Chapter 1. Guide to ToolBus Programming

Hayco de Jong Paul Klint Arnold Lankamp Pieter Olivier

2007-12-04 15:02:36 +0100 (Tue, 04 Dec 2007)

Table of Contents

2
2
3
4
5
6
6
6
7
8
14
14
20
28
33
33
33
34
35
35
36
37
38
38
39
39
42
43
47
48
48
49
50
50
51
51
51
51

Built-in functions	53
Synopsis of ToolBus primitives	58
Historical notes	60
Bibliography	60
То Do	61

Warning

This document is in state of creation and will further evolve. It provides a description of the classic C-based ToolBus as well as of the Java-based ToolBusNG. Eventually, it will exclusively focus on ToolBusNG. See the section called "*To Do*" (p. 61).

Introduction

Background and Motivation

Building large, heterogeneous, distributed software systems poses serious problems for the software engineer. Systems grow *larger* because the complexity of the tasks we want to automate increases. They become *heterogeneous* because large systems may be constructed by re-using existing software as components. It is more than likely that these components have been developed using different implementation languages and run on different hardware platforms. Systems become *distributed* because they have to operate in the context of local area networks.

Three aspects of heterogeneous, distributed, systems should be considered: *coordination*, *representation* and *computation*.

Coordination. Coordination is the way in which program and system parts interact with each other using, ordinary procedure calls, remote procedure calls (RPC), remote method invocation (RMI), and others.

Representation. Representation is the language and machine neutral format for data being exchanged between components.

Computation. Computation is done by specialized program code that carries out a specific task, e.g., providing a user-interface, providing database access, and the like.

Our key assumption is as follows:

Important

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

A system organization that respects this separation is shown in



Figure 1.1. Separating coordination from computation

We propose to get control over the possible interactions between software components (*tools*) by forbidding direct inter-tool communication. Instead, all interactions are controlled by a processoriented *script* that formalizes all the desired interactions among tools. This leads to a component interconnection architecture resembling a hardware communication bus, and therefore we call it a ``ToolBus".

ToolBus requirements

Given the motivation for the ToolBus we can briefly summarize the requirements that the ToolBus should satisfy:

- Provide a flexible interconnection architecture for software components that are not only written in different languages and executing on different hardware and software platforms, but are also running in a distributed fashion on a system of networked computers and devices. **Rationale**: it is more and more common that applications are built using existing commercial or open source components. This reduces implementation effort but increases the need for usable interconnection technology.
- Provide good control over the communication between components. **Rationale**: component integration requires both control over the communication between components and over the data that are being exchanged (see next requirement).
- Provide a uniform data exchange mechanism between heterogeneous components. **Rationale**: a common understanding about data formats is needed in order to exchange data between components.
- The description of communication should be based on existing concurrency theory and provide the option for formal verification of the cooperation between software components. **Rationale**: when components are integrated that run on different machines or on multi-core machines, it is unavoidable that concurrency is taken into account and to use existing theory to describe it. The long term perspective of checking formal aspects of these cooperation is appealing for certain, safety-critical, applications.
- Provide relatively simple application descriptions that can be understood by most programmers. **Rationale**: we don't want to frighten programmers by using formal notations.
- Provide multi-lingual support, at least C, Java, ASF+SDF, Tcl/Tk, and possibly Perl, Python and Ruby should be supported. Rationale: various tools of interest are currently implemented in the first four languages, and the last three languages are interesting for future developments.

The ToolBus architecture

The global architecture of the ToolBus is shown in Figure 1.2, "Global organization of the ToolBus" (p. 4). The ToolBus serves the purpose of defining the cooperation of a variable number of *tools* T_i (i = 1, ..., m) that are to be combined into a complete system. The internal behaviour or implementation of each tool is irrelevant: they may be implemented in different programming languages, be generated from specifications, etc. Tools may, or may not, maintain their own internal state. Here we concentrate on the external behaviour of each tool. In general an *adapter* will be needed for each tool to adapt it to the common data representation and message protocols imposed by the ToolBus.





The ToolBus itself consists of a variable number of processes P_i (i = 1, ..., n)¹The parallel composition of the processes P_i represents the intended behaviour of the whole system. Tools are external, computational activities, most likely corresponding with operating system level tasks. They come into existence either by an execution command issued by the ToolBus or their execution is initiated externally, in which case an explicit connect command has to be performed by the ToolBus. Although a one-to-one correspondence between tools and processes seems simple and desirable, we do not enforce this and permit tools that are being controlled by more than one process as well as clusters of tools being controlled by a single process.

Communication inside the ToolBus. Inside the ToolBus, there are two communication mechanisms available. First, a process can send a *message* (using snd-msg) which should be received, synchronously, by one other process (using rec-msg). Messages are intended to request a service from another process. When the receiving process has completed the desired service it may inform the sender, synchronously, by means of another message (using snd-msg). The original sender can receive the reply using rec-msg. By convention, part of the original message is contained in the reply (but this is not enforced).

Second, a process can send a *note* (using snd-note) which is broadcasted to other, interested, processes. The sending process does not expect an answer while the receiving processes read notes asynchronously (using rec-note). Notes are intended to notify others of state changes in the sending process. Sending notes amounts to *asynchronous selective broadcasting*. Processes will only receive notes to which they have *subscribed*.

¹By ``processes" we mean here computational activities *inside* the ToolBus as opposed to, for instance, processes at the operating system level. When confusing might arise, we will call the former ``ToolBus processes" and the latter ``operating system level tasks".



Figure 1.3. Communication between ToolBus and tools

Communication between ToolBus and tools. The communication between ToolBus and tools is based on handshaking communication between a ToolBus process and a tool. A process may send messages in several formats to a tool (snd-eval, snd-do, and snd-ack-event) while a tool may send the messages snd-event and snd-value to a ToolBus process. There is no direct communication possible between tools. These communication patterns are shown in Figure 1.3, "Communication between ToolBus and tools" (p. 5).

The execution and termination of the tools attached to the ToolBus can be explicitly controlled. It is also possible to connect or disconnect tools that have been executing independently of the ToolBus.

Knowledge separation. Equipped with the mechanisms provided by the ToolBus, careful control over application knowledge can be achieved as shown in Figure 1.4, "Knowledge separation in ToolBus-based application" (p. 5)where an application is depicted consisting of a user-interface (UI) and a database (DB). In a more conventional approach, elements of the user-interface, say a button, would be directly connected with functions in the database component and a strong coupling between the two components would be the result. Using the ToolBus, the two components can be completely oblivious of each other. It is only in the ToolBus script that they are configured to work together. The extra level of indirection introduced by the ToolBus thus leads to extra flexibility and decoupling.



Figure 1.4. Knowledge separation in ToolBus-based application

How to go from here?

After this brief motivation and explanation of the ToolBus architecture it is time to delve into more details. In the remainder of this chapter, we will have a look at the following topics:

- ToolBus scripts (or Tscripts, for short).
- How to write ToolBus tools.

- A brief peek at the ToolBus implementation.
- · Historical notes

Tscripts

Tscripts describe how the tools in an application cooperate. They allow the definition of a collection of concurrent processes that can communicate with each other and with the tools in the application.

Terms

Tscripts make heavy use of *terms*, simple prefix expressions that are used to exchange structured data between processes and tools. Terms are recursively defined as follows:

- A Boolean constant, integer constant, real constant, or string constant is a term, e.g., true, 37, 314e-12, or "rose".
- A value occurrence of a variable is a term, e.g., X, InitialAmount, or Highest-Bid.

Important

Variables always start with a capital letter. A value occurrence serves the purpose of using the current value of a variable.

• A result occurrence of a variable is a term, e.g., X?, InitialAmount? or Highest-Bid?.

Important

A result occurrence of a variable plays a role when this term is *matched* with another term. In the case that the match succeeds, the corresponding part of the other term is assigned to the result variable.

• A single *identifier* is a term, e.g., f, pair, or zero.

Important

Identifier always start with a lowercase letter.

- A function application is a term, e.g., pair ("rose", address ("STREE", 12345).
- A *list* is a term, e.g., [a, b, c] or [a, 1.25, "lost"].
- A placeholder is a term, e.g., <int> or add(<int>, <int>).

Matching

Term matching is used for several purposes in the ToolBus:

- To determine which actions can communicate with each other. For instance, a snd-msg and a rec-msg can only communicate if their arguments match.
- To transfer information between sender and receiver.
- To do case analysis, for instance, when receiving events from a tool.

Intuitively, the matching between two terms works as follows:

- Two terms match if they are structurally identical.
- For a value occurrence of a variable: use its current value.
- For a result occurrence of a variable: assign the matched subterm of the other term to the variable (but make this only permanent if the overall match succeeds).

This illustrated in Figure 1.5, "Example of term matching(p. 7) Before the match, two contexts are given. Each context associates some variables with a value. For instance, Context 1 associates the value 3 with variable X. For each context a term is given and the challenge is to match these two terms and to observe the effects on the two contexts. The matching of the two terms can be understood as follows:

- The top level function names are identical (both £) and both have the same number of arguments. The left term and the right term match if their arguments match.
- The first argument in the left term is X and 3 in the right term. Since, X has value 3 in Context 1, they match.
- The second argument in the left term is 4 and Z? in the right term. By assigning 4 to Z in Context 2 we achieve a match.
- The third argument in the left term is Y? and 5 in the right term. Here we achieve a match by assigning 5 to Y in Context 1.
- The fourth and last argument of both terms is 6 and thus matches.

The net result is that both terms match and that Context 1 and Context 2 are modified as shown at the bottom of the figure.

Context 1 X : 3 Y : 7	Before match	Context 2 Z : 17
f(X,4,У?,6)	Match and	f(3, <u>Z?,5,6</u>)
Context 1 X : 3 Y : 5	After successful match	Context 2 Z:4

Figure 1.5. Example of term matching

Types

The ToolBus uses a type system that is a compromise between the safety of static checking and the flexibility of dynamic typing. Another objective of the type system is to provide sufficient information to enable the automatic generation of adapter code for tools. Type are defined as follows:

- bool, int, real and str are the types of the elementary terms.
- list is the type of arbitrary lists.
- list(*Type*) is the type of lists with elements of type *Type*. For instance, list(int) is the type of lists of integers.

- *Id* is the type of all terms with function symbol *Id* (this allows the declaration of partial types). The type f, thus corresponds to the terms f, f(1), f("abc", 3) and the like.
- Id($Type_1$, ..., $Type_n$) is the type of terms with function symbol *Id* and the given types $Type_1$, ..., $Type_n$ as argument types. For instance, f(int,str) accepts f(3, "abc") but not f(3).
- [*Type*₁, ..., *Type*_n] is the type of a list of elements with the given types *Type*₁, ..., *Type*_n. For instance, [int, str] accepts [1, "abc"], but not [1,2,3].
- term is the type of an arbitrary term. And is used as escape from the more precise typing by the preceding types.

Types are used in the following ways:

- All variables have a type.
- Types are statically checked whenever possible. Only in the case of type term, dynamic checks are needed.
- Types play a role during matching: a match can also fail if the types of corresponding subterms are unequal. For instance, given I as int variable, S as str variable and T as term variable,
 - f(13) and f(I?) will match.
 - f(13) and f(S?) will fail.
 - f(13) and f(T?) will succeed.

Tscripts in detail

Overall structure

A Tscript can define the following ingredients:

- A *process definition* consisting of a process name, optional parameters and a process expression that describes the behaviour of this process.
- A *tool definition* consisting of a tool name and some operational details, such as the command to execute when the tool is started.
- A *ToolBus configuration* consisting of one or more process names (optionally followed by actual parameters) that will be created when the application is started. A Tscript may contain more than one ToolBus configuration.
- An *include file* that contains another Tscript that will be literally included.
- A constant definition.
- A conditional that allows the conditional inclusion or exclusion of parts of the Tscript.

A first example

Before delving into the details of Tscripts, it is good to have a look at the hello world application hellol.tb shown in Example 1.1, "hello1.tb" (p. 9).

Example 1.1. hello1.tb



Notes:

Here starts the definition of a process with name HELLO.

After the keyword is follows the process expression that defines the behaviour of this process.

3 The process expression consists of a single action that prints a string.

Define the initial ToolBus configuration, in this case only process HELLO will be started.

Running this example will yield the following command line dialog:

```
1> toolbus hello1.tb
Hello world, my first Tscript!
2>
```

Becoming more courageous, we show now a more ambitious Tscript hello2.tb in Example 1.2, "hello2.tb" (p. 10) that does not print the hello string itself, but executes a tool to compute it.

Example 1.2. hello2.tb

process HELLO is 🚺	
let H : hello, 🛛	
s:str 3	
in	
execute(hello, H?) 🕤 .	6
<pre>snd-eval(H, get_text) .</pre>	0
<pre>rec-value(H, text(S?)).</pre>	7
printf(S)	0
endlet	
<pre>tool hello is {command = "hello" }</pre>	Θ
toolbus(HELLO)	Θ

Notes:

- **D**efine a process HELLO.
- 2 Use a let ... in ... endlet construct to declare local variables. Variable H is declared with type hello.

B Variable S is declared with type str.

- Execute the hello tool (according to the tool definition at (p. 10)). The resulting tool identifier is assigned to variable H. Observe that the name of the tool and the type of H are identical.
- **5** Use the sequential composition operator . to combine atom actions into a larger process expression.
- Send an evaluation request to the tool we have just executed. H identifies the tool instance, and get_text is the term to be sent to the hello tool.
- In response to the evaluation request, the hello tool returns a value of the form text("Hello world from my first tool"). The actual text is extracted by the result variable S?.
- Print the string value of S.
- The definition for the hello tool. It contains the name of an executable program to be run when this tool is executed.
- The initial ToolBus configuration consisting of just the HELLO process.

All the Tscript primitives (including the ones that occur in these two simple examples) will now be described in more depth.

Process primitives

During execution, the ToolBus consists of a parallel composition of processes. The ToolBus configurations define the processes that are created at the start of the start of the application, but later on processes may die and new ones may be created.

Each process has a local state in the form of private local variables. These variables get their value through assignment and matching. They are only visible inside each process.

Processes are built-in up from atomic actions (detailed below) and atomic actions can be combined into process expressions using the following operators:

- Sequential composition P_1 . P_2 . First the actions in P_1 are executed and then the ones in P_2 .
- *Choice* $P_1 + P_2$. A choice is made between the first action in P_1 and the first action in P_2 . This choice is based on two criteria:

- An action to be selected must be *enabled*.
- If more than one action is enabled, a random choice is made.

There are various ways in which an action can be enabled (this depends on the precise action):

- An associated condition evaluates to true (see conditional and guarded command, below).
- An associated timing constraint is true.
- Required external tool results are available.
- Communication conditions are satisfied.

Once the choice for the first action has been made all remaining actions of the selected process expression P_1 or P_2 are executed as well.

- *Parallel composition* $P_1 \mid P_2$. The actions in P_1 and P_2 are executed in parallel. This means that the sequential order of the actions in P_1 respectively P_2 is respected but that apart from this constraint the actions can be executed in arbitrary order.
- *Iteration* $P_1 * P_2$. P_1 is executed repeatedly, until an action of P_2 is executed. Execution then continues with the remaining actions of P_2 .
- *Conditional* if T then P_1 else P_2 fi. The test T is evaluated and if the result is true then P_1 is executed, otherwise P_2 is executed. Note that the evaluation of the test does not count as a separate atomic action; the test is effectively attached to the first atom of P_1 respectively P_2 .
- *Guarded command* if T then P fi. The test T is evaluated and if the result is true then P_1 is executed, otherwise this command deadlocks.

Local variables

As we have seen local variables play a key role in the execution of Tscripts. They are defined using the let construct:

• let Var1 : Type1, ... in P endlet. Variables Var1, ... are declared with respective type Type1, These variables act as local variables during the execution of the process expression P. P may contain other let constructs.

Primitive actions

- *Deadlock* delta. This constant represents the process that cannot execute any further steps. During execution deadlock is always avoided as long as this is possible. A process that end in deadlock effectively terminates and disappears.
- *Silent step* tau. This constant represents one internal step in a process and resemble a dummy statement in a conventional programming language.
- Print printf. An action for generating formatted output.
- Assignment V := T. The term T is evaluated as expression (using the built-in functions) and the result is assigned to the local variable V.

Messages: synchronous communication primitives

Synchronous communication resembles an ordinary phone call: it involves two processes that can communicate at the same instant in time. In ToolBus terminology *messages* are used for synchronous communication. There are two primitives involved:

- snd-msg sends a message to another process.
- rec-msg receives a message from another process.

Two requirements have to be satisfied before communication can take place:

- The arguments of snd-msg and rec-msg match with each other.
- In addition, snd-msg respectively rec-msg are enabled in each process.

When communication takes place, the effects of the argument matching is recorded in the local state of each process and both continue execution. The observant reader may have noticed that sending and receiving is actually symmetric: by way of result variables in the arguments of snd-msg and rec-msg information may flow from sender to receiver and *vice versa*.

Notes: asynchronous broadcasting primitives

Asynchronous communication resembles conventional e-mail: it involves one sending and zero or more receiving processes that read the communicated information at a later instant in time. In ToolBus terminology *notes* are used for asynchronous communication. There four primitives involved:

- subscribe(*T*). Subscribes a process to notes that match the term *T*.
- unsubscribe(*T*). Unsubscribes a process from notes that match *T*.
- snd-note(T). Broadcast the term T to all subscribed processes. Effectively, T is placed in the private inbox of each subscribed process to be read at a later moment.
- rec-note(T). Receive a note that matches T. Effectively, the private inbox is searched for a note that matches T.
- no-note(*T*). There is no note that matches *T* in the private inbox.

Using named processes

A process definition associates a name *Pnm* (optimally followed by parameters) with a process expression *P*. These process names can be used in two ways in process expressions:

- An *inline process expression Prim(...)*: Effectively, this amount to macro substitution: *Prim* is replaced by the process expression *P* (after proper parameter substitution).
- A *process creation* create(*Pnm*(...), *Pid*?): a completely new process is created that runs in parallel with all other processes currently running in the ToolBus. The process identifier of this new process is assigned to *Pid*.

Tool primitives

There two possible scenarios for a ToolBus tool. In scenario 1, the tool is executed from the ToolBus, the tool receives a number of evaluation requests and/or generates an number of events, and finally, the ToolBus decides to terminate the execution of the tool. A variation of scenario 1 is that the tool decides to disconnect from the ToolBus and continues execution disconnect from the ToolBus application. In scenario 2, the tool is executed separately and starts its cooperation by requesting a connection with the ToolBus. Once connected, it follows the same steps as in scenario 1. The following primitives achieve this (also see Figure 1.3, "Communication between ToolBus and tools(p. 5) for the various communication patterns between ToolBus and tools):

- execute(*Tnm*, *Tid*?): Execute a tool with name *Tnm*. The result is a *tool identifier* that is bound to *Tid*. Tool identifiers are unique; if more than one instance of the same tool is executing they can be distinguished via their tool identifier. There are two additional constraints:
 - The Tscript should contain a tool definition for *Tram*.

- The variable Tid should have a type that corresponds with the tool name, i.e., it should be declared as *Tid* : *Trm*. Why? Well this in this way the implementation can track via the type of the tool identifier in each tool request, *which* tool it is and that information is essential for the automatic generation of adapter code.
- snd-terminate(*Tid*, *T*): terminates the execution of the tool instance *Tid*. The term *T* contains a reason for the termination and is usually printed by the tool on termination.
- rec-connect(*Tnm*, *Tid*?): receive a connection request for a tool with name *Tnm*. recconnect is very similar to execute. The only difference is the initiating party: for execute the ToolBus and for rec-connect the tool.
- rec-disconnect(*Tid*?): receive a disconnection request from a tool. It does not matter whether the connection with the tool was originally established via execute or rec-connect.
- snd-eval(Tid, T): send an evaluation request to a tool. All value occurrences in T are first replaced by their value before sending T to the tool. It is up to the tool to interpret the term. The usual scenario is that the outermost function symbol of T is identical to the name of a procedure in the tool and that procedure is called. The ToolBus can only send one evaluation request at a time. Only when the request is cancelled, or a value is returned by the tool, the next request can be sent to the tool.
- rec-value(*Tid*, *T*): receive a value from a tool in response to a previous snd-eval request. *T* has to match the value from the tool; this is useful for case distinctions. In many case, T consists of a single result variable, or a is a term that contains result variables.
- snd-do(*Tid*, *T*): send an evaluation request to a tool but do not expect a return value. Typically used to implement printing or logging activities.
- $rec-event(Tid, T_1, ...)$: receive an event from a tool. Events need not be handled oneby-one. The same tool may generate more than one event provided that the value of argument T_1 differs. T_1 thus serves as identification for this event.
- $snd-ack-event(Tid, T_1)$: acknowledge the completion of the handling of a previous event. Since, T_1 is identical to the T_1 in a preceding snd-event and is used to identify that event.

Timing primitives

Time can play an important role in applications, be it as ingredient in a protocol that prescribes certain time constraints, be it as watchdog that certain operations are carried out in time. The general approach in Tscripts is that a delay or timeout may be attached to every atom action. Delays and timeouts may be relative to the current time or they may be specified in absolute time. The primitives are as follows (for arbitrary atomic action A):

- *Relative delay*: A delay(*E*). Atom A can only become enabled after *E* seconds have passed.
- Absolute delay: A abs-delay(Year, Month, Day, Hour, Min, Sec). Atom A can only become enabled after the specified absolute date and time.
- *Relative timeout*: A timeout(*E*). Atom A is only enabled during the next *E* seconds.
- Absolute timeout: A abs-timeout(Year, Month, Day, Hour, Min, Sec). Atom A is only enabled until the specified absolute date and time.

Expressions

Terms can occur in Tscripts on various locations. In the majority of cases these terms are used as such; only variables are replaced by their value but no further evaluation of terms take place. There are, however, two exceptions to this general rule. In three cases, terms are evaluated:

- The test in if T then ... fi and if T then ... else ... fi.
- The right-hand side of the assignment V := T.
- In delays or timeouts.

The term is evaluated in a bottom-up manner, i.e., first arguments are evaluated and then the function is applied. Here are some examples:

- not(true) evaluates to false.
- add(mul(2,3), 4) evaluates to 10.
- greater(6,5) evaluates to true.
- first([9, 8, 7]) evaluates to 9.

A detailed overview of all built-in functions is given in XXX. They can be summarized as follows:

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

Examples of Tscripts

We are now ready to have a look at some larger examples of Tscripts.

Calculator Example

The calculator example illustrates how a calculator tool that can compute simple arithmetic expressions is shared by cooperating processes. The overall architecture is shown in Figure 1.6, "Architecture of the clock application" (p. 15). The application consists of the following 5 processes:

- CALC: the calculator *process* that regulates the access to the calculator *tool* calc, see Example 1.4, "Process CALC and tool calc" (p. 16).
- BATCH: a batch process that uses the tool batch to read an expression from file, calculate its value and write the result back to file, see Example 1.5, "Process BATCH and tool batch" (p. 16).
- UI: a user-interface process that uses the tool ui to allow a user to enter an expression and get its value back, see Example 1.6, "Process UI and tool ui'(p. 17) Observe that the processes BATCH and UI are both competing for the shared resource calculator (implemented by the process CALC and the tool calc).
- LOG: a logging process that maintains a log of all calculations that have been performed by the application, see Example 1.11, "Process LOG and tool log" (p. 19).
- CLOCK: a clock process that uses the tool clock to provide the current time, see Example 1.13, "Process CLOCK and tool clock" (p. 20).





The global structure of the Tscript calc.tb is sketched in Example 1.3, "Global structure of calc.tb" (p. 15).

Example 1.3. Global structure of calc.tb

```
process CALC is ...
tool calc is ...
    See Example 1.4, "Process CALC and tool
          calc" (p. 16)
process BATCH is ...
tool batch is ...
    See Example 1.5, "Process BATCH and tool
          batch" (p. 16)
process UI is ...
    See Example 1.6, "Process UI and tool
          ui" (p. 17)
process CALC-BUTTON is ...
    See Example 1.7, "Process CALC-BUTTON" (p. 18)
process LOG-BUTTON is ...
    See Example 1.8, "Process LOG-BUTTON" (p. 18)
process TIME-BUTTON is ...
    See Example 1.8, "Process LOG-BUTTON" (p. 18)
process QUIT-BUTTON is ...
    See Example 1.10, "Process QUIT-BUTTON" (p. 19)
process LOG is ...
    See Example 1.11, "Process LOG and tool
          log" (p. 19)
process CLOCK is ...
    See Example 1.13, "Process CLOCK and tool
          clock" (p. 20)
toolbus(CALC, BATCH, UI, LOG, CLOCK)
    See Example 1.14, "ToolBus configuration for calculator demo" (p. 20)
```



```
process CALC is
    let Tid : calc, E : str, V : term
    in
        execute(calc, Tid?).
        (2
            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"}
```

Notes:



Execute the calc tool. The tool identifier is assigned to the variable Tid, that is of type calc.

2 Begin of endless loop.

Receive a compute message.

Send an evaluation request to the calc tool and receive its value back.

Send a reply to the original compute request. By convention, the original message is included in the reply. Also send a note regarding this (expression, result) pair for the sake of logging.End of the endless loop.

Example 1.5. Process BATCH and tool batch

Notes:

Send an evaluation request to the batch tool and receive an expression back.



Ð

Communicate with the CALC process (and thus with the calc tool) to get the expression evaluated.

Send the result back to the batch tool.

The user-interface is shown in Figure 1.7, "The calc GUI" (p. 17) and behaves as follows:

• When the user presses Calc, a dialog window appears to enter an expression. The result is shown in a separate window. See Figure 1.8, "Dialog resulting from CALC-BUTTON" (p. 18).

- Pressing showLog displays all calculations so far.
- Pressing showTime displays the current time in a separate window.
- Pressing Quit end the application.

Figure 1.7. The calc GUI

Calc
showLog
showTime
Quit

Example 1.6. Process UI and tool ui



Notes:

B.

CALC-BUTTON and LOG-BUTTON are mutually exclusive and they can be activated indefinitely.

TIME-BUTTON is independent and can also be repeated indefinitely.

QUIT-BUTTON is also independent but be activated only once (for obvious reasons).

Also observe the extensive use of named process expressions like, for instance, CALC-BUTTON to give an high-level overview of the UI process. See Example 1.7, "Process CALC-BUTTON" (p. 18), Example 1.8, "Process LOG-BUTTON" (p. 18) Example 1.9, "Process TIME-BUTTON" (p. 19), and Example 1.10, "Process QUIT-BUTTON" (p. 19) for their definitions.

Example 1.7. Process CALC-BUTTON

```
process CALC-BUTTON(Tid : ui) is
   let N : int, E : str, V : term
    in
                                               0
            rec-event(Tid, N?, button(calc)) .
                                               0
            snd-eval(Tid, get-expr-dialog) .
                                                Ø
             rec-value(Tid, cancel)
            (
                                               8
             rec-value(Tid, expr(E?)) .
                snd-msg(compute, E) .
                rec-msg(compute, E, V?) .
                                               Ø
                 snd-do(Tid, display-value(V))
                                                ٠
            ) . snd-ack-event(Tid, N)
    endlet
```

Notes:

1 The Calc button is pressed; the ui tool generates an event.

- Ask the ui tool for an expression, see Figure 1.8, "Dialog resulting from CALC-BUTTON" (p. 18) for examples.
- B The user cancels the dialog; no further actions are needed.
- The user has entered an expression. Communicate with the CALC process to compute a value.
- Ask the ui tool to display the value.
- Acknowledge the event to the tool.

Figure 1.8. Dialog resulting from CALC-BUTTON



Example 1.8. Process 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
```

Example 1.9. Process TIME-BUTTON

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

Example 1.10. Process QUIT-BUTTON

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

Example 1.11. Process LOG and tool log

```
process LOG is
    let Tid : log, E : str, V : term, L : term
    in
          subscribe(compute(<str>, <term>)) .
            execute(log, Tid?).
                                                      rec-note(compute(E?, V?)) .
                (
                        snd-do(Tid, writeLog(E, V))
                +
                                                      Ø
                        rec-msg(showLog) .
                        snd-eval(Tid, readLog) .
                        rec-value(Tid, history(L?)) .
                        snd-msg(showLog, history(L))
                ) * delta
    endlet
```

Notes:

ព

Receive a note about a computation that has taken place and log it.



An alternative way to describe the LOG process is shown in Example 1.12, "Process LOG1: maintaining the log inside the ToolBus'(p. 19) Instead of running a separate log tool, the process LOG1 maintains the log in a local process variable TheLog. See the section called "Built-in functions" (p. 53) for a description of the function join that is used this example.

Example 1.12. Process LOG1: maintaining the log inside the ToolBus

```
process LOG1 is
    let TheLog : list, E : str, V : term
    in
            subscribe(compute(<str>, <term>)) .
            TheLog := [].
                        rec-note(compute(E?, V?)) .
                 (
                        TheLog := join(TheLog, [[E, V]])
                 +
                        rec-msg(showLog) .
                         snd-msg(showLog, TheLog)
                ) * delta
    endlet
```

19

Example 1.13. Process CLOCK and tool 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
```

The complete ToolBus configuration that describes the start of the calculator demo is shown in Example 1.14, "ToolBus configuration for calculator demo'(p. 20) It starts the mentioned processes in parallel and from that moment on user-interaction and the activities of the BATCH process will drive the execution.

Example 1.14. ToolBus configuration for calculator demo

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

Auction Example

In the *classical auction*, the auction master and all bidders are in the same room and interact with each other according to a fixed protocol. It is shown in Figure 1.9, "Classical auction" (p. 20) The steps in the protocol are:

- 1. The auction master introduces a new item for sale and sets an initial price for it.
- 2. Next, bidders raise their hand and shout a new bid that is to be acknowledged by the auction master.
- 3. Step 2 is repeated as long as new bids come in.
- 4. When no new bids are being made, the auction master asks for "any higher bid?" and waits during a fixed period.
- 5. If no new bids come in during this period, the auction master declares the item for sale to be sold to the highest bidder.
- 6. If a new bid comes in during this period, the procedure continues with step 2.

Figure 1.9. Classical auction



A striking aspect of the classical auction is that the auction master and the bidders can *see* each other. This is a great communication and synchronization tool. In the case of a *distributed auction*, auction master and bidders are on different locations and can only communicate via the Internet, see Figure 1.10, "Distributed auction" (p. 21).

Figure 1.10. Distributed auction



The communication and synchronization in a distributed auction has to be described explicitly and requires answers to questions like:

- How are bids synchronized?
- How to inform bidders about the highest bid?
- How to decide when bidding is over and the item is to be sold?
- How to handle bidders that come and go during the auction?

The auction application to be described answers these questions and has an architecture as shown in Figure 1.11, "Architecture of the auction example" (p. 21).

Figure 1.11. Architecture of the auction example



The auction application consists of a variable number of processes:

- Auction: the process Auction orchestrates the complete auction and is controlled by the tool master that enables the auction master to offer new items for sale and to monitor the progress of the auction.
- Bidder: for each new bidder that enters the auction a new process Bidder and tool bidder are created. The tool bidder keep the user informed and allows her or him to submit new bids.

The global structure of the Tscript auction.tb is sketched in Example 1.15, "Global structure of auction.tb" (p. 22).

Example 1.15. Global structure of auction.tb

```
process Auction is ...
See Example 1.16, "Process Auction" (p. 22)
tool master is ...
process ConnectBidder is ...
See Example 1.17, "Process ConnectBidder" (p. 23)
tool batch is ...
process OneSale is ...
See Example 1.18, "Process OneSale" (p. 24)
process Bidder is ...
See Example 1.22, "Process Bidder" (p. 27)
toolbus(Auction)
```

Example 1.16. Process Auction



Example 1.17. Process ConnectBidder

```
process ConnectBidder(Mid : master, Bid : bidder?) is
let Pid : int, Name : str
in
    rec-connect(Bid?) .
    create(Bidder(Bid), Pid?) .
    snd-eval(Bid, get-name) .
    rec-value(Bid, name(Name?)) .
    snd-do(Mid, new-bidder(Bid, Name))
endlet
```

Notes:



8

Receive a connection request from a new bidder tool.

Create a new Bidder process that orchestrates the behaviour of this bidder.

Ask bidder for its name.

Send the name of the bidder to the master tool.

Example 1.18. Process OneSale

```
process OneSale(Mid : master) is
                                    Ø
  let Descr : str,
                                    0
      InAmount : int,
                                    Ø
      Amount : int,
                                    8
      HighestBid : int,
                                    Ø
      Final : bool,
      Sold : bool,
                                    Θ
      Bid : bidder
                                                     \mathbf{Z}
  in rec-event(Mid, new-item(Descr?, InAmount?)) .
      HighestBid := InAmount .
                                                      8
      snd-note(new-item(Descr, InAmount)) .
      Final := false . Sold := false .
      (
                                                       Ø
         Here is the main logic of OneSale
      ) * if Sold then
             snd-ack-event(Mid, new-item(Descr, InAmount))
          fi
  endlet
```

Notes:

Π

Descr contains a textual description of the current item for sale.

- InAmount is the initial amount asked for the item.
- Amount is the value of the current bid.
- HighestBid is the highest bid so far for this item.
- Two Boolean values control the logic of the bidding process. Final is true when the call for final bids has been issued and Sold is true when the item has been sold.Bidder is a new bidder tool that has connected during the sale.
- Bidder is a new bidder tool that has connected during the salThe auction master wants to initiate the sale of a new item.
- Inform all connected bidders about the new item that is for sale.
- The detailed logic is explained in Example 1.19, "Process OneSale, main logic'(p. 25) Example 1.20, "Process OneSale, handling one bid(p. 25) and Example 1.21, "Process OneSale, handling other cases" (p. 26).

Example 1.19. Process OneSale, main logic

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

The main process logic consists of four parts:

Handle one incoming bid, see Example 1.20, "Process OneSale, handling one bid" (p. 25).

Start the "any higher bid" procedure, see Example 1.21, "Process OneSale, handling other cases" (p. 26).

Sell the item when no further bids are received, see Example 1.21, "Process OneSale, handling other cases" (p. 26).

Connect a bidder during the sale, see Example 1.21, "Process OneSale, handling other cases" (p. 26).

Example 1.20. Process OneSale, handling one bid

```
( if not(Sold) then
                                            O
   rec-msg(bid(Bid?, Amount?)) .
                                            2
    snd-do(Mid, new-bid(Bid, Amount)) .
    if less-equal(Amount, HighestBid) then
                                            Ø
        snd-msg(Bid, rejected)
    else
                                            R
        HighestBid := Amount .
                                            ٨
        snd-msg(Bid, accepted) .
                                            B
        snd-note(update-bid(Amount)) .
        snd-do(Mid, update-highest-bid(Bid, Amount)) .
        Final := false
    fi
  fi
+ if not(or(Final, Sold)) then
                                  ... fi
+ if and(Final, not(Sold)) then ... fi
+ ConnectBidder(Mid, Bid?) ...
) * if Sold then ... fi
```

Notes:

Receive a bid from a bidder.

Inform the auction master about the new bid.

Reject the bid when it is too low.

Remember this bid as the highest bid so far.

5 Inform the bidder that his bid is accepted.

6 Inform all connected bidders that there is higher bid.

7 Update the status of the auction master.

Example 1.21. Process OneSale, handling other cases

```
( if not(Sold) then ... fi
+ if not(or(Final, Sold))
                           then
                                                  0
      snd-note(any-higher-bid) delay(sec(10)) .
     snd-do(Mid, any-higher-bid(10)) .
                                                   Ø
     Final := true
 fi
+ if and(Final, not(Sold)) 🕤 then
                                                  0
      snd-note(sold(HighestBid)) delay(sec(10)) .
                                                   Ø
     Sold := true
   fi
                                                   2
+ ConnectBidder(Mid, Bid?) .
                                                   ً₿
  snd-msg(Bid, new-item(Descr, HighestBid)) .
                                                   Ø
  Final := false
 * if Sold then ... fi
```

Notes:

The item is not yet sold, but we have not yet asked for a final bid.

Wait for 10 seconds and then ask for final bids.

Inform the auction master and remember that we asked for final bids.

The item is not yet sold but we have already asked for final bids.

G Wait another 10 seconds and inform all bidders that the item has been sold.

B Record that the item is sold.

During the sale a new bidder wants to connect.

8 Inform the new bidder about the progress of the auction.

Restart the bidding procedure; this overrules the call for final bids.

```
Example 1.22. Process Bidder
```

```
process Bidder(Bid : bidder) is
  let Descr : str, Amount : int, Acceptance : term
  in
                                                  Ø
     subscribe(new-item(<str>, <int>)) .
     subscribe(update-bid(<int>)) .
     subscribe(sold(<int>)) .
     subscribe(any-higher-bid) .
                                                  2
     ( ( rec-msg(Bid, new-item(Descr?, Amount?))
                                                  Ø
         + rec-note(new-item(Descr?, Amount?))
         + rec-disconnect(Bid) 🕤. delta
        ) .
                                                  Ø
       snd-do(Bid, new-item(Descr, Amount)) .
                                                  0
       ( rec-event(Bid, bid(Amount?)) .
                                                  71
         snd-msg(bid(Bid, Amount)) .
         rec-msg(Bid, Acceptance?) .
         snd-do(Bid, accept(Acceptance)) .
         snd-ack-event(Bid, bid(Amount))
       ) *
                                                  ً₿
     + rec-note(update-bid(Amount?)) .
       snd-do(Bid, update-bid(Amount))
     + rec-note(any-higher-bid) .
       snd-do(Bid, any-higher-bid)
                                                  Ø
     + rec-disconnect(Bid) . delta
     )
      *
                                                  0
       rec-note(sold(Amount?))
       snd-do(Bid, sold(Amount))
     )* delta
  endlet
```

Notes:

Subscribe to all relevant notes.

2 Get information about the current item for sale, directly after connecting.

Get information about the current item for sale during regular progress of the auction.

Disconnect between sales.

G Inform the bidder tool by updating the information about the item for sale.

6 This bidder want to bid on the current item for sale.

Pass this bid on and await its acceptance. The result is returned to the bidder tool and the event is acknowledged.

8 Handle the informative notes about the update of the highest bid and the request for any higher bids.

Disconnect during a sale.

The item is sold.

Wave Example

Have you ever considered a (guitar) string that is attached at both ends and wondered how the movements of the strings can be simulated? Although this is a completely atypical application of the ToolBus, it is fun to do so we will delve into the details.

In mathematical physics, the vibrating string is described by the so-called one-dimensional wave equation that describes a discrete approximation of the continuous string. The discretization is achieved by sampling the amplitude of the string at certain points i = 1, ..., N, where *N* is the number of points. The amplitude at point *i* at time *t* can now be described by $y_{i+1}(t)$ (see also Figure 1.12, "One-dimensional wave equation" (p. 28)) that is defined as follows:

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

In other words, the amplitude at point i and time t depends on:

- the current amplitude,
- the previous amplitude at this point,
- the current amplitude of the left neighbour, and
- the current amplitude of the right neighbour.

It also depends on the function *F* defined as follows:

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

where

- #x is the (small) interval between sampling points, and
- c is a constant representing the propagation velocity of the wave.

Figure 1.12. One-dimensional wave equation



After these preparations, we have to define the architecture of a ToolBus application that can simulate the behaviour of a string. The key idea is to use a separate ToolBus process to represent the behaviour of each sampling point. The architecture is shown in Figure 1.13, "Architecture of the wave example" (p. 29) and consists of the following processes and tools:

- Process Pend models an end point of the string, see Example 1.26, "Process Pend" (p. 31) .Two instances are used to model the left and right end point.
- Process P models one sampling point, see Example 1.27, "Process P"(p. 32) N-1 instances are used to model all intermediate points.
- The auxiliary process F computes the function F discussed above, see Example 1.28, "Auxiliary process F" (p. 33).
- Process makeWave constructs *N* connected instances of processes *P* and two end points *Pend*, see Example 1.24, "Process MakeWave" (p. 30).
- The tool display visualizes the simulation.

MakeWave Pend Pend P P P P ... ToolBus ToolBus master Tools

Figure 1.13. Architecture of the wave example

The global structure of the Tscript wave.tb is shown in Example 1.23, "Global structure of wave.tb" (p. 29).

Example 1.23. Global structure of wave.tb

```
process MakeWave ... See Example 1.24, "Process MakeWave" (p. 30)process Pend ...See Example 1.26, "Process Pend" (p. 31)process P ...See Example 1.27, "Process P" (p. 32)process F ...See Example 1.28, "Auxiliary process F" (p. 33)toolbus(MakeWave(...))
```

```
Example 1.24. Process MakeWave
```

```
process MakeWave(N : int) is
  let Tid : display, Id : int, I : int, L : int, R : int
  in
                                     0
     execute(display, Tid?) .
     snd-do(Tid, mk-wave(N)) .
                                     2
     create(Pend(Tid, 0, 1), Id?).
     L := sub(N,1).
     create(Pend(Tid, N, L), Id?) .
                                     Ø
     I := 1 .
     if less(I, N) then
        L := sub(I, 1) . R := add(I, 1) .
        create(P(Tid, L, I, R, 1.0, 1.0), Id?) .
        I := add(I, 1)
     fi *
     shutdown("end") delay(sec(60)) <sup>①</sup>
  endlet
tool display is { command = "wish-adapter -script ui-wave.tcl"}
```

Notes:

0

Z

Execute the display tool and initialize it to show *N* points. Figure 1.14, "User-interface of wave demo" (p. 30) shows the display tool in action during a later stage of the simulation. Create the two end points with index 0 and *N*.

Create the intermediate points 1,..., *N*-1 in a loop.

Run the demo for one minute.

Figure 1.14. User-interface of wave demo



Example 1.25. ToolBus configuration for wave

```
toolbus(MakeWave(8))
```

Example 1.26. Process Pend

```
process Pend(Tid : display, I 1: int, NB 2: int) is
let W : real
in
( rec-msg(NB, I, W?) || snd-msg(I, NB, 0.0) || 3
snd-do(Tid, update(I, 0.0))
) * delta
endlet
```

Notes:

I is the index of this end point.

NB is the index of the neighbouring point.

Receive the amplitude of the neighbour and send our own zero amplitude to the neighbour. Since a parallel operator | | is used, these communications can appear in any order. Display the zero amplitude of this end point on the display.



Example 1.27. Process P

```
process P(Tid : display, 🚺
                         0
         L : int,
                         Ø
         I : int,
                         8
         R : int,
         Dstart : real, 5
         Estart : real 🔒
         ) is
 let AL : real, AR : real, D : real, D1 : real, E : real
  in
    D := Dstart . E := Estart .
                                    0
     ( ( rec-msg(L, I, AL?)
      || rec-msg(R, I, AR?)
                                    8
       || snd-msg(I, L, E)
       || snd-msg(I, R, E)
       || snd-do(Tid, update(I, E)) 🖸
       ) .
      D1 := E .
                                    Θ
      F(E, D, AL, AR, E?).
      D := D1
     ) * delta
  endlet
```

Notes:

1

0	Tid is the tool identifier of the display tool.
2)	L is the index of the left neighbour of this point.
Ð	I is the index of this point.
÷	R is the index of the right neighbour of this point.
5	Dstart is the previous amplitude of this point.
8	Estart is the current amplitude of this point.
7	Receive amplitudes from our neighbours.
8	Send our own amplitude to our neighbours.
9	Show current amplitude on the display.
8	Compute a new value for our amplitude by applying F

Example 1.28. Auxiliary process F



Notes:

- Recall that we are computing $2z_1 z_2 + (c \# t/\# x)^2 (z_3 2z_1 + z_4)$ and that the main challenge is to write this formula in prefix form.
- **?** Take an arbitrary (small) value for $(c \# t/\# x)^2$.



- $(c \# t / \# x)^2 * ...$
- **5** ... + $(z_3 2z_1 + z_4)$

Executing ToolBus and tools

The ToolBus interpreter (**ToolBus**) and all tools have some standard program arguments in common, but they have some specific arguments as well. In this section we describe all possible program arguments and the way to execute **ToolBus** and tools.

Common arguments

ToolBus and tools have the following optional arguments in common:

- -help: prints a description of all arguments of the ToolBus or tool.
- -Pport_name: defines the "well known socket" port_name to which all tools temporarily connect in order to set up their own private socket that connects them permanently to the ToolBus interpreter. When omitted, socket 8998 will be used.

Warning

Not yet implemented.

Note that explicit arguments defining the sockets are *only* needed when several ToolBus interpreters are running simultaneously on the *same* host machine.

ToolBus arguments

The *script_name* (see below) given as argument to the ToolBus is always preprocessed by a preprocessor before it is parsed as a Tscript. In this way, directives like, e.g., #define, #include and #ifdef can be used freely in Tscripts. The following preprocessor arguments are accepted by the **ToolBus** command:

- -Idir: append directory dir to the list of directories searched for include files.
- -Dname: defines name with the string "1" as its definition.
- Dname=defn: defines name with defn as definition.

Other arguments specific for the ToolBus command are:

Warning

The following arguments will probably be supported differently in ToolBusNG.

- -viewer: execute the ToolBus viewer that enables step-by-step execution and inspection of the state of each process state.
- -gentifs: only generate tool interfaces for all tools used in the script in a language independent format. For a script file named script.tb the tool interfaces are written to script.tifs.Do not execute the script.

Warning

Not yet implemented.

• -fixed-seed: use a fixed seed for the random generator used by the interpreter for scheduling processes and selecting alternatives in processes. By default, the random generator is initialized with the current time the **ToolBus** command is given. Using the -fixed-seed option makes the execution of the script reproducible across multiple runs of the **ToolBus** command.

Warning

Not yet implemented.

• -Sscript_name: any other argument is the name of the ToolBus script to be interpreted.

As an example, consider first

toolbus -Shello.tb

which starts interpreting the script hello.tb. Next, consider

toolbus -Imy-include-dir -DCNT=33 -Swave.tb

which searches the directory my-include-dir for files used in #include directives in the script wave.tb and it will define the name CNT with value 33. All occurrences of CNT in the script will be replaced by this value before parsing it as a Tscript. Finally,

toolbus -gentifs -Shello.tb

produces the tool interfaces file hello.tifs.

Tool arguments

Warning

This section needs some work. Arguments specific for tools are:

- -TB_HOST *host_name*: defines the host machine *host_name* on which the ToolBus interpreter is running and to which the tool should be connected. When omitted, the ToolBus interpreter should be running on the same host as the tool.
- -TB_TOOL_NAME tool_name: the tool name as defined in the Tscript (added automatically, when a tool is executed by the ToolBus).
- -TB_TOOL_ID *Id*: internal tool identifier of this tool execution (added automatically, when a tool is executed by the ToolBus).

The execution of a tool can start in two ways:

- The tool is started by an execute command in the Tscript.
- The initiative to execute the tool is taken outside the ToolBus. This requires that the script contains a rec-connect for this particular tool.

When ToolBus and tool are running on different host machines, it is important to define the host machine on which the ToolBus interpreter is running when starting the execution of the tool. As an example, consider the **hello** application described in Example 1.2, "hello2.tb" (p. 10) The **hello** tool will be executed by the ToolBus using the command

hello -P8998 -TB_HOST host1.institute.nl

when running on machine host1.institute.nl. Suppose, we replace the explicit execute in Example 1.2, "hello2.tb" (p. 10) by a rec-connect as shown in Figure~\ref{fig:hello3.tb}. We may then manually start the **hello** tool by typing

hello

where we use the default values for the input/output sockets and assume that tool and ToolBus interpreter are both running on the same host (i.e., hostl.institute.nl). Starting the execution from *another* host is achieved by typing (on, say, host2.institute.nl):

hello -TB_HOST host1.institute.nl

ToolBus tools

There are some general issues to understand about ToolBus tools and we cover them here. First, the global structure of a tool is explained in the section called "*The global structure of a ToolBus tool*" (p. 35) Next, we describe how tool adapters work in the section called "*Adapters for tools and languages*" (p. 36) Finally, we cover in the section called "*Automatic generation of tool interfaces*" (p. 37) the automatic generation of tool interfaces that is needed for some tool implementation languages.

The global structure of a ToolBus tool



Figure 1.15. Global tool organization

In its simplest form, a tool is a box connected via an input and an output port to a ToolBus. In the most general case, a tool has

• one input port from the ToolBus to the tool and can receive tree structures (terms) via this port;

- one output port from the tool to the ToolBus and can send terms to the ToolBus via this port;
- zero or more *term ports* to receive terms from other sources;
- zero or more *character ports* to receive character data from other sources.

This global, architectural, structure of a tool is shown in Figure 1.15, "Global tool organization" (p. 35). With each input port, an *event handler* is associated that takes care of the processing of the data received via that port and is responsible for returning a result (if any). One tool may thus contain several event handlers. When a request is received, the following steps are taken:

- The data received are parsed to check that they form a legal ToolBus term *T*. (If this is impossible, a warning message is generated).
- The event handler is called with *T* as argument.
- The event handler can do arbitrary processing needed to decompose *T*, to determine what has to be done, and perform any desired computation.
- The event handler returns either:
 - a legal ToolBus term representing a reply to be sent back to the ToolBus.
 - NULL indicating that there is no reply.

The global mode of operation of a tool is now:

- receive data on any input port and respond to this by sending some term (or NULL) to the ToolBus; or
- take the initiative to send a term to the ToolBus (typically to inform the ToolBus about some external event).

A tool is thus on the one hand a reactive engine that responds to a request from the ToolBus and returns the result back to the ToolBus in the form of a term (e.g., calculate the value of some expression), but on the other hand it can also take the initiative to send a term to the ToolBus (e.g., generate an event when a user pushes some button).

Adapters for tools and languages



Figure 1.16. Two organizations of a tool adapter

The main purpose of adapters is to act as small *wrappers* around existing programs or programming languages in order to transform them into tools that can be connected to the ToolBus. There exist two global strategies for constructing adapters:

• The adapter and the program to be adapted are executed as separate (Unix) processes. This structure is sketched in Figure 1.16, "Two organizations of a tool adapter? (p. 36) The advantage of this approach is that no access is needed to the source code of the program: it can remain a black box. Another advantage is that adapters may be reused for the adaptation of different programs. A possible disadvantage is some loss in efficiency.

In this category a further subdivision is possible:

- The program is executed once as a child process of the adapter and all snd-eval/snd-do requests are directed to this child process. The program can thus maintain an internal state between requests.
- The same program is executed as a child process of the adapter for each snd-eval/snd-do request.
- A different program is executed as a child process of the adapter for each snd-eval/snd-do request.
- Integrate the adapter and the software to be adapted into a single (Unix) process. This approach permits the most detailed adaptation of the program and is also the most efficient solution. This approach leads, however, to potentially less reusable adapters than the previous approach.

In order to achieve some uniformity, the current collection of adapters have the following optional program arguments in common:

- -cmd: the (default) program to be executed by the adapter. All arguments of the adapter that follow -cmd are interpreted as the name and arguments of the program to be executed.
- All tool arguments, see the section called "Tool arguments" (p. 34).

Automatic generation of tool interfaces



Figure 1.17. Automatic generation of tool interfaces

The interface code for each tool depends on the particulars of the Tscript in which it is used. Changing the number of arguments in an evaluation request to the tool, or adding a new request, requires making changes to the interface code that are easily forgotten and therefore error prone. Another observation is that the interface code for different tools has a lot in common. An obvious solution to both problems

is to *generate* tool interfaces automatically, given a Tscript. This generation process is shown in Figure 1.17, "Automatic generation of tool interfaces" (p. 37) and consists of two steps:

• Generate a language-independent description of all tool interfaces used in the script. This amounts to a static analysis of all tool communication in the script. It is achieved by using the -gentifs option of the ToolBus interpreter. For instance,

toolbus -gentifs hello2.tb

will create a file hello2.tifs containing the tool interfaces.

• Use the language independent interface description to generate a tool interface for a specific tool in a specific implementation language. The generator **tifstoc** exists for generating C tool interfaces. It is called as follows:

```
tifstoc -tool Name TifsFile
```

and generates a file named Name.tif.c. For the hello example, we would have, for instance:

tifstoc -tool hello hello2.tifs

The resulting file hello.tif.c is shown in Example 1.30, "The generated file hello.tif.c" (p. 45).

In Figure 1.17, "Automatic generation of tool interfaces" (p. 37) it is also shown how tool interface generators for *other* languages (e.g., Java, Cobol) fit into this scheme. In addition to **tifstoc**, we used to support the generation of Java interfaces by way of **tifstojava**. In the current ToolBus implementation, this is no longer necessary, see the section called "*Writing ToolBus tools in Java*" (p. 47).

Writing ToolBus tools in C

Although ToolBus tools can be implemented in many languages (including Java, C++, Tcl/Tk, ASF+SDF, and others) we start explaining how tools can be written in C. In other languages identical notions will be used with only minor adjustments to language-specific features and limitations. Writing tools in C amounts to:

• ATerms: the essential data type that is used to exchange information between tool and TB

Note

Add ref to ATerms

- The global structure of a ToolBus tool, see the section called "*The global structure of a ToolBus tool*" (p. 35).
- The ToolBus Application Programmer's Interface, see the section called "*The ToolBus API*" (p. 39).
- Compiling ToolBus tools written in C, the section called "*Compiling ToolBus tools written in C*" (p. 42).
- Generating tool interfaces with {\tt tifstoc}, the section called "Automatic generation of tool interfaces" (p. 37).

The include file atb-tool.h

Each tool needs to include the file atb-tool.h which defines some basic types as well as the set of library functions available. It consists of

• An include of <aterm1.h>.

- Defines ATBhandler: the type of event handlers.
- Defines the prototypes of all library functions

The tool library libATB.a

When compiling tools, the library libATB.a must be specified in order to make the tool library available (using the -lATB option of the C compiler). It provides the following functions:

- ATBinit: tool initialization, see the section called "ATBinit" (p. 39).
- ATBconnect: to connect the tool to the ToolBus, the section called "ATBconnect" (p. 40).
- ATBdisconnect: to disconnect the tool from the ToolBus, the section called "ATBconnect" (p. 40).
- ATBeventloop: a standard event loop for a tool, see the section called "*ATBeventloop*" (p. 40).
- ATBpostEvent send an event to the ToolBus, see the section called "ATBpostEvent" (p. 41).

In the following section, we will describe these functions.

The ToolBus API

During the initialization of each tool, some preparations have to made before the tool can be properly connected to the ToolBus. These preparations include

- Defining the *name* of the tool as it is known from a tool declaration in a Tscript.
- Parsing standard program arguments that are passed to the tool when it is started.
- Creating a pair of socket connections with a ToolBus interpreter.
- Starting an event loop.

During execution of the event loop, the tool can either *receive* terms from the ToolBus or it can take the initiative to *send* terms to the ToolBus. It is thus possible for a tool to both respond to ToolBus requests and *asynchronously* send terms to the ToolBus.

ATBinit

The initialization of the ToolBus API is achieved by

```
int ATBinit(int argc, char *argv[], ATerm *bottomOfStack).
```

This initializes the ToolBus API as well as the ATerm library that is used by it.

The standard program arguments that are passed (via *argc* and *argv*) are fully described in the section called "*Executing ToolBus and tools*" (p. 33) Particularly important is that the tool is initialized with a proper name. It should be literally equal (including the case of letters) to a tool name as appearing in a tool declaration in the Tscript. This is important since the tool name will be used when the tool is connected to the ToolBus. Note that ATBinit also initializes the ATerm library (hence the *bottomOfStack* argument, see

Note

Ref to Section \ref{ATinit}) in ATerm manual.

The return value indicates whether or not the ToolBus host could be found: 0 indicates that all is well, and -1 indicates an error, in which case the standard variable errno of the C run-time system is set to indicate which error.

ATBconnect

A tool can be connected to the ToolBus using *term ports* that can be using for sending and receiving data in the form of complete terms. Two aspects of term ports are important: the input channel used for the actual data transfer and the *handler* that takes care of processing input terms when they arrive. The connection is established as follows:

Here, toolname is the tool name to be used, host is the machine where the ToolBus is executing, port is the file descriptor of the channel to be used, and h is the handler to associated with this connection. If value NULL is passed as toolname or host, default values are used that are taken from argv that was passed to ATBinit. The same is true when -1 is passed as value for port. The return value of ATBconnect is either -1 (failure) or a positive number (the connection succeeded and the result is the file descriptor of the resulting socket connection with the ToolBus). Handlers for term ports are functions from ATerm to ATerm and have the type:

ATerm some_handler(int conn, ATerm input)

The argument *conn* is the connection along which the input term was received and *input* is the actual term received. The term returned by the handler is the reply to be sent to the ToolBus in response to this input event, or NULL if no reply is needed. In this fashion, an arbitrary number of term input ports can be set up which will be read in parallel: as soon as a term arrives at one of the ports the associated handler is activated. A connection can be terminated as follows:

```
void ATBdisconnect(int conn)
```

where *conn* is a connection that has been created earlier using ATBconnect.

ATBeventloop

Many tools first establish a number of term ports and then enter an infinite loop that processes input events. The function

```
int ATBeventloop(void)
```

captures this idea. It never returns, unless something goes wrong. We can now give a skeleton that many tools have in common:

return 0;

ATBpostEvent

}

So far, we have seen primitives for tools that only receive terms from the ToolBus. In the case of events that are generated by a tool, a term needs to be sent from the tool to the ToolBus. This can be achieved using

int ATBpostEvent(int conn, ATerm term)

which sends *term* along the port *conn*. Failure is indicated by the return value -1. A typical usage is:

ATBpostEvent(conn, ATmake(button("ok"))).

Primitives for advanced control flow

Tool programming amounts, in essence, to event driven programming: most of the time a tool is awaiting the arrival of data on one of its ports and when the data are there, a reply is sent to the ToolBus by the handler associated with that port. In computation-intensive tools, the need may arise to check for the availability of incoming data from the ToolBus during computations. In those cases, ATBeventloop may not offer enough flexibility. More customized control flow can be achieved using the following functions. Observe that these function are parameterized with a specific ToolBus connection (as returned by ATBconnect) and can be used to handle situations where a single tool is connected with *more than one* ToolBus.

Checking if there is input awaiting on a ToolBus connection is done by:

ATbool ATBpeekOne(int conn)

This function returns ATtrue if incoming data from a ToolBus are available on the connection conn.

Similarly, the availability of data on any connection may be checked by:

```
int ATBpeekAny(void)
```

If input is waiting, the appropriate connection is returned. Otherwise -1 is returned. The sequence of activities needed for handling (once) the data available from a specific connection is captured by the function

void ATBhandleOne(int conn)

This amounts to calling the handler associated with connection *conn* with the available data as input term. Similarly, the data from *any* connection is handled by

int ATBhandleAny(void)

which returns -1 if anything goes wrong.

Finally, the function

int ATBgetDescriptors(fd_set *set)

Gathers all ToolBus connection file descriptors in a single descriptor set. The return value indicates the maximum value of any descriptor in the set.

Control flow patterns

Given, the control flow primitives in the previous section, we can express various common control flow patterns.

The function ATBeventloop can be expressed with the primitives just introduced:

```
int ATBeventloop(void)
{ int conn;
  while(ATtrue)
   {
        n = ATBhandleAny();
        if(n < 0)
        return -1;
    }
}</pre>
```

Another style mixes the handling of input from the ToolBus, with other computations:

```
while(ATtrue)
{
    if(n = ATBpeekAny() >= 0) /* if there is an incoming event */
    ATBhandleOne(n); /* handle it */
    else {
        ... /* perform other computation */
    }
}
```

In some tools, a mixture of passively awaiting input and actively sending terms to the ToolBus can be seen. Using ATBwriteTerm, the most general global event loop of a tool becomes:

```
while(ATtrue)
{
... ATBwriteTerm(c1,e1); ...; ATBwriteTerm(cn,en); ...
ATBhandleAny();
}
```

In other words, each iteration starts by sending zero or more terms to the ToolBus (using ATBwriteTerm) and ends with processing one event coming from some port (using ATBhandleAny). The Tscript being used should, of course, be able to receive such events.

Compiling ToolBus tools written in C

When compiling a tool written in C the following questions should be answered:

- Where is the include file aterm1.h (or aterm2.h if you use the more sophisticated parts of the ATerm library)?
- Where is the include file atb-tool.h?
- Where is the ATerm library libATerm.a?
- Where is the ToolBus API library libATB.a?
- Which other libraries are needed to compile the tool?

The answers to these questions are clearly system dependent. There are two strategies to answer them. Strategy 1: find the desired locations on your system and hard code them in the compilation command. This will lead to a call to the C compiler with the following arguments:

- -Idir-where-aterm1.h-is
- -Idir-where-ATB-tool.h-is
- hello.c -o hello
- -Ldir-where-libATerm.a-is
- -lATerm

- -Ldir-where-libATB.a-is
- -latb
- other libraries.

Strategy 2: write a make file that encodes this information. As a result, the location information is hardwired in the make file rather than in a command that has to be repeated over and over again.

Automatic generation of C tool interfaces

As already explained in the section called "Automatic generation of tool interface"s(p. 37) tool interfaces can be generated from a given Tscript for a given tool name. The ToolBus can generate a language-independent .tifs file, when it is started with the -gentifs option. In the case of C, the command **tifstoc** generates a tool interface in C for use with the ATerm library. The generated interface consists of two files:

- a C source file (hello2.tif.c in the example below), and
- a C header file (hello2.tif.h in the example below).

In the header file a number of interface functions is declared, one for each element in the input signature of the tool. It is up to the writer of the tool to provide an implementation for these functions. The generated C file contains a handler function that analyzes incoming terms from the ToolBus, and delegates actual processing to the appropriate interface function.

We will use the Tscript hello2.tb shown earlier in Example 1.2, "hello2.tb" (p. 10) and describe all the steps needed to write and compile the hello tool.

Step 1: generate tifs

Using the command:

```
toolbus -gentifs hello2.tb
```

we generate a file called hello2.tifs. It contains information amount the interfaces for all tools that are used in a given Tscript.

Warning

The -gentifs flag is not yet implemented.

Step 2: generate C tool interface

Using the command:

```
tifstoc -tool hello test.tifs
```

we generate two files:

- the header file hello.tif.h, see Example 1.29, "The generated header file hello.tif.h" (p. 44).
- the source file hello.tif.c, see Example 1.30, "The generated file hello.tif.c" (p. 45).

Example 1.29. The generated header file hello.tif.h

```
**
 * This file is generated by tifstoc. Do not edit!
 * Generated from tifs for tool 'hello' (prefix='')
 */
#ifndef _HELLO_H
#define _HELLO_H
#include <atb-tool.h>
/* Prototypes for functions called from the event handler */
ATerm get_text(int conn);
void rec_terminate(int conn, ATerm);
extern ATerm hello_handler(int conn, ATerm sigs);
```

#endif

Only the functions get_text and rec_terminate together with a simple main function have to be implemented to build a fully functional ToolBus tool.

```
Example 1.30. The generated file hello.tif.c
```

```
/**
 * This file is generated by tifstoc. Do not edit!
 * Generated from tifs for tool 'hello' (prefix='')
 */
#include "hello.tif.h"
#define NR_SIG_ENTRIES
                        2
                                                O
static char *signature[NR_SIG_ENTRIES] = {
  "rec-eval(<hello>,get_text)",
  "rec-terminate(<hello>,<term>)",
};
/* Event handler for tool 'hello' */
                                                [2]
ATerm hello_handler(int conn, ATerm term)
 ATerm in, out;
  /* We need some temporary variables during matching */
  ATerm t0;
  if(ATmatch(term, "rec-eval(get_text)")) {
   return get_text(conn);
  if(ATmatch(term, "rec-terminate(<term>)", &t0)) {
   rec_terminate(conn, t0);
   return NULL;
  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%\n\tl\n",
                result);
   return NULL;
  }
  ATerror("tool hello cannot handle term %t", term);
  return NULL; /* Silence the compiler */
/* Check the signature of the tool 'hello' */
                                                Ø
ATerm hello_checker(int conn, ATerm siglist)
 return ATBcheckSignature(siglist, signature, NR_SIG_ENTRIES);
}
```

Notes:

An array of signature definitions (signature) that contains the argument and return types of each interface function.



Step 3: write main

As mentioned earlier, the only thing needed to implement the actual hello tool, is the implementation of the two interface functions get_txt and rec_terminate, and the implementation of main to get things going. We will first take a look at the initialization stuff that the main function of the hello tool has to do, see Example 1.31, "main function of hello tool" (p. 46).

Example 1.31. main function of hello tool

```
#include <stdlib.h>
#include "hello.tif.h"
int main(int argc, char *argv[])
                                                          ATerm bottomOfStack;
                                                          2
 ATBinit(argc, argv, &bottomOfStack);
                                                          Ø
 if(ATBconnect(NULL, NULL, -1, testing_handler) >= 0) {
                                                          8
   ATBeventloop();
  } else {
   fprintf(stderr,
            "Could not connect to the ToolBus, giving up!\n");
   return -1;
 return 0;
```

Notes:



The variable bottomOfStack is needed by the ATerm library to determine where to look for the stack. It is passed as argument to ATBinit.

The variables argc and argv are passed unchanged to ATBinit, so the ToolBus library can look for default values for things like the ToolBus' well-known socket address and the ToolBus host name.

The call to ATBCONNECT connects to a running ToolBus, and requires four arguments: a character string representing the tool name, a character string representing the host name of the ToolBus to connect to, the port number of the ToolBus to connect to, and a handler function. Passing NULL, NULL, and -1 respectively as the tool name, the host name, and the port number cause the defaults for these values to be used instead.

When all goes well, the call to ATBeventloop starts the main ToolBus eventloop and the tool will be ready to receive requests from the ToolBus.

Step 4: implement interface functions

Finally, we only need the implementation of the two interface functions get_txt and rec_terminate, see Example 1.32, "Implementation of interface functions of hello tool" (p. 47)

Example 1.32. Implementation of interface functions of hello tool

```
ATerm get_text(int conn)
{
    return
    ATmake(
        "snd-value(text(\"Hello World, my first tool in C!\n\"))"
    );
}
void rec_terminate(int conn, ATerm msg){
    exit(0);
}
```

Notes:



get_text: generate the greeting text. The conn argument identifies the ToolBus connection, making it possible to distinguish *which* ToolBus made the request. This enables connecting to more than one ToolBus at the same time.

Mandatory function that is called to terminate the tool.

Writing ToolBus tools in Java

Now we will show how ToolBus tools can be implemented in Java. The overall organization is shown in Figure 1.18, "Global organization of a tool implemented in Java(p. 47) The actual communication between ToolBus and tool is taken care of by an instance of the class ToolBridge that takes care of low-level communication details. The ToolBridge is used by AbstractTool, an abstract class that defines the possible interactions between ToolBus and tool. The actual tool, in the figure MyTool, extends AbstractTool, gives implementations for its abstract methods, and implements tool-specific behaviour.

Compared to writing a tool in C, using Java is simpler because there is no need for generating a tool interface using a tif file.²



Figure 1.18. Global organization of a tool implemented in Java

²This is due to the use of Java reflection in the class ToolBridge. In older versions, a generation step was also needed for Java tools.

Before showing the Java implementation of the hello tool, we first explain the AbstractTool class.

The AbstractTool class

public	abstract class AbstractTool implements IOperations {
public	AbstractTool(){ }
public	<pre>void connect(String[] args) throws Exception{ }</pre>
public	ToolBridge getToolBridge(){ } 3
public	PureFactory getFactory(){ }
public	void sendEvent(ATerm aTerm){ } 5
public	void disconnect(ATerm aTerm){ } 6
public	abstract void receiveAckEvent(ATerm aTerm); 🔽
public	abstract void receiveTerminate(ATerm aTerm); 🛚

Notes:

The default constructor of AbstractTool.

- Connect to the ToolBus. The argument args contains the required information for running a tool (name, id and additionally the host and port of the ToolBus, depending on how this tool is connected to the ToolBus). An exception is thrown when something goes wrong during parsing of the arguments or while establishing the connection.
- Returns a reference to the tool bridge that this tool instance is using.
- Returns a reference to the ATerm factory that is being used.
- **5** Send an event to the ToolBus.
- Send a disconnect request to the ToolBus. The argument aTerm gives additional information about the request.

Receive an acknowledgement (in response to a previous event generated by this tool instance by way of sendEvent). The argument aTerm gives further details about the acknowledgement.
 Receive a request from the ToolBus to terminate the execution of this tool instance.

Tool definition in the hello script

The hello script from Example 1.2, "hello2.tb" (p. 10) can be used as is, except that the definition of the hello tool has to be changed to reflect the Java implementation:

tool hello is { kind = "javaNG" class = "toolbus.tool.java.hello.HelloTool"}

The hello example in Java

Example 1.33. Hello tool implemented in Java

```
package toolbus.tool.java.hello;
import toolbus.adapter.AbstractTool;
import aterm.ATerm;
import aterm.ATermFactory;
public class HelloTool extends AbstractTool 🕕
 public HelloTool(String[] args){ 🛛
    super();
    try {
      connect(args);
    } catch(Exception ex){
      throw new RuntimeException(ex);
  }
 protected ATerm getText(){
   ATermFactory factory = getFactory();
   return factory.make("text(<str>)",
                             "Hello world in Java!\n");
  1
  public void receiveAckEvent(ATerm aTerm){
  // Left blank intentionally.
  1
  public void receiveTerminate(ATerm msg){
     System.out.print("rec-terminate received: " + msg);
  public static void main(String[] args) {
   new HelloTool(args);
```

Notes:

• The class HelloTool extends AbstractTool and provides an implementation for the hello tool.

The constructor for HelloTool first calls its super class and then attempt to make a connection. The arguments of the constructor are passed to the connect call.

The evaluation request snd-eval(H, get_text) in the Tscript is implemented by the method getText. It constructs the required ATerm that text("Hello world in Java!\n") and returns it as result. This result will passed to the ToolBus and will be accepted by the atom rec-value(H, text(S?)) in the Tscript. Note that in this example H is the tool identifier of the hello tool.

Writing ToolBus tools in other languages

Writing tools in Tcl/TK

Writing ToolBus tools in Tcl is greatly simplified by the **wish-adapter** to be explained in the section called "*wish-adapter*" (p. 50) Next, a small set of predefined Tcl functions is described that are always loaded by the **wish-adapter** and can be used in any Tcl script, see the section called "*Predefined Tcl functions*" (p. 50) Finally, we present the Tcl version of the hello tool in the section called "*The hello example in Tcl*" (p. 51).

wish-adapter

The purpose of the wish-adapter is to execute Tcl/Tk's windowing shell wish as a tool. For instance,

wish-adapter -script calculator.tcl

executes wish as a TooBus tool and executes the Tcl script calculator.tcl.

In addition to the common tool arguments, wish-adapter has the following specific arguments:

- -wish Name: Use Name rather than wish as Tcl/Tk's windowing shell.
- -lazy-exec: Postpone execution of wish until needed.
- -script: The Tcl script to be executed.
- -script-args: The arguments for the Tcl script to be executed. These arguments are available to the Tcl script throught the variables *argc* and *argv*.

Various communication patterns are supported by wish-adapter. Communication is described here from the point of view of the ToolBus, i.e., snd- and rec- mean, respectively, send by ToolBus and receive by ToolBus. The communication patterns are:

- snd-do(*Tid*, *Fun*(A₁, ..., A_n)) perform the Tcl function call *Fun*(A₁, ..., A_n). Here *Tid* is a tool identifier (as produced by execute or rec-connect) for an instance of the **wish-adapter**.
- snd-do(*Tid*, *Fun*(A₁, ..., A_n)): perform the Tcl function call *Fun*(A₁, ..., A_n). Here *Tid* is a tool identifier (as produced by execute or rec-connect) for an instance of the **wish-adapter**. Note that the function *Fun* must send an answer back to the ToolBus (using TBsend "snd-eval(...)").
- rec-value(*Tid*, *Res*): the return value for a previous evaluation request.
- rec-event(Tid, A_1 , ..., A_n): event generated by wish.
- snd-ack-event(*Tid*, A₁): acknowledgement of a previously generated event.
- snd-terminate(*Tid*, *A*₁): terminate execution of wish-adapter.

The command **wish** is executed once, an initial Tcl script is read, and all further requests are directed to this incarnation of **wish**. A small set of Tcl procedures is available for unpacking and packing ToolBus terms (see below).

Predefined Tcl functions

The following Tcl functions are predefined and can be used freely in Tcl script executed via the wish-adapter:

• TBstring *Str*: converts a Tcl string to a ToolBus string by surrounding it with double quotes and escaping double quotes occurring inside *Str*.

- TCLstring *Str*: converts a ToolBus string into a Tcl string by removing surrounding double quotes.
- TBlist *List*: converts a Tcl list to a ToolBus list by separating the elements with commas and surrounding the list by curly braces.
- TBerror *Msg*: constructs an error message that can be sent to the ToolBus.
- TBsend *Trm*: send *Trm* back to the ToolBus.
- TBevent *Event*: send event Event to the ToolBus.
- TBrequire *ToolName ProcName Nargs*: check that the Tcl code for *ToolName* contains a procedure declaration for *ProcName* with *Nargs* formal parameters. This function is mainly used by the **wish-adapter** to check compatibility of the Tcl code with the expected input signature of the tool.

Warning

All communication between **wish-adapter** and a tool written in Tcl is done via standard input/output. Only use the standard error stream for print statements in the Tcl script, since using standard output will disrupt the communication with the ToolBus.

The hello example in Tcl

Writing the hello tool in Tcl requires two steps:

- Write the required Tcl code hello.tcl. The result is shown in Example 1.34, "hello.tcl: the hello tool in Tcl" (p. 51).
- Replace hello's tool definition in hello2.tb (see Example 1.2, "hello2.tb" (p. 10)) by:

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

Example 1.34. hello.tcl: the hello tool in Tcl

```
# hello.tcl -- hello tool in Tcl/Tk
proc get-text {} {
   TBsend "snd-value(text(\"Hello World, my first ToolBus tool in Tcl!\n\"))"
}
proc rec-terminate { n } {
   exit
}
```

Python

A Python adapter is only available for older versions of the ToolBus. It is currently not supported.

Perl

A Perl adapter is only available for older versions of the ToolBus. It is currently not supported.

Reference Information

The syntax of Tscripts

A Tscript may contain directives like, e.g., #define, #include and #ifdef that are replaced by a preprocessor similar to the C preprocessor. We summarize the most frequently used directives:

- #define Identifier Token-sequence causes the preprocessor to replace all occurrences of Identifier by Token-sequence.
- #include "Filename" will be replaced by the entire contents of the named file.
- #ifdef and #ifndef can be used for the conditional incorporation or exclusion of parts of a script.

The syntax of Tscripts (without preprocessor directives) is as follows:

Warning

This definition is slightly out-of-date.

```
exports
 SORTS BOOL NAT INT SIGN EXP UNSIGNED-REAL REAL STRING ID
       NAME VNAME BSTR TERM TERM-LIST VAR GEN-VAR TYPE ATOM
       ATOMIC-FUN PROC PROC-APPL FORMALS TIMER-FUN
       FEATURE-ASG FEATURES TB-CONFIG DEF T-SCRIPT
 lexical syntax
        [ \t\n]
                                               -> LAYOUT
        "%%" ~[\n]*
                                              -> LAYOUT
        [0-9]+
                                              -> NAT
       NAT
                                              -> INT
       SIGN NAT
                                              -> INT
       [+\-]
                                              -> SIGN
        [eE] NAT
                                              -> EXP
        [eE] SIGN NAT
                                              -> EXP
       NAT "." NAT
                                              -> UNSIGNED-REAL
       NAT "." NAT EXP
                                              -> UNSIGNED-REAL
       UNSIGNED-REAL
                                              -> REAL
       SIGN UNSIGNED-REAL
                                              -> REAL
        [a-z][A-Za-z0-9\-]*
                                              -> ID
        "\"" ~[\"]* "\""
                                              -> STRING
        [A-Z][A-Za-z0-9\-]*
                                              -> NAME
        [A-Z][A-Za-z0-9\-]*
                                              -> VNAME
        [a-z][a-z\-]*
                                              -> ATOMIC-FUN
                                              -> TIMER-FUN
       delay
        abs-delay
                                               -> TIMER-FUN
                                               -> TIMER-FUN
        timeout
       abs-timeout
                                              -> TIMER-FUN
 context-free syntax
                                              -> BOOL
       true
        false
                                              -> BOOL
                                              -> TERM
       BOOL
                                               -> TERM
        TNT
        REAL
                                               -> TERM
       STRING
                                              -> TERM
       TERM
                                              -> TYPE
                                              -> VAR
        VNAME
        VNAME ":" TYPE
                                               -> VAR
        VAR
                                               -> GEN-VAR
```

```
VAR "?"
                                            -> GEN-VAR
      GEN-VAR
                                            -> TERM
       "<" TERM ">"
                                            -> TERM
                                            -> TERM
      ID
      ID "(" TERM-LIST ")"
                                            -> TERM
       {TERM ","}*
                                            -> TERM-LIST
       "[" TERM-LIST "]"
                                            -> TERM
      NAME
                                            -> VNAME
      ATOMIC-FUN "(" TERM-LIST ")"
                                            -> ATOM
      delta
                                            -> ATOM
       tau
                                            -> ATOM
       create "(" NAME "(" TERM-LIST ")" ","
                 TERM ")"
                                            -> ATOM
       ATOM TIMER-FUN "(" TERM ")"
                                            -> ATOM
       VNAME ":=" TERM
                                            -> ATOM
                                            -> PROC
      ATOM
      PROC "+" PROC
                                            -> PROC {left}
      PROC "." PROC
                                            -> PROC {right}
                                            -> PROC {right}
      PROC "||" PROC
      PROC "*" PROC
                                            -> PROC {left}
       "(" PROC ")"
                                            -> PROC {bracket}
       if TERM then PROC else PROC fi -> PROC
       if TERM then PROC fi
                                            -> PROC
       execute(TERM-LIST)
                                           -> PROC
      let {VAR ","}* in PROC endlet
                                            -> PROC
                                            -> PROC-APPL
      NAME
      NAME "(" TERM-LIST ")"
                                            -> PROC-APPL
      PROC-APPL
                                            -> PROC
       "(" {GEN-VAR ","}* ")"
                                            -> FORMALS
                                            -> FORMALS
      process NAME FORMALS is PROC
                                            -> DEF
      ID "=" STRING
                                            -> FEATURE-ASG
      "{" { FEATURE-ASG ";"}* "}"
tool ID FORMALS is FEATURES
                                            -> FEATURES
                                            -> DEF
       toolbus "("{PROC-APPL ","}+ ")"
                                            -> TB-CONFIG
      DEF* TB-CONFIG
                                            -> T-SCRIPT
priorities
     PROC "*" PROC -> PROC > PROC "." PROC -> PROC >
     PROC "+" PROC -> PROC > PROC " | | " PROC -> PROC
```

Built-in functions

Tscripts provide a limited form of built-in functions that are summarized here. Recall that built-in functions are only evaluated at the following syntactic positions in a Tscript:

- The right-hand side of an assignment.
- The test in an if-then or if-then-else construct.
- The expression in time-related constructs.

Boolean functions

Table 1.1. Doolean functions

Function	Result type	Description
<pre>not(<bool>1)</bool></pre>	<bool></bool>	$\neg < bool >_1$
and($_1$, $_2$)	<bool></bool>	$_1$ and $_2$
or(<bool>1, <bool>2)</bool></bool>	<bool></bool>	<bool>1 OR <bool>2</bool></bool>
equal(<term>1, <term>2)</term></term>	<bool></bool>	$< term >_1 = < term >_2$
<pre>not-equal(<term>1, <term>2)</term></term></pre>	<bool></bool>	$< term >_1 NE < term >_2$

Integer functions

Table 1.2. Integer functions

Function	Result type	Description
$add(\langle int \rangle_1, \langle int \rangle_2)$	<int></int>	$\langle int \rangle_1 + \langle int \rangle_2$
<pre>sub(<int>1, <int>2)</int></int></pre>	<int></int>	$\langle int \rangle_1 - \langle int \rangle_2$
<pre>mul(<int>1, <int>2)</int></int></pre>	<int></int>	<int>1 TIMES <int>2</int></int>
div(<int>1, <int>2)</int></int>	<int></int>	<int>1/<int>2</int></int>
$mod(\langle int \rangle_1, \langle int \rangle_2)$	<int></int>	<int>1 mod <int>2</int></int>
abs(<int>1)</int>	<int></int>	<int>1 </int>
$less(_1, _2)$	<bool></bool>	$_1 < _2$
<pre>less-equal(<int>1, <int>2)</int></int></pre>	<bool></bool>	<int>1 LEQ <int>2</int></int>
greater($_1$, $_2$)	<bool></bool>	$_1>_2$
greater-equal(<int>1, <int>2)</int></int>	<bool></bool>	<int>1 GEQ <int>2</int></int>

Real functions

Table 1.3. Real functions

Function	Result type	Description
<pre>radd(<real>1, <real>2)</real></real></pre>	<real></real>	<real>1 + <real>2</real></real>
$rsub(\langle real \rangle_1, \langle real \rangle_2)$	<real></real>	<real>1 - <real>2</real></real>
<pre>rmul(<real>1, <real>2)</real></real></pre>	<real></real>	<real>1 × <real>2</real></real>
<pre>rdiv(<real>1, <real>2)</real></real></pre>	<real></real>	<real>1/<real>2</real></real>
<pre>mod(<real>1, <real>2)</real></real></pre>	<real></real>	<real>1 mod <real>2</real></real>
rabs(<real>1)</real>	<real></real>	<pre> <real>1</real></pre>
<pre>rless(<real>1, <real>2)</real></real></pre>	<bool></bool>	<real>1 < <real>2</real></real>
rless-equal(<real>1, <real>2)</real></real>	<bool></bool>	<real>>1#<real>2</real></real>
$rgreater(_1, _2)$	<bool></bool>	<real>1><real>2</real></real>
<pre>rgreater-equal(<real>1, <real>2)</real></real></pre>	<bool></bool>	<real>>1#<real>2</real></real>

Goniometric functions

Table 1.4. Goniometric functions

Function	Result type	Description
<pre>sin(<real>1)</real></pre>	<real></real>	<pre>sin(<real>1)</real></pre>
cos(<real>1)</real>	<real></real>	<pre>cos(<real>1)</real></pre>
atan(<real>1)</real>	<real></real>	<i>tan</i> ⁻¹ (<real>₁) in the range [- #/2, #/2]</real>
<pre>atan2(<real>1, <real>2)</real></real></pre>	<real></real>	$tan^{-1}(_1/_2)$ in the range [-#, #]
exp(<real>1)</real>	<real></real>	$e^{<\operatorname{real}>_1}$
log(<real>1)</real>	<real></real>	natural logarithm $ln(_1)$, with $_1 > 0$
<pre>log10(<real>1)</real></pre>	<real></real>	base 10 logarithm $log_{10}(_1)$, with $_1 > 0$
<pre>sqrt(<real>1)</real></pre>	<real></real>	# <real>1, with <real>1 $# 0$</real></real>

Functions on lists

Table 1.5. Functions on lists

Function	Result type	Description
<pre>first(<list>1)</list></pre>	<term></term>	First element of <list>1; The empty list [] when applied to non-list or empty list.</list>
<pre>next(<list>1)</list></pre>	<list></list>	Remaining elements of <list>1.</list>
<pre>join(<term>1, <term>2)</term></term></pre>	<list></list>	Concatenation of <term>1 and <term>2. When both arguments are lists their elements are spliced into a new list. A non-list argument is included as single element in the new list.</term></term>
<pre>size(<list>1)</list></pre>	<int></int>	The number of elements in <list>1.</list>

Table 1.6. Functions on lists as arrays

Function	Result type	Description
<pre>index(<list>1, <int>1)</int></list></pre>	<term></term>	The <int>1-th element of <list>1, if it exists; otherwise [].</list></int>
<pre>replace(<list>1,<int1>,<term>1)</term></int1></list></pre>	<list></list>	If the <int>1-the element exists, replace it by <term>1 and returned the modified list; otherwise return <list>1 unmodified.</list></term></int>

Function	Result type	Description
<pre>get(<list>1, <term>1)</term></list></pre>	<term></term>	<pre>If <list>1contains a pair [<term>1, <term>1'] then return <term>1'; otherwise [].</term></term></term></list></pre>
<pre>put(<list>1, <term>1, <term>2)</term></term></list></pre>	<list></list>	<pre>If <list>1contains a pair [<term>1, <term>1'] then replace it by [<term>1, <term>2]; otherwise add a new pair [<term>1, <term>2] to <list>1.</list></term></term></term></term></term></term></list></pre>

Table 1.7. Functions on lists as symbol tables

Table 1.8. Functions on lists as multi-sets

Function	Result type	Description
<pre>member(<term>1, <list>1)</list></term></pre>	<bool></bool>	<term>1 IN <list>2 (membership in multi-set)</list></term>
<pre>subset(<list>1,<list>2)</list></list></pre>	<bool></bool>	list>1 SUBSET <list>2 (subset on multi-set)</list>
<pre>diff(<list>1,<list>2)</list></list></pre>	<list></list>	list>1 DIFF <list>2 (difference on multi-set)</list>
<pre>inter(<list>1,<list>2)</list></list></pre>	<list></list>	list>₁ INTER <list>₂ (intersection on multi-set)</list>

Functions on terms

Table 1.9	. Functions	on ATerms
-----------	-------------	-----------

Function	Result type	Description
<pre>is-bool(<term>)</term></pre>	<bool></bool>	If <term> is of type bool then true; otherwise false.</term>
<pre>is-int(<term>)</term></pre>	<bool></bool>	If <term> is of type int then true; otherwise false.</term>
<pre>is-real(<term>)</term></pre>	<bool></bool>	If <term> is of type real then true; otherwise false.</term>
is-str(<term>)</term>	<bool></bool>	If <term> is of type str then true; otherwise false.</term>
<pre>is-bstr(<term>)</term></pre>	<bool></bool>	If <term> is of type bstr then true; otherwise false.</term>
<pre>is-appl(<term>)</term></pre>	<bool></bool>	If <term> is an application then true; otherwise false.</term>
<pre>is-list(<term>)</term></pre>	<bool></bool>	If <term> is a list then true; otherwise false.</term>
<pre>is-empty(<term>)</term></pre>	<bool></bool>	If <term> is equal to [] then true; otherwise false.</term>
<pre>is-var(<term>)</term></pre>	<bool></bool>	If <term> is a variable then true; otherwise false.</term>
<pre>is-var(<term>)</term></pre>	<bool></bool>	If <term> is a variable then true; otherwise false.</term>
<pre>is-result-var(<term>)</term></pre>	<bool></bool>	If <term> is a result variable then true; otherwise false.</term>
<pre>is-formal(<term>)</term></pre>	<bool></bool>	If <term> is a formal variable then true; otherwise false.</term>
<pre>fun(<term>)</term></pre>	<str></str>	If <term> is a function application then its function symbol; otherwise " ".</term>
args(<term>)</term>	<list></list>	If <term> is a function application then its argument; otherwise [].</term>

Time-related functions

Table 1.10. Time-related functions

Function	Result type	Description
current-time	<list></list>	Six-tuple describing the current absolute time
sec(<int>)</int>	<int></int>	Convert <int> into seconds</int>

Miscellaneous functions

Table 1.11. Miscellaneous functions

Function	Result type	Description
process-id	<int></int>	Process id of the current process
process-name	<str></str>	Name of the current process
<pre>quote(<term>)</term></pre>	<term></term>	Quoted (unevaluated) term; only variables are replaced by their values
functions	<list></list>	List of all built-in functions

Synopsis of ToolBus primitives

In the following two sections all primitives are summarized that can occur in a Tscript.

Process-related primitives in Tscripts

Primitive	Synopsis	See
delta	Inaction (deadlock)	
tau	Internal step	
$P_1 + P_2$	Choice between P_1 and P_2	-
P ₁ . P ₂	P_1 followed by P_2	
$P_1 \parallel P_2$	P_1 parallel with P_2	
$P_1 * P_2$	Repeat P_1 until P_2	
if T then P fi	Guarded command	
if T then P_1 else P_2 fi	Conditional	
create(<i>Pnm</i> (<i>T</i> ,), <i>Pid</i> ?)	Create new process	
V := T	Assign T (seen as expression) to V	
<pre>snd-msg(T)</pre>	Send synchronous message	
rec-msg(T)	Receive a synchronous message	
<pre>snd-note(T)</pre>	Broadcast an asynchronous note	
rec-note(T)	Receive an asynchronous note	
no-note(T)	No note available	
subscribe(T)	Subscribe to notes	
A delay(T)	Relative delay of atom execution	
A abs-delay(T)	Absolute delay of atom execution	
A timeout(T)	Relative timeout of atom execution	
A abs-timeout(T)	Absolute timeout of atom execution	
shutdown(T)	Terminate ToolBus application	
printf(<i>S</i> , <i>T</i> ,)	Print terms according to format string S	
read(T_1 , T_2)	Give prompt T_1 and read term that should match with T_2	
process Pnm(F,) is P	Define process Prim	
let F, in P endlet	Declare local variables in P	
ToolBus(<i>Pnm</i> (<i>T</i> ,),)	Define initial ToolBus process configuration	

Table 1.12. Process-related primitives in Tscripts

Tool-related primitives in Tscripts

Primitive	Synopsis	See
rec-connect(Tid?)	Receive connection request from tool	
rec-disconnect(Tid?)	Receive disconnection request from tool	
<pre>execute(Tnm(T,), Tid?)</pre>	Execute a tool	
<pre>snd-terminate(Tid, T)</pre>	Terminate execution of a tool	
<pre>snd-eval(Tid, T)</pre>	Send evaluation request to tool	
<pre>snd-cancel(Tid)</pre>	Cancel previous evaluation request	
rec-value(<i>Tid</i> , <i>T</i>)	Receive answer to evaluation request	
snd-do(<i>Tid</i> , <i>T</i>)	Send evaluation request to tool (no return value)	
<pre>rec-event(Tid, T,)</pre>	Receive event from tool	
<pre>snd-ack-event(Tid, T)</pre>	Acknowledge previous event from tool	
<pre>tool Tnm is { Feat, }</pre>	Define tool Tram	
host = <i>Str</i>	Host feature in tool definition	
command = Str	Command feature in tool definition	

Table 1.13. Tool-related primitives in Tscripts

Historical notes

The first generation ToolBus is described in [BK94] (p. 60) In addition to the design, the complete C implementation is discussed in detail. The second generation ("discrete time") ToolBus includes timing primitives as well as built-in functions. It has been formally described using ASF+SDF, see [BK95] (p. 60)and [BK98] (p. 60) In [Oli00] (p. 61)a framework for the debugging of ToolBus applications is presented. Initial thoughts about a next generation ToolBus were published in [dJK03] (p. 60). [dJ07] (p. 60) describes architectural aspects of ToolBus-based applications.

Warning

Add: theses of Peter Heibrink, Arnold Lankamp, Dennis Hendriks.

Bibliography

- [BK94] J.A. Bergstra and P. Klint. *The toolbus: a component interconnection architecture*. Technical ReportP9408. University of Amsterdam, Programming Research Group. 1994.
- [BK95] J.A. Bergstra and P. Klint. *The discrete time toolbus*. Technical ReportP9502. University of Amsterdam, Programming Research Group. 1995.
- [BK98] J.A. Bergstra and P. Klint. *The discrete time ToolBus -- a software coordination architecture*. 205--229. *Science of Computer Programming*. 31. 2-3. July 1998.
- [dJK03] H.A. de Jong and P. Klint. *Toolbus: the next generation*. 220--241. Formal Methods for Components and Objects. F.S. de Boer, M. Bonsangue, S. Graf, and W.P de Roever. Lecture Notes in Computer Science. 2852</seriesvolnum>2003">seriesvolnum>2852</seriesvolnum>2003. Springer.
- [dJ07] H.A. de Jong. Flexible Heterogeneous Software Systems. PhD thesis. University of Amsterdam. 2007.

[Oli00] P.A. Olivier. A Framework for Debugging Heterogeneous Applications. PhD thesis. University of Amsterdam. 2000.

To Do

- What do we do with the other adapters?
- Describe current viewer.
- Describe console commands.
- Do we describe the global structure of the Java implementation (or partially refer to online docs)?
- Describe viewer interface.