next state
- A select system call waits for i/o on any socket or expiration of timer





Introduction to the ToolBus Coordination Architecture

Implementation considerations

- Terms are linearized before sending and parsed when receiving them
- There is a separate transport layer that provides byte level messages of given length (to avoid system dependent segmentation of the byte stream)



- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

Recall the Hello World script

process HELLO is printf("Hello world, my first Tscript!\n")

toolbus(HELLO)



Introduction to the ToolBus Coordination Architecture
Hello World: string generated by tool

```
process HELLO is
```

let H : hello,

```
S:str
```

```
in
```

```
execute(hello, H?) .
snd-eval(H, get-text) .
rec-value(H, text(S?)) .
printf(S)
endlet
```

tool hello is {command = "hello" }
toolbus(HELLO)

How can we implement this tool?

```
The Meta-Environment
```

Introduction to the ToolBus Coordination Architecture

First version of a hello tool (C)



hello_handler

```
get-text
ATerm hello_handler(int conn, ATerm inp)
{ ATerm arg, isig, osig;
 if(ATmatch(inp, "rec-eval(get-text)"))
  return ATmake("snd-value(text(\"Hello World, my first ToolBus tool in C!\n\"))");
 if(ATmatch(inp, "rec-terminate(<term>)", &arg))
                                                                  terminate
  exit(0);
 if(ATmatch(inp, "rec-do(signature(<term>,<term>))", &isig, &osig)){
  return NULL:
                                                          receive input signature
 ATerror("hello: wrong input %t received \n", inp);
 return NULL:
```



Observations

- Tool consists of main and an event handler
- All processing is routed via the event handler
- Event handler does repetitive (and error prone) decoding of requests using ATmatch
- Event handler takes care of standard messages for termination, signature handling, etc.
- Why not automate some of these tasks?







Introduction to the ToolBus Coordination Architecture





Generated file hello.tif.c





Generated file hello.tif.c





Generated file hello.tif.c

```
ATerm hello_handler(int conn, ATerm term)
{ ...
 if(ATmatch(term, "rec-do(signature(<term>,<term>))", &in, &out)) {
  ATerm result = hello_checker(conn, in);
  if(!ATmatch(result, "[]"))
   ATfprintf(stderr, "warning: not in input signature:\n\t%t\n\tl\n", result);
  return NULL:
 ATerror("tool hello cannot handle term %t", term);
 return NULL; /* Silence the compiler */
```



```
process CALC is
  let Tid : calc, E : str, V : term
  in
     execute(calc, Tid?).
         rec-msg(compute, E?).
         snd-eval(Tid, expr(E)) . rec-value(Tid, val(V?)) .
         snd-msg(compute, E, V) . snd-note(compute(E, V))
       )* delta
  endlet
tool calc is { command = "calc"}
```











- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

ToolBus Adapters

• Needed to adjust existing programs/libraries to the ToolBus

Library:

Separate program:



A selection of adapters

- wish-adapter: execute Tcl/Tk windowing shell
- tcltk-adapter: ditto but uses the Tcl/Tk library[†]
- java-adapter: java program as tool
- perl-adapter: perl program as $tool^{\dagger}$
- python-adapter: python program as $tool^{\dagger}$
- gen-adapter: arbitrary Unix command as tool[†]

† = not yet supported in Java-based ToolBus

The wish-adapter

- Execute Tcl/Tk's windowing shell as a tool
- Ex. wish-adapter -script calculator.tcl
 - -script: The Tcl script to be executed
 - -script-args: Arguments for the Tcl script
- The command wish is executed once and all further requests are directed to this instance of wish



The wish-adapter

- snd-eval(*Tid*, *Fun*($A_1, ..., A_n$)): perform the Tcl function call *Fun* $A_1 ... A_n$
- rec-value(*Tid*, *Res*?): return value for previous eval request
- rec-event(*Tid*, $A_1, ..., A_n$): event generated by wish
- snd-ack-event(*Tid*, A₁): ack previous event
- snd-terminate(*Tid*, A₁): terminate wish-adapter

The gen-adapter

- Execute arbitrary Unix command as tool
- Example: gen-adapter -cmd ls -l
- snd-eval(*Tid*, cmd(*Cmd*, input(*Str*): execute the Unix command *Cmd* with *Str* as standard input
- rec-value(*Tid*, output(*Res*?)): receive the standard output *Res* from a previous command
- snd-terminate(*Tid*, *Arg*): terminate execution of gen-adapter

- The problem: component interconnection
- Implementation issues
 - Overview of the ATerm API
 - Brief overview of the ToolBus implementation
 - Writing ToolBus tools
 - Using ToolBus adapters
- Conclusions

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

- The problem: component interconnection
- History & requirements
- Terms, types & matching
- The ToolBus architecture
- ToolBus scripts (Tscripts)
- Larger examples
- Implementation issues
- Conclusions

The Meta-Environment

Conclusions

- ToolBus is an effective technology for coordination and composition of tools
- ToolBus fits in the popular model of serviceoriented architectures
- ToolBus enables incremental software renovation



A Legacy System

- A complete blackbox:
- Subsystems unknown
- Subsystem depencies unknown





Analyze and decompose in major subsystems





Introduction to the ToolBus Coordination Architecture





Separate renovation strategy per subsystem



The Renovation Process





Further reading

- See at http://www.meta-environment.org (Documentation menu entry):
 - Guide to ToolBus Programming
 - The ATerm Programming Guide
 - Further references can be found in *Bibliography*

