


Assured Reconfiguration of Embedded Real-Time Software


By E. Strunk and J. Knight



Cut and paste job presented by Foo
without permission of authors

Claims

- ? They present a comprehensive approach to assured reconfiguration
- ? Provide a framework for formal verification that allows system developer to use a set of application-level properties to show general reconfiguration properties
- ? State their approach allows the system design to be much simpler and easier to analyze



What *you* should learn from this paper?


- ? When designing a system, you will have high-level properties/goals you want to satisfy
- ? It may be difficult to prove them formally in your design
- ? Instead, use low level properties that can be easier to guarantee and use these properties to prove the high-level ones
- ? You must be able to formally describe these properties in a language of your choosing
- ? In essence, that is what they did here, so this paper is a good example of how to do that

Definitions

- ? Specified set S of functionality levels
- ? Reconfiguration defines transitions between pairs of members of S
- ? Informally, *the process through which a system halts operation under its current source specification S_i and begins operation under a different target specification S_j*


Terms used

- S : the set $\{S_1, S_2, \dots, S_n\}$ of service specifications of the system
- E : a set of possible functions from factors that affect which specification is an appropriate instantiation of S_j to possible values of those factors
- $\text{Env}(t) \rightarrow e \in E$: function that returns the value of the environmental state at time t
- $\text{Choose}(S_x, e \in E) \rightarrow S_y$: function that returns appropriate target elements of S , given source elements of S and E .



More terms used

- Pre_x : the precondition that must be satisfied for the system to operate under S_x
- Inv_{ij} : an invariant that must hold during the transition from S_i to S_j
- $\uparrow S_x$: the event marking the beginning of the system's operation under specification S_x . This event occurs when the system first operates according to the function set out in S_x , and concurrently satisfies Pre_x .
- $\downarrow S_x$: the event marking the end of the system's operation under specification S_x (we define the occurrence of a reconfiguration signal to imply $\downarrow S_x$)
- T_{ij} : the maximum allowable time between $\downarrow S_i$ and $\uparrow S_j$




Conditions that must hold for reconfiguration R

P1. $@(\uparrow R, b) = @(\downarrow S_i, a)$

(R begins at the same time the system is no longer operating under S_i)

Note that this property defines the beginning of **R** rather than specifying a requirement that should be met.



P2. $@(\downarrow R, b) = @(\uparrow S_j, c)$

(R ends at the same time the system becomes compliant with S_j)

P3. $\exists t: \text{time s.t. } @(\uparrow R, b) \leq t \leq @(\downarrow R, b) \wedge$
 $(\text{Choose}(S_i, \text{Env}(t)) = S_j) < t, t >$

*(S_j is the proper choice for the target specification at some point during **R**)*

Continued

P4. $@(\downarrow R, b) - @(\uparrow R, b) \leq T_{ij}$

(R takes less than or equal to T_{ij} time units)

P5. $\text{Inv}_{ij} < @(\uparrow R, b), @(\downarrow R, b) >$

*(The transition invariant holds during **R**)*

P6. $\text{Pre}_j < @(\downarrow R, b), @(\downarrow R, b) >$

*(The precondition for S_j is true at the time **R** ends)*

P7. $@(\downarrow S_j, c-1) < @(\uparrow R, b) \Rightarrow @(\uparrow S_j, c) = @(\downarrow R, b)$

*(The lifetime of **R** is bounded by any two occurrences of the same specification)*

Architecture for Assurance

? Claim: If system implementer builds system using this architecture and shows the low-level properties required of the architectural elements, he will know the high-level properties that assure reconfiguration have been met

? 2 major components

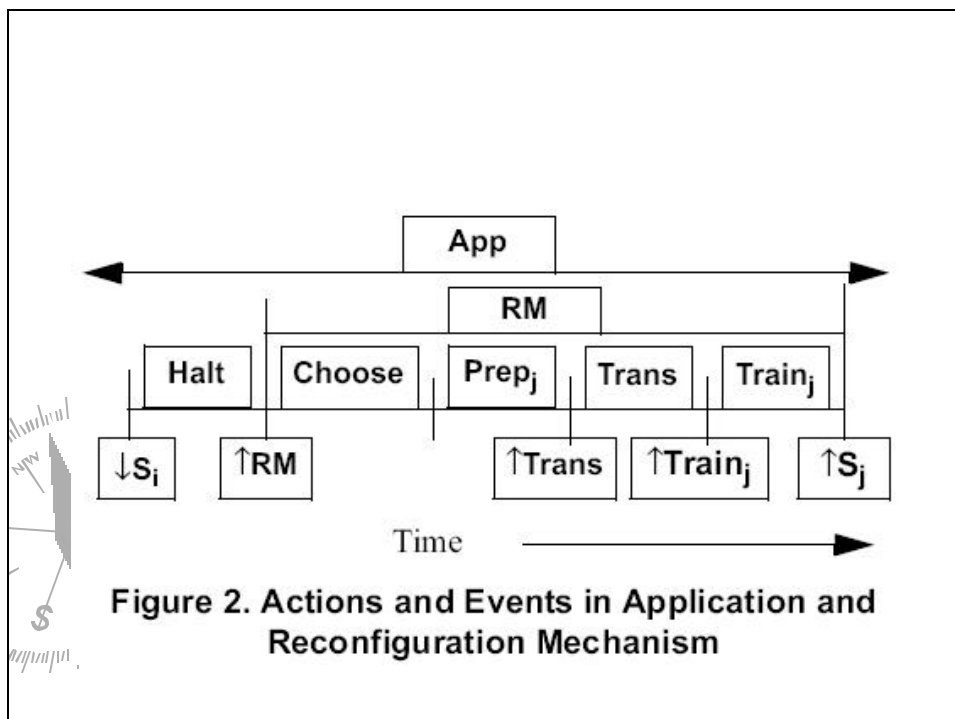
- *Application* that performs computations associated with members of S
- *Reconfiguration mechanism* that ensures reconfiguration can be carried out

Building blocks (events, actions,...)

App	Action representing operation of the application. Operates continuously throughout R .
RM	Action representing operation of the reconfiguration mechanism.
$\uparrow RM$	Event that App signals fault at top level
$\downarrow RM$	$\equiv \downarrow Trans$ (below)
$\downarrow S_i$	Event that a reconfiguration signal is sent to S_i
Halt	Action of App causing S_i to meet Post
$\uparrow Halt$	$\equiv \downarrow S_i$
Post	Predicate that must be true of App in order for reconfiguration to begin. This condition protects data, and ensures that the software is fail-stop.
Choose	Action of RM determining member of S for S_j

More building blocks

- Pre_{transj}** Predicate that must be true of **App** before passing control back to **App**
- Pre_j** Action of **App** causing **Pre_{transj}** to be met
- Trans** Action of **App** effecting functional transition
- ↑Trans** Event that **RM** instructs **App** to transition to S_j
- ↓Trans** Event that **App** acknowledges to **RM** that transition is complete
- Train_j** Action of **App** initializing S_j or training its data
- ↑Train_j** \equiv **↓Trans**
- ↓Train_j** \equiv **↑S_j**
- ↑S_j** S_j has now initialized or trained all data and is operational.



Application

- ? Application consist of modules, each with interface to support reconfiguration assurance
- ? Properties of the modules compose properties of app
- ? Each interface function takes a module-specific *service level parameter* that instructs interface to provide a level of function
- ? Composition of different module service levels allow system to operate under different specifications

- ? Now we state the necessary design-level properties that must be shown, in order to meet the required high-level reconfiguration properties

Module capabilities


- ? Interface to the set of functions contained within the module
- ? Set of possible values for the service level parameter
- ? Set of persistent data structures
 - Data which is relevant to preconditions, postconditions, and invariants
- ? Module postcondition
 - Basic coherency condition representing min state requirement for application to continue some form of operation
- ? Mechanism through which reconfiguration signals are handled
- ? Mechanism through which reconfiguration signals are propagated to calling functions
- ? Set of module transition conditions
- ? Set of module preconditions
- ? Mechanism through which a module's service transition condition is guaranteed to be met
- ? Timing guarantees on interface functions related to reconfiguration
- ? Set of assured reconfiguration invariants

Module properties shown using capabilities

M1	If none of a module's functions is currently executing, that module's postcondition is met.
M2	Each module has a function halt that, when called, will cause its postcondition to be met.
M3	Each module function either: always leaves its data in a consistent state; or when interrupted, calls the module halt function to leave the state consistent with the module postcondition.
M4	If a function is interrupted, its caller is interrupted with no intervening calls to any function other than halt .
M5	There is a method to meet the transition condition for each specification level.
M6	Each function always meets its timing constraint.
M7	The invariant for normal operation is stricter than the generic reconfiguration invariant, which is stricter than the specific reconfiguration invariant: $Inv_1 \Rightarrow Inv_R \Rightarrow Inv_i$

Application capabilities

- ? Modules that compose the system are contained within a separate top-level structure called *monitoring layer*
- ? Monitoring layer includes
 - Facility to activate reconfiguration mechanism
 - State variable **config** representing current operation specification
 - Capability to cause operation under current specification
 - Max time **train_time** that training of data might take for each member of S




Application properties

<i>App1</i>	If App is not reconfiguring, it will function in accordance with the specification represented by config 's value.
<i>App2</i>	Every operation is called from some sequence of functions initiated by the monitoring layer.
<i>App3</i>	The postcondition is the conjunction of module postconditions.
<i>App4</i>	The transition condition is the conjunction of module transition conditions.
<i>App5</i>	An interrupt of the monitoring layer will cause an immediate transfer of control to reconfig, which is the functional equivalent of the action RM : $@(\downarrow \text{Halt}, a) = @(\uparrow \text{RM}, b)$
<i>App6</i>	config 's value is invariant outside of RM : $\text{config} = x \wedge @(\downarrow \text{RM}, b), @(\downarrow \text{RM}, b) \Rightarrow \text{config} = x \wedge @(\downarrow \text{RM}, b), @(\uparrow \text{RM}, b+1) \Rightarrow$
<i>App7</i>	There are no circular dependencies among module reinitialization functions.
<i>App8</i>	If the transition precondition holds at the time the transition is completed, Pre_j will be met within train_time(j) time units: $\text{Pre}_{\text{Trans}(j)} \wedge @(\uparrow \text{Train}_j, b), @(\uparrow \text{Train}_j, b) \Rightarrow @(\downarrow \text{Train}_j, b) \leq @(\uparrow \text{Train}_j, b) + \text{train_time}(j)$
<i>App9</i>	The transition takes no real time: $@(\uparrow \text{Trans}, b) = @(\downarrow \text{Trans}, b)$ This property is true of the structure rather than a particular application, and can be stated as an axiom.




Reconfiguration Mechanism capabilities

- ? Implementation of Choose()
- ? Mechanism through which each module is ordered to meet its precondition for the new service level



Reconfiguration Mechanism properties

RM1	choose will be executed immediately when RM is called. choose is equivalent to the action Choose. This means that $@(\uparrow RM, b) = @(\uparrow Choose, b)$.
RM2	If the postcondition of choose is met, the new operational specification is the correct one and is stored in config.
RM3	If config's value is invariant outside of RM, then config's value is invariant outside of Choose: $App6 \Rightarrow \{ config = x < @(\downarrow Choose, b), @(\downarrow Choose, b) > \Rightarrow config = x < @(\downarrow Choose, b), @(\uparrow Choose, b+1) > \}$
RM4	$App7 \Rightarrow$ RM calls all the prep functions of the modules in an order in which no dependencies are violated.
RM5	Exactly the prep functions are called between choose and transition to the monitoring layer; this implies: $@(\downarrow Choose, b) = @(\uparrow Prep, c) \leq @(\downarrow Prep, c) = @(\uparrow Trans, b)$ and $(Post < @(\uparrow Prep, c), @(\uparrow Prep, c) > \wedge RM4 \wedge App4) \Rightarrow Pre_{trans} < @(\downarrow Prep, c), @(\downarrow Prep, c) >$
RM6	If each function meets its timing constraint, then App can halt, RM can execute, and App can train within the allotted time: $M6 \Rightarrow @(\downarrow Halt, a) - @(\uparrow Halt, a) + @(\downarrow RM, b) - @(\uparrow RM, b) + @(\downarrow Train, c) - @(\uparrow Train, c) \leq T_i$
RM7	The invariants that must hold during transition hold at the appropriate times: $Inv_i < @(\uparrow Halt, a), @(\uparrow Halt, a) > \wedge Inv_{ix} < @(\uparrow Halt, a), @(\downarrow Choose, b) > \wedge Inv_{ij} < @(\downarrow Choose, b), @(\uparrow S_j, c) >$
RM8	RM begins before transition and ends at the time of transition; training begins at the time of transition: $@(\uparrow RM, b) \leq @(\uparrow Trans, b) = @(\downarrow RM, b) = @(\uparrow Train, b)$



? Now they show, using design-level properties previously stated, that a system using their architecture and satisfying architectural properties will cause high-level properties to hold



P1. $@(\uparrow R, b) = @(\downarrow S_i, a)$

(R begins at the same time the system is no longer operating under S_i)

P1 is definitional only, and does not impose specific requirements.

P2. $@(\downarrow R, b) = @(\uparrow S_j, c)$


(R ends at the same time the system becomes compliant with S_j)

This property is definitional and so by itself requires no proof. It implies that the system must at some point transition to S_j , but in our model this property is subsumed by P6 since Pre_j can in general be satisfied only after a functional reconfiguration takes place.

P7. $@(\downarrow S_j, c-1) < @(\uparrow R, b) \Rightarrow @(\uparrow S_j, c) = @(\downarrow R, b)$

(The lifetime of R is bounded by any two occurrences of the same specification)

P7 is definitional only, and does not impose specific requirements.



P4. $@(\downarrow R, b) - @(\uparrow R, b) \leq T_{ij}$


(*R* takes less than or equal to T_{ij} time units)

This property can be shown using P1, P2, the definitions of Halt and Train, App5, RM8, and RM6.

P5. $\text{Inv}_{ij} < @(\uparrow R, b), @(\downarrow R, b) >$

(*The transition invariant holds during R*)

This property can be shown using P1, P2, M7, and RM7.



P4

P1. $@(\uparrow R, b) = @(\downarrow S_i, a)$

(*R* begins at the same time the system is no longer operating under S_i)


P2. $@(\downarrow R, b) = @(\uparrow S_j, c)$

(*R* ends at the same time the system becomes compliant with S_j)

App5 | An interrupt of the monitoring layer will cause an immediate transfer of control to reconfig, which is the functional equivalent of the action **RM**: $@(\downarrow \text{Halt}, a) = @(\uparrow \text{RM}, b)$

RM6 | If each function meets its timing constraint, then App can halt, RM can execute, and App can train within the allotted time: $M6 \Rightarrow @(\downarrow \text{Halt}, a) + @(\uparrow \text{Halt}, a) + @(\downarrow \text{RM}, b) - @(\uparrow \text{RM}, b) + @(\downarrow \text{Train}_j, c) - @(\uparrow \text{Train}_j, c) \leq T_{ij}$

RM8 | *RM* begins before transition and ends at the time of transition; training begins at the time of transition:
 $@(\uparrow \text{RM}, b) \leq @(\uparrow \text{Trans}, b) = @(\downarrow \text{RM}, b) = @(\uparrow \text{Train}_j, b)$



$\downarrow S_i$ | Event that a reconfiguration signal is sent to S_i

Halt | Action of **App** causing S_i to meet **Post**

$\uparrow \text{Halt} \equiv \downarrow S_i$

Trans | Action of **App** effecting functional transition

$\uparrow \text{Trans}$ | Event that **RM** instructs **App** to transition to S_j

$\downarrow \text{Trans}$ | Event that **App** acknowledges to **RM** that transition is complete

Train_j | Action of **App** initializing S_j or training its data


$\uparrow \text{Train}_j \equiv \downarrow \text{Trans}$

$\downarrow \text{Train}_j \equiv \uparrow S_j$

P5

M7 | The invariant for normal operation is stricter than the generic reconfiguration invariant, which is stricter than the specific reconfiguration invariant: $Inv_i \Rightarrow Inv_{ix} \Rightarrow Inv_j$

RM7 | The invariants that must hold during transition hold at the appropriate times:
 $Inv_i < @(\uparrow Halt, a), @(\uparrow Halt, a) > \wedge Inv_{ix} < @(\uparrow Halt, a), @(\downarrow Choose, b) > \wedge Inv_j < @(\downarrow Choose, b), @(\uparrow S_j, c) >$



P6. $Pre_j < @(\downarrow R, b), @(\downarrow R, b) >$

(The precondition for S_j is true at the time R ends)


The proof of P6 is more complex because it requires that a series of predicates be satisfied. The proof is aided by using the following lemmas:

L6.1. An interruption will cause the application to meet its postcondition: $\exists t: \text{time s.t. } @(\downarrow S_i, a) = t \Rightarrow \exists t: \text{time s.t. } Post < t, t >$

This lemma can be shown through induction using App2, App3, M1, M2, M3, and M4.

L6.2. An interruption will cause control to be transferred back to the monitoring layer: $\exists t: \text{time s.t. } @(\downarrow S_i, a) = t \Rightarrow \exists t: \text{time s.t. } t = @(\downarrow Halt, a)$.


This lemma can be shown through induction using App2 and M4.



Together with a second application of M1, these lemmas imply that an interruption will cause control to be transferred to the monitoring layer at the same time the postcondition is met. This being true,

L6.3. $\text{Post} < @(\uparrow \text{RM}, b), @(\uparrow \text{Pre}_j, c) >$

which follows from App5, RM1, RM5, and M1. Using RM5 again, then App9 and App8, we see that at some time t between $@(\downarrow \text{Trans}, b)$ and $@(\downarrow \text{Trans}, b) + \text{train_time}$, Pre_j is satisfied. Because RM2 and RM3 tell us that config holds the correct value, and RM5 and App9 tell us that at $\downarrow \text{Trans}$ control is passed back to the monitoring layer, and using App1 this means that the system is in functional compliance with S_j , we know that **App** is operating according to S_j , so $t = \uparrow S_j$; and using P2, P6 is shown.



P3. $\exists t: \text{time s.t. } @(\uparrow R, b) \leq t \leq @(\downarrow R, b) \wedge$
 $(\text{Choose}(S_i, \text{Env}(t)) = S_j) < t, t >$

(S_j is the proper choice for the target specification at some point during R)

We present the full proof of P3 to give the reader an example of their construction, and then outline subsequent proofs so that the reader can construct the full chain of reasoning for himself.

For brevity, we write $(\text{Choose}(S_i, \text{Env}(t)) = S_j) < t, t >$ as the predicate $\text{valid}_j(t)$. We first establish that $\text{valid}_j(t)$ is true for the time Choose ends: $\text{valid}_j(@(\downarrow \text{Choose}, b))$. RM1 says that


$\text{choose} \equiv \text{Choose}$

\Rightarrow (by the relationship we have assigned functional and sequence properties)

$\text{choose.post} < @(\downarrow \text{Choose}, b), @(\downarrow \text{Choose}, b) >$

\Rightarrow (RM2)

$(\text{Choose}(S_i, \text{Env}(t)) = S(\text{config})) < @(\downarrow \text{Choose}, b), @(\downarrow \text{Choose}, b) >$



App6 and RM3 tell us that $(\text{config} = j) < @(\downarrow\text{Choose}, b)$,
 $@(\downarrow\text{Choose}, b) >$ since S_{config} will be the new operational
specification and config will remain constant until $\uparrow S_j$.
Together with the above statement we see

$\text{validj}(@(\downarrow\text{Choose}, b))$.

Intuitively, it is obvious that $@(\uparrow R, b) \leq @(\downarrow\text{Choose}, b) \leq$
 $@(\downarrow R, b)$. More rigorously:

$\exists t: \text{time s.t. } @(\downarrow\text{Choose}, b) = t \wedge \text{validj}(t)$

\Rightarrow (axiomatic basis of time in RTL)

$\exists t: \text{time s.t. } @(\uparrow\text{Choose}, b) \leq t \leq @(\downarrow\text{Choose}, b) \wedge$
 $\text{validj}(t)$

\Leftrightarrow (RM1 and RM5)

$\exists t: \text{time s.t. } @(\uparrow R, b) \leq t \leq @(\uparrow\text{Trans}, b) \wedge \text{validj}(t)$

\Leftrightarrow (App5 and RM8)

$\exists t: \text{time s.t. } @(\downarrow\text{Halt}, a) \leq t \leq @(\uparrow\text{Train}_i, c) \wedge \text{validj}(t)$

\Rightarrow (axiomatic basis of time in RTL)


$\exists t: \text{time s.t. } @(\uparrow\text{Halt}, a) \leq t \leq @(\downarrow\text{Train}_i, c) \wedge \text{validj}(t)$

\Rightarrow (definitions of Halt and Train)

$\exists t: \text{time s.t. } @(\downarrow S_i, a) \leq t \leq @(\uparrow S_j, c) \wedge \text{validj}(t)$

\Rightarrow (P1 and P2, substituting for $\text{validj}(t)$)

$\exists t: \text{time s.t. } @(\uparrow R, b) \leq t \leq @(\downarrow R, b) \wedge$
 $(\text{Choose}(S_i, \text{Env}(t)) = S_j) < t, t >$




Example: RIPS system

? NASA's Runway Incursion Prevention
System

? Runs on aircraft to prevent collisions with
objects on runway – gives advice to pilots

? Specific to RIPS, the Runway Safety Monitor
(RSM) algorithm



Service specifications

? Their model of the algo contains 4 major operational specifications

- S1 – monitors runway and surroundings
- S2 – monitors runway (easier than S1)
- S3 – halts and alerts pilots
- S4 – gives aircraft command to climb and alerts pilot



Main modules

? GEOM – computes basic geometric functions


? IZ – sets up geometry specific to RSM

? ALG – analyzes incoming data

? GEOM is protected by a layer that checks its outputs to give them strong correctness arguments

? Two sets of persistent data

- Incursion zone structure (belongs to IZ)
- Data structures for system interface (ALG)



- ? In order to indicate how all of the high-level reconfiguration properties can be proved, they choose 3 representative design-level properties.
- ? The representative module property is M1: If none of a module's functions is currently executing, that module's postcondition is met
- ? Disallow data structure access through any function outside the module interface
- ? Any failure of a check causes reconfiguration signal to propagate through modules
- ? Modules will cause data structures to meet their postconditions