

NASA Technical Memorandum 104196

IN-62
83750
p-161

Formal Design and Verification of a Reliable Computing Platform For Real-Time Control

Phase 2 Results

Ricky W. Butler
Ben L. Di Vito

(NASA-TM-104196) FORMAL DESIGN AND
VERIFICATION OF A RELIABLE COMPUTING
PLATFORM FOR REAL-TIME CONTROL. PHASE 2:
RESULTS (NASA) 161 p CSCL 09B

N92-22320

Unclas
G3/62 0083750

January 1992

NASA

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23665-5225

1

Contents

1	Introduction	1
1.1	Design of the Reliable Computing Platform	2
1.2	Overview of Results	4
1.3	Previous Efforts	6
2	Specification Hierarchy and Verification Approach	6
2.1	The State Machine Approach to Specification	6
2.2	Specifying Behavior in the Presence Of Faults	7
2.3	The Specification Hierarchy	8
2.4	Extended State Machine Model	10
2.5	The Proof Method	11
3	US/RS Specification	13
3.1	Preliminary Definitions	13
3.2	US Specification	14
3.3	RS Specification	14
3.4	Actuator Outputs	17
3.5	Generic Fault-Tolerant Computing	18
3.5.1	State Model for Transient Fault Recovery	18
3.5.2	Transient Recovery Axioms	19
3.5.3	Sample Interpretations of Theory	21
4	RS to US Proof	22
5	DS Specification	24
6	DS to RS Proof	28
6.1	DS to RS Mapping	29
6.2	The Proof	30
7	DA Specification	32
7.1	Clock Synchronization Theory	33
7.2	The DA Formalization	36
8	DA to DS Proof	42
8.1	DA to DS Mapping	42
8.2	The Proof	43
8.2.1	Decomposition Scheme	43
8.2.2	Proof of com_broadcast_2	44
9	Implementation Considerations	48
9.1	Restrictions Imposed by the DA Model	48
9.2	Processor Scheduling	49
9.3	Hardware Protection Features	50

9.4 Voting Mechanisms	51
10 Future Work	52
10.1 Further Refinement	52
10.2 Task Scheduling and Voting	54
10.3 Actuator Outputs	54
10.4 Development of a Detailed Reliability Model	54
11 Concluding Remarks	54
A Appendix — L^AT_EX-printed Specification Listings	62
B Appendix — L^AT_EX-printed Supplementary Specification Listings	108
C Appendix — Results of Proof Chain Analysis	115

1 Introduction

NASA is engaged in a major research effort towards the development of a practical validation and verification methodology for digital fly-by-wire control systems.¹ Researchers at NASA Langley Research Center (LaRC) are exploring formal verification as a candidate technology for the elimination of design errors in such systems. In previous reports [1, 2, 3], we put forward a high level architecture for a *reliable computing platform* (RCP) based on fault-tolerant computing principles. Central to this work is the use of formal methods for the verification of a fault-tolerant operating system that schedules and executes the application tasks of a digital flight control system. Phase 1 of this effort established results about the high level design of RCP. This report presents our Phase 2 results, which carry the design, specification, and verification of RCP to lower levels of abstraction.

The major goal of this work is to produce a verified real-time computing platform, both hardware and operating system software, which is useful for a wide variety of control-system applications. Toward this goal, the operating system provides a user interface that “hides” the implementation details of the system such as the redundant processors, voting, clock synchronization, etc. We adopt a very abstract model of real-time computation, introduce three levels of decomposition of the model towards a physical realization, and rigorously prove that the decomposition correctly implements the model. Specifications and proofs have been mechanized using the EHDM verification system [4].

A major goal of the RCP design is to enable the system to recover from the effects of transient faults. More than their analog predecessors, digital flight control systems are vulnerable to external phenomena that can temporarily affect the system without permanently damaging the physical hardware. External phenomena such as electromagnetic interference (EMI) can flip the bits in a processor’s memory or temporarily affect an ALU. EMI can come from many sources such as cosmic radiation, lightning or High Intensity Radiated Fields (HIRF). There is growing concern over the effects of HIRF on flight control systems. In the FAA Digital Systems Validation Handbook - volume II [5], we find:

A number of European military aircraft fatal accidents have been attributed to High Energy Radio Frequency (HERF).² A digital fly-by-wire military Tornado aircraft and crew were lost during a tactical training strafing attack in Germany. The loss was attributed to HERF when the aircraft flew through a high intensity Radio Frequency (RF) field. The civil/military aviation industry has very limited experience or data directed to accidents caused by electromagnetic transients and/or radiation. The present criteria, specifications, and procedures are being reevaluated. The HERF fields apparently upset the digital flight control system of the Tornado which was qualified to a very low electromagnetic Environment (EME) standard.

While composite materials may offer significant advantages in strength, weight, and cost, they provide less electromagnetic shielding than aluminum. The use

¹In fly-by-wire aircraft the direct mechanical and hydraulic linkages between the pilot and actuators of the system are replaced with digital computers. These digital computers are being used to control life critical functions such as the engines, sensors, fuel systems and actuators.

²The term HERF has largely been replaced in current usage by the newer term HIRF.

of solid-state digital technology in flight-critical systems create major challenges to prevent transient susceptibility and upset in both civil and military aircraft. Therefore, the Civil Aviation Authority (CAA), United Kingdom (U.K.) and the Federal Aviation Administration (FAA), United States (U.S.) voiced concern relative to emerging technology aircraft and systems.

The RCP system is designed to automatically flush the effects of transients periodically, as long as the effect of a transient is not massive, that is, simultaneously affecting a majority of the redundant processors in the system.³ Of course, there is no hope of recovery if the system designed to overcome transient faults contains a design flaw. Consequently, a major emphasis in this work has been the development of techniques that mathematically show when the desired recovery properties have been achieved. The advantages of this approach are significant:

- Confidence in the system does not rely primarily on end-to-end testing, which can never establish the absence of some rare design flaw (yet more frequent than 10^{-9} [6]) that can crash the system [7].
- Minimizes the need for experimental analysis of the effects of EMI or HIRF on a digital processor. The probability of occurrence of a transient fault must be experimentally determined, but it is not necessary to obtain detailed information about how a transient fault propagates errors in a digital processor.
- The role of experimentation is determined by the assumptions of the mathematical verification. The testing of the system can be concentrated at the regions where the design proofs interface with the physical implementation.

1.1 Design of the Reliable Computing Platform

Traditionally, the operating system function in flight control systems has been implemented as an executive (or main program) that invokes subroutines implementing the application tasks. For ultra-reliable systems, the additional responsibility of providing fault tolerance and undergoing validation makes this approach questionable. We propose a well-defined operating system that provides the applications software developer a reliable mechanism for dispatching periodic tasks on a fault-tolerant computing base that *appears* to him as a single ultra-reliable processor.

Our system design objective is to minimize the amount of experimental testing required and maximize our ability to reason mathematically about correctness. The following design decisions have been made toward that end:

- the system is non-reconfigurable
- the system is frame-synchronous
- the scheduling is static, non-preemptive
- internal voting is used to recover the state of a processor affected by a transient fault

³Future work will concentrate on the massive transient and techniques to detect and restart a massively upset system.

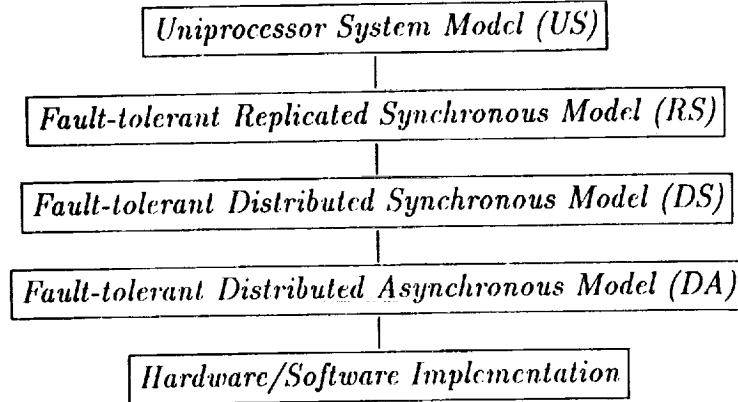


Figure 1: Hierarchical Specification of the Reliable Computing Platform.

A four-level hierarchical decomposition of the reliable computing platform is shown in figure 1.

The top level of the hierarchy describes the operating system as a function that sequentially invokes application tasks. This view of the operating system will be referred to as the *uniprocessor model*, which is formalized as a state transition system in section 3.2 and forms the basis of the specification for the RCP.

Fault tolerance is achieved by voting results computed by the replicated processors operating on the same inputs. Interactive consistency checks on sensor inputs and voting of actuator outputs require synchronization of the replicated processors. The second level in the hierarchy describes the operating system as a synchronous system where each replicated processor executes the same application tasks. The existence of a global time base, an interactive consistency mechanism and a reliable voting mechanism are assumed at this level. The formal details of the model, specified as a state transition system, are described in section 3.3.

Although not anticipated during the Phase 1 effort, another layer of refinement was inserted before the introduction of asynchrony. Level 3 of the hierarchy breaks a frame into four sequential phases. This allows a more explicit modeling of interprocessor communication and the time phasing of computation, communication, and voting. The use of this intermediate model avoids introducing these issues along with those of real time, thus preventing an overload of details in the proof process.

At the fourth level, the assumptions of the synchronous model must be discharged. Rushby and von Henke [8] report on the formal verification of Lamport and Melliar-Smith's [9] interactive-convergence clock synchronization algorithm. This algorithm can serve as a foundation for the implementation of the replicated system as a collection of asynchronously operating processors. Dedicated hardware implementations of the clock synchronization function are a long-term goal.

Final realization of the reliable computing platform is the subject of the Phase 3 effort. The research activity will culminate in a detailed design and prototype implementation.

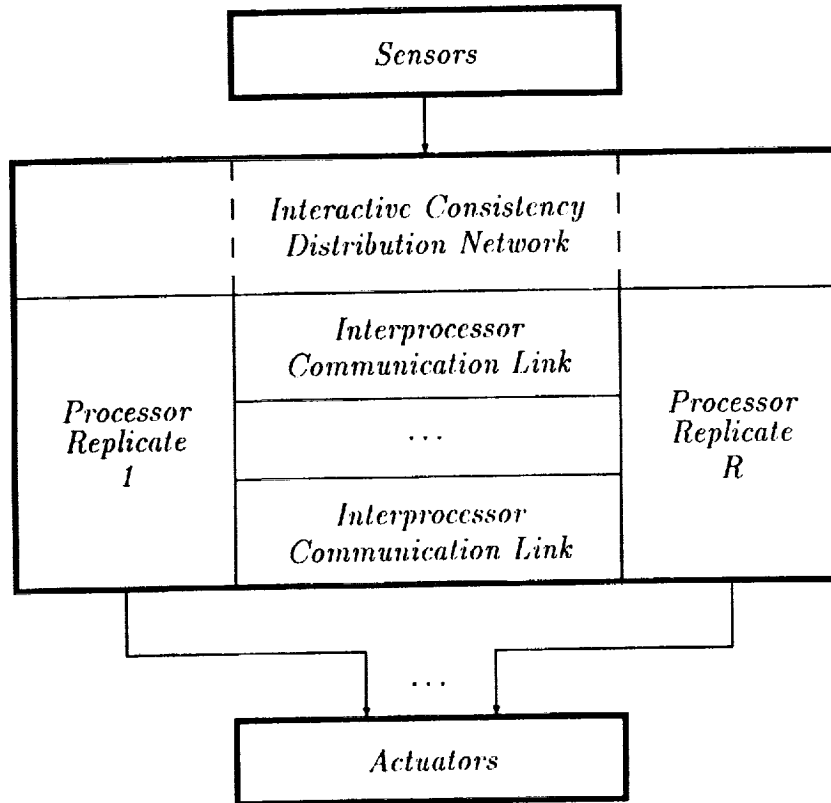


Figure 2: Generic hardware architecture.

Figure 2 depicts the generic hardware architecture assumed for implementing the replicated system. Single-source sensor inputs are distributed by special purpose hardware executing a Byzantine agreement algorithm. Replicated actuator outputs are all delivered in parallel to the actuators, where force-sum voting occurs. Interprocessor communication links allow replicated processors to exchange and vote on the results of task computations. As previously suggested, clock synchronization hardware may be added to the architecture as well.

1.2 Overview of Results

Before presenting the complete details, we provide an overview of the major formalizations and results for the reliable computing platform. In accordance with accepted terminology, we consider a *fault* to be a condition in which a piece of hardware is not operating within its specifications due to physical malfunction, and an *error* to be an incorrect computation result or system output. When a fault occurs, errors may or may not be produced. Although fault-tolerant architectures offer a high degree of immunity from hardware faults, there is a limit to how many simultaneous faults can be tolerated. Unless this limit is exceeded during system operation, the system will mask the occurrence of errors so that the system as a whole produces no computation errors. If the limit is exceeded, however, the system might produce erroneous results.

The primary mechanism for tolerating faults is voting of redundant computation results. Voting can take place at a number of locations in the system and associated with each choice are various tradeoffs. If voting occurs only at the actuators and the internal state of the system (contained in volatile memory) is never subjected to a vote, a single transient fault can permanently corrupt the state of a good processor. This is an unacceptable approach since field data indicates that transient faults are significantly more likely than permanent faults [10]. An alternative voting strategy is to vote the entire system state at frequent intervals. This approach quickly purges the effects of transient faults from the system; however, the computational overhead for this approach may be prohibitive. There is a trade-off between the rate of recovery from transient faults and the frequency of voting. The more frequent the voting, the faster the recovery from transients, but at the price of increased computational overhead. We observe that voting need only occur for a system state that is not recoverable from sensor inputs. A sparse voting approach can accomplish recovery from the effects of transient faults at greatly reduced overhead, but involves increased design complexity. The formal models presented here provide an abstract characterization of the voting requirements for a fault-tolerant system that purges the effects of transient faults.

The proofs we construct are implicitly conditional to account for the situation of limited fault tolerance. The main results we establish can be expressed by the following formula:

$$W(r_1, \dots, r_n) \supset s = V(r_1, \dots, r_n)$$

where W is a predicate to define a minimal working hardware subset over time, s is the uniprocessor model's system results, r_1, \dots, r_n are the results of the replicated processors, and V is a function that selects the properly voted values at each step. Moreover, asynchronous operation is assumed at the lowest specification layer. In this case, we further establish that if the minimal working hardware includes an adequate number of nonfaulty clocks, and clock synchronization is maintained, then the voted outputs continue to match those of higher level specifications. Thus, as long as the system hardware does not experience an unusually heavy burst of component faults, the proof establishes that no erroneous operation will occur at the system level. Individual replicates may produce errors, but they will be out-voted by replicates producing correct results.

If the condition W were true 100% of the time, the system would never fail. Unfortunately, real devices are imperfect and this cannot be achieved in practice. The design of the fault-tolerant architecture must ensure that condition W holds with high probability; typically, the goal is $P(W) \geq 1 - 10^{-9}$ for a 10 hour mission. This condition provides a vital connection between the reliability model and the formal correctness proofs. The proofs conditionally establish that system output is not erroneous as long as W holds, and the reliability model predicts that W will hold with adequately high probability.

In the formal development to follow, we model the possible occurrence of component hardware faults and the unknown nature of computation results produced under such conditions. It is important to note that this modeling is for specification purposes *only* and reflects no self-cognizance on the part of the running system. We assume a nonreconfigurable architecture that is capable of masking the effects of faults, but makes no attempt to detect or diagnose those faults. Each replicate is computing independently and continues to operate the best it can under faulty conditions; it has no knowledge of its own faultiness or that of

its peers. Wherever the formal specifications consider the two cases of whether a processor is faulty or not, it is important to remember that this case analysis is not performed by the running system. Also, it is important to realize that transient-fault recovery is a process that is continually in effect, even when there have been no fault occurrences. Each processor in the system continually votes and replaces its state with voted values. Thus, the transient fault recovery process does not require fault detection.

1.3 Previous Efforts

Many techniques for implementing fault-tolerance through redundancy have been developed over the past decade, e.g. SIFT [11], FTMP [12], FTP [13], MAFT [14], and MARS [15]. An often overlooked but significant factor in the development process is the approach to system verification. In SIFT and MAFT, serious consideration was given to the need to mathematically reason about the system. In FTMP and FTP, the verification concept was almost exclusively testing.

Among previous efforts, only the SIFT project attempted to use formal methods [16]. Although the SIFT operating system was never completely verified [17], the concept of Byzantine Generals algorithms was developed [18] as was the first fault-tolerant clock synchronization algorithm with a mathematical performance proof [9]. Other theoretical investigations have also addressed the problems of replicated systems [19].

Some recent work at SRI International has focused on problems related to the style of fault-tolerant computing adopted by RCP. Rushby has studied a fault masking and transient recovery model and created a formalization of it using EHD [20, 21]. In addition, Shankar has undertaken the formalization of a general scheme for modeling fault-tolerant clock synchronization algorithms [22, 23].

2 Specification Hierarchy and Verification Approach

This section outlines the general methods used in the RCP specifications and proofs. Detailed discussions of the actual specifications appear in later sections.

2.1 The State Machine Approach to Specification

The specification of the Reliable Computing Platform (RCP) is based upon a state-machine method. The behavior of the system is described by specifying an initial state and the allowable transitions from one state to another. The specification of the transition must determine (or constrain) the allowable destination states in terms of the current state and current inputs. One way of doing this is to specify the transition as a function:

$$f_{tran} : state \times input \rightarrow state$$

This is an appealing method when it can be used. A second method is to specify the transition as a mathematical relation between the current state, the input and the new state. One way

to specify a mathematical relation is to define it using a function from the current state, the current input and the new state to a boolean:

$$R : state \times input \times state \rightarrow boolean$$

The function R is true precisely when the relation holds and false, otherwise. The meaning is as follows: a transition from the current state to the new state can occur only when the relation is true. Although the concept is simple it is somewhat awkward to use at first. Consider the function g defined by $g(x) = (x + 4)^2$.

In relational form this function might be expressed by:

$$R(x, y) = [y = (x + 4)^2]$$

The latter form is more awkward than the former when a purely functional relationship exists between x and y . However, a relational approach has some advantages over a functional approach for the specification of complex system behavior. In particular, nondeterminism can be accommodated in a specification by only partially constraining system behavior. For example, if R is changed to the following:

$$R(x, y) = [x > 0 \supset y = (x + 4)^2]$$

the value of y is specified only for positive values of x . In other cases, any value of y would stand in the relation R to x . Such partially constrained specifications are very natural for modeling fault-tolerant systems. It allows us to say nothing about the behavior of failed components, thereby enabling proved results to hold no matter what behavior is exhibited by failed components during system operation.

The relation R would be described as follows in the EHDM specification language:

```
R: function[number, number -> bool] =
  (LAMBDA x,y: (x > 0 IMPLIES y = (x+4)*(x+4)))
```

The first line declares that R is a function from $number \times number$ to the set of booleans (`bool`). The second line uses lambda notation to define the body of the function.

It should also be noted that the modeling approach used in this paper is not based upon a *finite* state machine technique. Some of the components of the state takes values from infinite domains. Therefore, verification tools such as STATEMATE [24] or MCB [25] are not applicable to our specifications.

2.2 Specifying Behavior in the Presence Of Faults

The specification of the RCP system is given in relational form. This enables one to leave unspecified the behavior of a faulty component. Consider the example below.

```
Rtran : function[State, State -> bool] =
  (λ s, t : nonfaulty(s(i)) ⊃ t(i) = f(s(i)))
```

In the relation R_{tran} , if component i of state s is nonfaulty, then component i of the next state t is constrained to be equal to $f(s(i))$. For other values of i , that is, when $s(i)$ is faulty, the next state value $t(i)$ is unspecified. Any behavior of the faulty component is acceptable in the specification defined by R_{tran} .

An alternative approach is to define the transition as a partially-specified function:

$$f_{tran} : \text{function}[\text{State} \rightarrow \text{State}]$$

$$\text{tran_ax} : \text{Axiom } \text{nonfaulty}(s(i)) \supset f_{tran}(s)(i) = g(s(i))$$

This approach does not fit within the definitional structure of EHD. Therefore, one must use an axiom to specify properties of a total, but partially defined function. This leads to a large number of axioms at the base of the proofs and significantly increases the possibility of inconsistency in the axiom set.

2.3 The Specification Hierarchy

The RCP specification consists of four separate models of the system: Uniprocessor System (US), Replicated Synchronous (RS), Distributed Synchronous (DS), Distributed Asynchronous (DA). Each of these specifications is in some sense complete; however, they are at different levels of abstraction and describe the behavior of the system with different degrees of detail. The US model is the most abstract and defines the behavior of the system using a single uninterpreted definition. The RS level supplies more detail. The computation is replicated on multiple processors but the data exchange and voting is captured in one transition. The next level, the DS level, introduces even more detail. Explicit buffers for data exchange are modeled and the transition of the RS level is decomposed into 4 sub-transitions. The DA level introduces time, and different clock times on each of the separate processors.⁴

1. **Uniprocessor System layer (US).** As in the Phase 1 report [1], this constitutes the top-level specification of the functional system behavior defined in terms of an idealized, fault-free computation mechanism. This specification is the correctness criterion to be met by all lower level designs. The top level of the hierarchy describes the operating system as a function that performs an arbitrary, application-specific computation.
2. **Replicated Synchronous layer (RS).** This layer corresponds to level 2 of the Phase 1 report. Processors are replicated and the state machine makes global transitions as if all processors were perfectly synchronized. Interprocessor communication is hidden and not explicitly modeled at this layer. Suitable mappings are provided to enable proofs that the RS layer satisfies the US layer specification. Fault tolerance is achieved using exact-match voting on the results computed by the replicated processors operating on the same inputs. Exact match voting depends on two additional system activities: (1) single source input data must be sent to the redundant sites in a consistent manner to ensure that each redundant processor uses exactly the same inputs during its

⁴Due to the difficulties associated with reasoning about asynchronous systems, it was desirable to perform as much of the design and verification using a synchronous model as possible. Thus, only at level 4 is time explicitly introduced.

computations, and (2) the redundant processing sites must synchronize for the vote. *Interactive consistency* can be achieved on sensor inputs by use of Byzantine-resilient algorithms [18], which are probably best implemented in custom hardware. To ensure absence of single-point failures, electrically isolated processors cannot share a single clock. Thus, a fault-tolerant implementation of the uniprocessor model must ultimately be an asynchronous distributed system. However, the introduction of a fault-tolerant clock synchronization algorithm, at the DA layer of the hierarchy, enables the upper level designs to be performed as if the system were synchronous.

3. **Distributed Synchronous layer (DS).** Next, the interprocessor communication mechanism is modeled and transitions for the RS layer machine are broken into a series of subtransitions. Activity on the separate processors is still assumed to occur synchronously. Interprocessor communication is accomplished using a simple mailbox scheme. Each processor has a mailbox with bins to store incoming messages from each of the other processors of the system. It also has an outgoing box that is used to broadcast data to *all* of the other processors in the system. The DS machine must be shown to implement the RS machine.
4. **Distributed Asynchronous layer (DA).** Finally, the lowest layer relaxes the assumption of synchrony and allows each processor to run on its own independent clock. Clock time and real time are introduced into the modeling formalism. The DA machine must be shown to implement the DS machine provided an underlying clock synchronization mechanism is in place.

The basic design strategy is to use a fault-tolerant clock synchronization algorithm as the foundation of the operating system. The synchronization algorithm provides a global time base for the system. Although the synchronization is not perfect it is possible to develop a reliable communications scheme where the clocks of the system are skewed relative to each other, albeit within a strict known upper bound. For all working clocks p and q , the synchronization algorithm provides the following key property:

$$|c_p(T) - c_q(T)| < \delta$$

assuming that the number of faulty clocks, say m , does not exceed $(nrep-1)/3$, where $nrep$ is the number of replicated processors. This property enables a simple communications protocol to be established whereby the receiver waits until $maxb + \delta$ after a pre-determined broadcast time before reading a message, where $maxb$ is the maximum communication delay.

Each processor in the system executes the same set of application tasks every cycle. A cycle consists of the minimum number of frames necessary to define a continuously repeating task schedule. Each frame is `frame_time` units of time long. A frame is further decomposed into 4 phases. These are the `compute`, `broadcast`, `vote` and `sync` phases. During the `compute` phase, all of the applications tasks scheduled for this frame are executed. The results of all tasks that are to be voted this frame are then loaded into the outgoing mailbox. During the next phase, the `broadcast` phase, the system merely waits a sufficient amount of time to allow all of the messages to be delivered. As mentioned above, this delay must be greater than $maxb + \delta$. During the `vote` phase, each processor retrieves all of the replicated data

from each processor and performs a voting operation. Typically, this operation is a majority vote on each of the selected state elements. The processor then replaces its local memory with the voted values. It is crucial that the vote phase is triggered by an interrupt and all of the vote and state-update code be stored in ROM. This will enable the system to recover from a transient even when the program counter has been affected by a transient fault. Furthermore, the use of ROM is necessary to ensure that the code itself is not affected by a transient.⁵ During the final phase, the sync phase, the clock synchronization algorithm is executed. Although conceptually this can be performed in either software or hardware, we intend to use a hardware implementation.

2.4 Extended State Machine Model

Formalizing the behavior of the Distributed Asynchronous layer requires a means of incorporating time. We accomplish this by formulating an extended state machine model that includes a notion of local clock time for each processor. It also recognizes several types of transitions or operations that can be invoked by each processor. The type of operation dictates which special constraints are imposed on state transitions for certain components.

The time-extended state machine model we use allows for autonomous local clocks on each processor to be modeled using snapshots of clock time coinciding with state transitions. Clock values represent the time at which the last transition occurred (time current state was entered). If a state was entered by processor p at time T and is occupied for a duration D , the next transition occurs for p at time $T + D$ and this clock value is recorded for p in the next state.⁶ A function $c_p(T)$ is assumed to map local clock values for processor p into real time. $c_p(T)$ is a specification-only function; it is not implemented by the system.

Clocks may become skewed in real time. Consequently, the occurrence of corresponding events on different processors may be skewed in real time. A state transition for the DA state machine corresponds to an aggregate transition in which each processor experiences a particular event, such as completing one phase of a frame and beginning the next. Each processor may experience the event at different real times and even different clock times if duration values are not identical.

The DA model is based on a specialized kind of state machine tailored to the needs of an asynchronous system of replicated processors. The intended interpretation is that each component of the state models the local state of one processor and its associated hardware. Each processor is assumed to have a local clock running independently of all the others. Interprocessor communication is achieved by one class of transition that performs a simultaneous broadcast of a portion of the local state variables to all the other processors. Broadcast values are assumed to arrive in the destination mailboxes within a bounded amount of real time $\max b$.

The four classes of transitions are defined as follows:

⁵In the design specifications, these implementation details are not explicitly specified. However, it is clear that in order to successfully implement the models and prove that the implementation performs as specified, such implementation constructs will be needed. These issues will be explored in detail in future work.

⁶We will use the now standard convention of representing clock time with capital letters and real time with lower case letters.

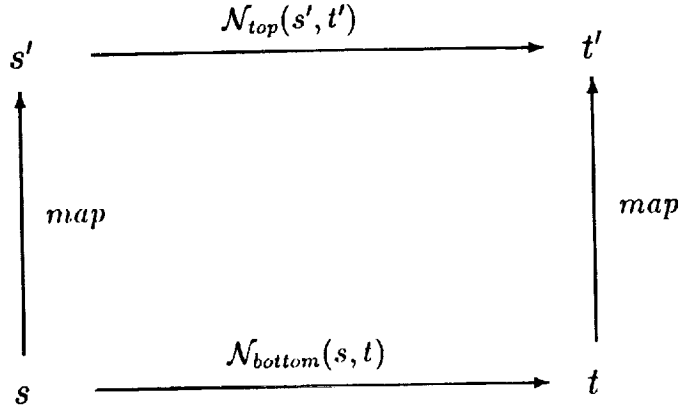


Figure 3: States, transitions, and mappings.

1. **L**: Purely local processing that involves no broadcast communication or reading of the mailboxes.
2. **B**: Broadcast communication where a send is initiated when the state is entered and must be completed before the next transition.
3. **R**: Local processing that involves no send operations, but does include reading of mailbox values.
4. **C**: Clock synchronization operations that may cause the local clock to be adjusted and appear to be discontinuous.

We make the simplifying assumption that the duration spent in each state, except those of type C, is nominally a fixed amount of clock time. Allowances need to be made, however, for small variations in the actual clock time used by real processors. Thus if ν is the maximum rate of variation and D_I, D_A are the intended and actual durations, then $|D_A - D_I| \leq \nu D_I$ must hold.

2.5 The Proof Method

The proof method is a variation of the classical algebraic technique of showing that a homomorphism exists. Such a proof can be visualized as showing that a diagram “commutes” (figure 3). The system is described at two levels of abstraction, which will be referred to as the top and bottom levels for convenience. The top level consists of a current state s' , a destination state, t' and a transition that relates the two. The properties of the transition are given as a mathematical relation, $\mathcal{N}_{top}(s', t')$. Similarly, the bottom level consists of a state s , a destination state, t and a transition that relates the two. The properties of the transition are given as a mathematical relation, $\mathcal{N}_{bottom}(s, t)$. The state values at the bottom level are related to the state values at the top level by way of a mapping function, map . To establish that the bottom level implements the top level one must show that the diagram commutes:

$$\mathcal{N}_{bottom}(s, t) \supset \mathcal{N}_{top}(map(s), map(t))$$

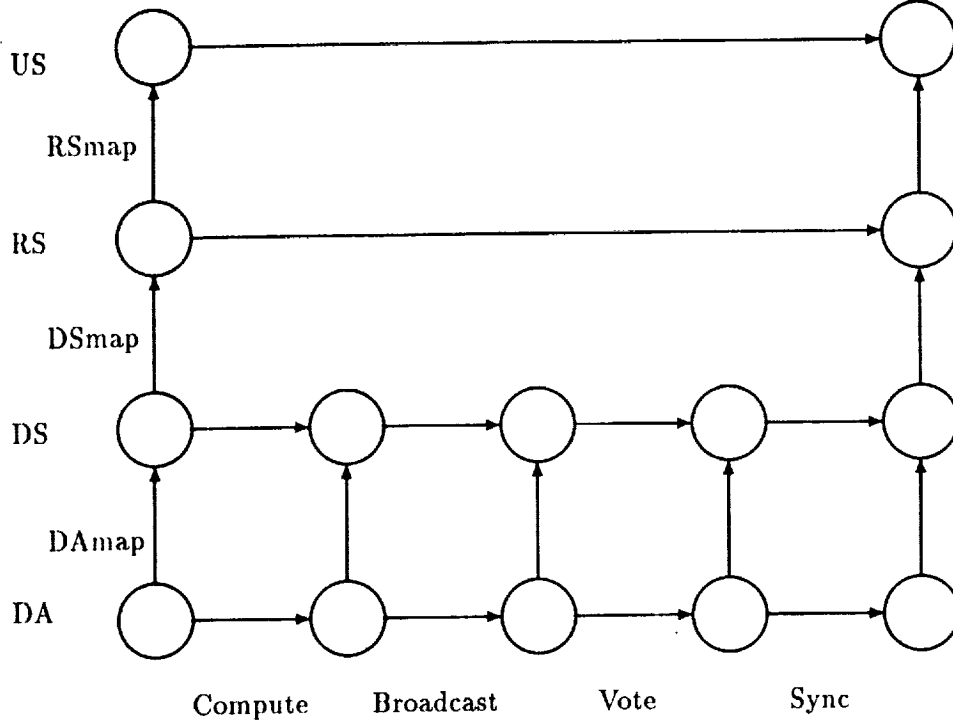


Figure 4: The RCP state machine and proof hierarchy

where $map(s) = s'$ and $map(t) = t'$ in the diagram. One must also show that initial states map up:

$$\mathcal{I}_{bottom}(s) \supset \mathcal{I}_{top}(map(s))$$

An additional consideration in constructing such proofs is that only states reachable from an initial state are relevant. Thus, it suffices to prove a conditional form of commutativity that assumes transitions always begin from reachable states. A weaker form of the theorem is then called for:

$$reachable(s) \wedge \mathcal{N}_{bottom}(s, t) \supset \mathcal{N}_{top}(map(s), map(t))$$

This form enables proofs that proceed by first establishing state invariants. Each invariant is shown to hold for all reachable states and then invoked as a lemma in the main proof.

Figure 4 shows the complete state machine hierarchy and the relationships of transitions within the aggregate model. By performing three layer-to-layer state machine implementation proofs, the states of DA, the lowest layer, are shown to correctly map to those of US, the highest layer. This means that any implementation satisfying the DA specification will likewise satisfy US under our chosen interpretation.

3 US/RS Specification

Up to now we have dealt only with general methods. Next we present the RCP specifications as developed using the EHDM language. An index at the end of this report indicates page numbers where each specification identifier and special symbol is defined in the text. The complete EHDM specifications can be found in Appendix A.

3.1 Preliminary Definitions

The US and RS specifications are expressed in terms of some primitive type definitions. First, we must establish a “domain” or type to represent the complete computation state of a processor. This domain is called **Pstate**. It is declared in EHDM as

Pstate: Type (* computation state of a single processor *)

Thus, all of the state information subject to computation has been collapsed into a single type **Pstate**. Similarly, **inputs** denotes the domain of external system inputs (sensors), and **outputs** the domain of output values that will be sent to the actuators of the system. These domains are named by the following EHDM declarations:

inputs: Type (* type of external sensor input *)

outputs: Type (* actuator output type *)

The number of processors in the system is declared as an arbitrary, positive constant, **nrep**:

nrep: nat (* number of replicated processors *)

The constraint on **nrep**'s value is expressed by the following axiom

processors_exist_ax: Axiom nrep > 0

is a requirement that the system have at least one processor. Nearly all symbolic constants we introduce will have similar constraints imposed on them.

At the RS level and below, information is exchanged among processors via some interprocessor communication mechanism. Additional types are needed to describe the information units involved, being based on a mailbox model of communication. First, we introduce a domain of values for each bin in the mailboxes:

MB : Type (* mailbox exchange type *)

Then we construct a type for a complete mailbox on a processor:

MBvec: Type = array [processors] of MB

This scheme provides one slot in the mailbox array for each replicated processor.

3.2 US Specification

The US specification is very simple:

```
s, t: Var Pstate
u: Var inputs
 $\mathcal{N}_{us}$ : Definition function[Pstate, Pstate, inputs  $\rightarrow$  bool] =
  ( $\lambda s, t, u : t = f_c(u, s)$ )
```

The function \mathcal{N}_{us} defines a mathematical relation between a current state and a final state, i.e., it defines the transition relation. For this model, the transition condition is captured by a function: $f_c(u, s)$, i.e., the computation performed by the uniprocessor system is deterministic and thus can be modeled by a function $f_c : \text{inputs} \times \text{Pstate} \rightarrow \text{Pstate}$. To fit the relational, nondeterministic state machine model we let the state transition relation $\mathcal{N}_{us}(s, t, u)$ hold iff $t = f_c(u, s)$.

External system outputs are selected from the values computed by f_c . The function $f_a : \text{Pstate} \rightarrow \text{outputs}$ denotes the selection of state variable values to be sent to the actuators. The type **outputs** represents a composite of actuator output types.

Although there is no explicit mention of time in the US model, it is intended that a transition correspond to one frame of the execution cycle (i.e., the schedule).

The uninterpreted constant `initial_proc_state` represents the initial Pstate value from which computation begins.

```
initial_us: function[Pstate  $\rightarrow$  bool] = ( $\lambda s : s = \text{initial\_proc\_state}$ )
```

`initial_us` is expressed in predicate form for consistency with the overall relational method of specification, although in this case the initial state value is unique.

3.3 RS Specification

At the RS layer of design, the state is replicated and a postprocessing step is added after computation. This step represents the voting of state variables and thus may be selectively applied. It suffices to encapsulate the entire voting process under a single function of the global state. Nonetheless, it is better to split voting into two parts to facilitate refinement to the DS layer. Another difference introduced at this layer is that the state transition relation needs to be conditioned on the nonfaulty status of each processor.

The global state at this level has type **RSstate**. This is a vector of length `nrep` where each component of the vector defines the state of a specific processor. Each processor in the system can be faulty or nonfaulty as a function of time measured in frames. The local processor "state" must not only reflect the computation state but indicate whether or not a processor is faulty. Such status information about faultiness is included for the purpose of modeling system behavior. An actual system component would be unable to maintain this status and it is understood that this part of the state exists only to model operational behavior and is not an implemented part of the system. Specification of the state type is as follows:

```

rs_proc_state: Type = Record healthy : nat,
                    proc_state : Pstate
end record

```

```

RSstate: Type = array [processors] of rs_proc_state

```

The state of a single processor is given by a record named `rs_proc_state`. The first field of the record is `healthy`, which is 0 when a processor is faulty. Otherwise, it indicates the (unbounded) number of state transitions since the last transient fault. Its value is one greater than the number of prior nonfaulty frames. A permanent fault is indicated by a perpetual value of 0. A processor that is *recovering* from a transient fault is indicated by a value of `healthy` less than the recovery period, denoted by the constant `recovery_period`. This constant is determined by details of the application task schedule and the voting pattern used for transient recovery. A processor is said to be *working* whenever `healthy` \geq `recovery_period`. The second field of the record is the computation state of the processor. It takes values from the same domain as used in the US specification. The complete state at this level, `RSstate`, is a vector (or array) of these records.

Two uninterpreted functions are assumed to express specifications that involve selective voting on portions of the computation state. Their role is described more fully in section 3.5.

```

f_s: function[Pstate → MB] (* state selection for voting *)
f_v: function[Pstate, MBvec → Pstate] (* voting and overwriting *)

```

These two functions split up the selective voting process to mirror what happens in the RCP architecture. First, f_s is used to select a subset of the state components to be voted during the current frame. The choice of which components to vote is assumed to depend on the computation state. It maps into the type `MB`, which stands for a mailbox item. Second, the function f_v takes the current state value and overwrites selected portions of it with voted values derived from a vector of mailbox items. Voting is performed on a component-by-component basis, that is, applied to each task state separately, rather than applied to entire mailbox contents. Note that selection via f_s need not be a mere projection, but could involve more complex data transformations such as adding checksums to ensure integrity during transmission.

Given this background, the transition relation, $\mathcal{N}_{r,s}$, can be defined:

```

N_{r,s}: Definition function[RSstate, RSstate, inputs → bool] =
  (λ s, t, u : (∃ h : (∀ i :
    (s(i)).healthy > 0
    ⊃ good_values_sent(s, u, h(i)) ∧ voted_final_state(s, t, u, h, i)))
  ∧ allowable_faults(s, t))

```

This relation is defined in terms of three subfunctions: `good_values_sent`, `voted_final_state`, and `allowable_faults`. The first aspect of this definition to note is that the relation holds only when `allowable_faults` is true. This corresponds to the “Maximum Fault Assumption” discussed in [1], namely that a majority of processors have been working up to the current time. The next thing to notice is that the transition relation is defined in terms of a conjunction `good_values_sent(s,u,h(i)) ∧ voted_final_state(s,t,u,h,i)`. The meaning is intuitive: the

outputs produced by the good processors are contained in the vector h (i.e., $h(i)$ is derived from the value produced on processor i), and the final state t is obtained by voting the h values. Let us look at the `voted_final_state` relation first.

```
voted_final_state: function[RSstate, RSstate, inputs, MBmatrix, processors → bool]
  = (λ s, t, u, h, i : t(i).proc_state = f_v(f_c(u, s(i).proc_state), h(i)))
```

Processor i is initially in state $s(i)$. If it is nonfaulty ($s(i).healthy > 0$), then its transition to the state $t(i)$ observes the following constraint:

$$t(i).proc_state = f_v(f_c(u, s(i).proc_state), h(i))$$

Otherwise, the behavior of the processor is not defined (i.e., a known mathematical relation is not given). The change to the processor state is defined using two functions: f_c, f_v . The function f_c is the same function used in the US specification. The function f_v operates on the updated computation state and values obtained from the other processors to produce a new state. The idea is that the new state is obtained by replacing local values with voted values.

The values sent by the other processors must satisfy the following relation:

```
good_values_sent: function[RSstate, inputs, MBvec → bool] =
  (λ s, u, w : (∀ j :
    (s(j)).healthy > 0 ⊃ w(j) = f_s(f_c(u, s(j).proc_state))))
```

This relation constrains the $h(i)$ values used in the definition of the \mathcal{N}_r transition relation. Although this function is called with $h(i)$ as an argument, its formal parameter is named w . There is one w value for each processor, which is used to model that processor's mailboxes. If the sending processor j is nonfaulty ($s(j).healthy > 0$), then the value in the receiving mailbox w is given by

$$f_s(f_c(u, s(j).proc_state)).$$

The function f_s selects which portion of the total state is to be voted. Note that since it is a function of the (complete) state, it can differ as a function of the frame, i.e., different data are voted during different frames.

The `allowable_faults` function is defined as follows:

```
allowable_faults: function[RSstate, RSstate → bool] =
  (λ s, t : maj_working(t)
    ∧ (∀ i : t(i).healthy > 0 ⊃ t(i).healthy = 1 + s(i).healthy))
```

This function enforces the restriction imposed by the Maximum Fault Assumption, namely that all reachable states must have a majority of working processors. The condition is expressed in terms of the function `maj_working` and its subordinates:

```
maj_condition: function[set[processors] → bool] =
  (λ A : 2 * card(A) > card(fullset[processors]))
```

`working_proc`: function[RSstate, processors \rightarrow bool] =
 ($\lambda s, p : (s(p)).healthy \geq recovery_period$)

`working_set`: function[RSstate \rightarrow set[processors]] =
 ($\lambda s : (\lambda p : working_proc(s, p))$)

`maj_working`: function[RSstate \rightarrow bool] =
 ($\lambda t : maj_condition(working_set(t))$)

The `working_set` function gives the set of working processors for the current replicated state. The cardinality of this set is then the number of working processors. (Note that sets are usually represented in EHDMM by predicates on the element type. Thus, $(\lambda x : P(x))$ denotes the set $\{x | P(x)\}$.) The relation `allowable_faults` is defined whenever the destination state contains a majority of working processors. It also states that if a processor is nonfaulty for the current frame then the next state's value of `healthy` equals the previous state's value plus one.

The initial state predicate `initial_rs` sets each element of the RS state array to the same value with the `healthy` field equal to `recovery_period` and the `proc_state` field equal to `initial_proc_state`.

`initial_rs`: function[RSstate \rightarrow bool] =
 ($\lambda s : (\forall p : s(p).healthy = recovery_period \wedge s(p).proc_state = initial_proc_state)$)

The constant `recovery_period` is the number of frames required to fully recover a processor's state after experiencing a transient fault. By initializing all `healthy` fields to this value, we are starting the system with all processors *working*.

3.4 Actuator Outputs

The nature of actuator outputs in the RCP application deserves special attention. In the uniprocessor case, an output is produced during each frame and sent to the actuators and no ambiguity exists. In a replicated system, however, multiple actuator values are produced and sent during each frame. Each nonfaulty processor p sends actuator values given by $f_a(rs(p).proc_state)$. There are `nrep` sets of actuator values delivered in parallel, some of which may be copies of previous values for processors that have failed in such a way as to stop generating new values.

It is understood that actuator outputs may be sent through one or more hardware *voting planes* before arriving at the actuators themselves. Other types of signal transformations may be applied to actuator lines between the output drivers and termination points. Additionally, some kind of *force-sum voting* typically is applied at the actuators to mask the presence of errors in one or more channels. All of this activity seeks to ensure that actuators perform as directed by a consensus of processors. These special-purpose requirements of the application leave us unable to completely reflect the proper constraints in the correctness criteria. However, we can use the majority function to map replicated output values into the single actuator output value that would be produced by an ideal uniprocessor. This captures the effect of voting planes and approximates the effect of force-sum voting at the actuators.

To show that replicated actuator outputs can be mapped into a single actuator output, we reason as follows. At the RS level, there are $nrep$ actuator values given by $f_a(rs(p).proc_state)$ for $p = 1, \dots, nrep$. In section 4, a property of RS states is described that asserts that a majority exists among the $proc_state$ values. In other words, a majority of values in $\{rs(p).proc_state\}$ equal $maj(rs)$. Therefore, a majority of $f_a(rs(p).proc_state)$ values exists and is equal to $f_a(maj(rs))$. Since $maj(rs)$, the mapped value of an RS state, is equal to the corresponding US state, this shows that a majority of RS actuator outputs match the value produced by the fault-free US machine.

Note that various additional requirements may be necessary, but are regarded as peculiar to the nature of an RCP application. Hence they must be imposed as correctness criteria beyond those necessary to show that one state machine properly implements another. The intended use of replicated actuator outputs is not contained in the state machine models and may necessitate the use of additional, application-specific correctness conditions.

3.5 Generic Fault-Tolerant Computing

To model a very general class of fault-tolerant, real-time computing schemes, we seek to parameterize the specifications as much as possible. This parameterization takes the form of a set of uninterpreted constants, types, and functions along with axioms to constrain their values. Some instances have already been introduced. The function f_c , for example, represents any computation that can be modeled as a function mapping from inputs and current state into a new state. As hardware redundancy and transient fault recovery are added to the specifications, additional types and functions are needed to express system behavior.

3.5.1 State Model for Transient Fault Recovery

Thus far, we have not concerned ourselves with the internal structure of the computation state $Pstate$. However, to capture the concept of recovering this state information piecewise, it is necessary to make some minimal assumptions about the structure of a $Pstate$ value.

```
control_state: Type (* portion of state used to control or schedule
                    computation activities, e.g., frame counter *)
cell: Type (* index for components of computation state *)
cell_state: Type (* information content of computation state components *)
```

We assume the state contains a control portion, used to schedule and manage computation, and a vector of *cells*, each individually accessible and holding application-specific state information. A sample instantiation of these types is that found in our previous report [1]: the control state is a frame counter and the cells represent the outputs of task instances in the task schedule. Unlike our previous model, however, the more general framework allows a system to maintain state information further back than just the previous execution of a schedule cell.

Also assumed is the existence of access functions to extract and manipulate these items from a $Pstate$ value.

$\text{succ: function}[\text{control_state} \rightarrow \text{control_state}]$ (* next control state *)
 $f_k: \text{function}[\text{Pstate} \rightarrow \text{control_state}]$ (* extracts control state *)
 $f_i: \text{function}[\text{Pstate, cell} \rightarrow \text{cell_state}]$ (* extracts cell (e.g. task) state *)

As described in section 3.3, two additional functions are assumed to express specifications that involve selective voting on portions of the computation state. The functions $f_s: \text{Pstate} \rightarrow \text{MB}$ and $f_v: \text{Pstate} \times \text{MBvec} \rightarrow \text{Pstate}$ were introduced to model the selective voting process applied by each processor. f_s selects which portions of the computation results are subject to voting. f_v takes these selected values from the replicated processors and replaces the required portions of the current state with voted values.

For every voting scheme used for transient fault recovery within RCP, we must be able to determine when the state components have been recovered from voted values. This condition is expressed in terms of the current control state and the number of nonfaulty frames since the last transient fault. Two uninterpreted functions are provided for this purpose.

$\text{rec: function}[\text{cell, control_state, nat} \rightarrow \text{bool}]$

The predicate $\text{rec}(c, K, H)$ is true iff cell c 's state should have been recovered when in control state K with healthy frame count H . Recall that we use a healthy count of one to indicate that the current frame is nonfaulty, but the previous frame was faulty. This means that $H - 1$ healthy frames have occurred prior to the current one.

$\text{dep: function}[\text{cell, cell, control_state} \rightarrow \text{bool}]$

The predicate $\text{dep}(c, d, K)$ indicates that cell c 's value in the next state depends on cell d 's value in the current state, when in control state K . This notion of dependency is different from the notion of computational dependency; it determines which cells need to be recovered in the current frame on the recovering processor for cell c 's value to be considered recovered at the end of the current frame. If cell c is voted during K , or its computation takes only sensor inputs, there is no dependency. If c is not computed during K , c depends only on its own previous value. Otherwise, c depends on one or more cells for its new value.

One derived function is used in the axioms. It asserts that two states X and Y agree on all the corresponding cells on which cell c depends.

$\text{dep_agree: function}[\text{cell, control_state, Pstate, Pstate} \rightarrow \text{bool}] =$
 $(\lambda c, K, X, Y : (\forall d : \text{dep}(c, d, K) \supset f_i(X, d) = f_i(Y, d)))$

3.5.2 Transient Recovery Axioms

Having postulated several functions that characterize a generic fault-tolerant computing application, it is necessary to introduce axioms that sufficiently constrain these functions. Once concrete definitions for the functions have been chosen, these axioms must be proved to follow as theorems for the RCP results to hold for a given application. The eight axioms are presented below.

$\text{succ_ax: Axiom } f_k(f_c(u, \text{ps})) = \text{succ}(f_k(\text{ps}))$

The first axiom states the simple condition that f_c computes the successor of its control state component.

Three axioms give properties of the function rec .

$$\text{full_recovery: Axiom } H \geq \text{recovery_period} \supset \text{rec}(c, K, H)$$

$$\text{initial_recovery: Axiom } \text{rec}(c, K, H) \supset H > 2$$

$$\text{dep_recovery: Axiom } \text{rec}(c, \text{succ}(K), H + 1) \wedge \text{dep}(c, d, K) \supset \text{rec}(d, K, H)$$

First, we require that after the recovery period has transpired, all cells should be considered recovered by rec . Second, it takes a minimum of two frames to recover a cell. (This is necessary because one frame is used to recover the control state. In some applications, it may be possible to recover cells in one frame, but our proof approach does not accommodate those cases and the more conservative minimum of two is used.) Third, if cell c is to be recovered in the next state, all cells it depends on must be recovered in the current state.

components_equal: Axiom

$$f_k(X) = f_k(Y) \wedge (\forall c: f_t(X, c) = f_t(Y, c)) \supset X = Y$$

This axiom, which is a type of *extensionality* axiom, requires that the control state and cell state values form an exhaustive partition of a Pstate value.

Two axioms capture the key conditions for recovery of individual state components.

control_recovered: Axiom

$$\text{maj_condition}(A) \wedge (\forall p: p \in A \supset w(p) = f_s(\text{ps})) \supset f_k(f_v(Y, w)) = f_k(\text{ps})$$

cell_recovered: Axiom

$$\begin{aligned} &\text{maj_condition}(A) \\ &\wedge (\forall p: p \in A \supset w(p) = f_s(f_c(u, \text{ps}))) \\ &\wedge f_k(X) = K \wedge f_k(\text{ps}) = K \wedge \text{dep_agree}(c, K, X, \text{ps}) \\ &\supset f_t(f_v(f_c(u, X), w), c) = f_t(f_c(u, \text{ps}), c) \end{aligned}$$

The first axiom requires that the control state component be recovered after every frame. Thus, f_v must vote the control state unconditionally and update the Pstate value accordingly. The conditions in the antecedent state that for a majority of processors, their mailbox items must match the value selected by the function f_s . The other axiom gives the required condition for recovering an individual cell state value. All cell values that c depends on must already agree with the majority value. After voting with f_v , the function f_t must extract a cell state that matches that of the consensus.

vote_maj: Axiom $\text{maj_condition}(A) \wedge (\forall p: p \in A \supset w(p) = f_s(\text{ps}))$

$$\supset f_v(\text{ps}, w) = \text{ps}$$

The final axiom expresses the additional requirement on f_v that if a majority of processors agree on selected mailbox values derived from state ps , then f_v applied to ps preserves the value ps . In other words, once a Pstate value has been fully recovered, it will stay that way in the face of subsequent voting.

3.5.3 Sample Interpretations of Theory

The proofs of section 4 make use of the foregoing axioms to establish that the RS specification correctly implements the US specification. A valid interpretation of the model provides definitions for the uninterpreted types and functions that are ultimately used to prove the axioms as theorems of the interpreted theory. To maintain the generality of our model and its applicability to a wide range of designs, we do not provide any standard interpretations. Nevertheless, it is desirable to carry out the exercise to establish that the axioms are consistent and can be satisfied for reasonable interpretations.

Two sample interpretations were constructed based on voting schemes introduced in the Phase 1 report [1]. Definitions for the basic concepts of a static, task-based scheduling system were formalized first. Included were the notions of cells as being derived from a frame, subframe pair, and state components to record both the frame counter as well as task outputs. Task execution according to a fixed, repeating schedule was assumed. Definitions were also provided for the *continuous voting* and *cyclic voting* schemes [1]. In both cases, the transient recovery axioms were proved using EHDM. A preliminary form of these specifications are given in Appendix B.

Carrying out the proofs required several changes to the module structure embodied in the specifications of Appendix A. For this reason, the specifications in Appendix B have not yet been integrated with the specifications of Appendix A. Additional work is required to integrate these provisional interpretations into the existing framework. The proofs conducted thus far were performed simply to demonstrate that the axioms could be satisfied and are thus consistent.

The continuous voting scheme requires that all state components are voted during each frame. Hence transient recovery is nearly immediate. Formalizations for this case are very simple and the proofs are trivial. The cyclic voting scheme represents the typical case where state components are voted in the frame they are produced. A cell's value is not voted during frames where it is not recomputed. Formalization in this case is somewhat more involved and the proofs require a bit more effort. The proofs and supporting lemmas comprise about two pages of EHDM specifications. A few selected definitions for the cyclic voting functions are shown below.

```
f_s: function[Pstate → MB] =
  ( λ ps : ps with [(control) := ps.control, (cells) :=
    cell_apply(( λ c : ps.cells(c),
                ps.control,
                null_cell_array,
                num_cells))])
```

```
f_v: function[Pstate, MBvec → Pstate] =
  ( λ ps, w : ps with [(control) := k_maj(w), (cells) :=
    cell_apply(( λ c : t_maj(w, c),
                ps.control,
                ps.cells,
                num_cells))])
```

```

rec: function[cell, control_state, nat → bool] =
  ( λ c, K, H : H
    > 1 + ( if K = cell.frame(c)
            then schedule_length
            else mod_minus(K, cell.frame(c))
            end if))

```

```

dep: function[cell, cell, control_state → bool] =
  ( λ c, d, K : cell.frame(c) ≠ K ∧ c = d)

```

A few supporting definitions are omitted; these functions are presented merely to show the general order of complexity involved.

4 RS to US Proof

Proving that the RS state machine correctly implements the US state machine involves introducing a mapping between states of the two machines. The function `RSmap` defines the required mapping, namely the majority of `Pstate` values over all the processors.

```

RSmap: function[RSstate → Pstate] = ( λ rs : maj(rs))

```

```

maj: function[RSstate → Pstate]

```

```

maj_ax: Axiom ( ∃ A :
  maj_condition(A) ∧ ( ∀ p : p ∈ A ⊃ (rs(p)).proc_state = us))
  ⊃ maj(rs) = us

```

The two theorems required to establish that RS implements US are the following.

```

frame_commutates: Theorem reachable(s) ∧  $\mathcal{N}_{rs}(s, t, u) \supset \mathcal{N}_{us}(RSmap(s), RSmap(t), u)$ 

```

```

initial_maps: Theorem initial_rs(s) ⊃ initial_us(RSmap(s))

```

The theorem `frame_commutates`, depicted in figure 5, shows that a successive pair of reachable RS states can be mapped by `RSmap` into a successive pair of US states. The theorem `initial_maps` shows that an initial RS state can be mapped into an initial US state.

The notion of state reachability is used to express the theorem `frame_commutates`. This concept is formalized as follows:⁷

```

rs_measure: function[RSstate, nat → nat] == ( λ rs, k : k)
reachable_in_n: function[RSstate, nat → bool] =
  ( λ t, k : if k = 0
            then initial_rs(t)
            else ( ∃ s, u : reachable_in_n(s, k - 1) ∧  $\mathcal{N}_{rs}(s, t, u)$ )
            end if) by rs_measure
reachable: function[RSstate → bool] = ( λ t : ( ∃ k : reachable_in_n(t, k)))

```

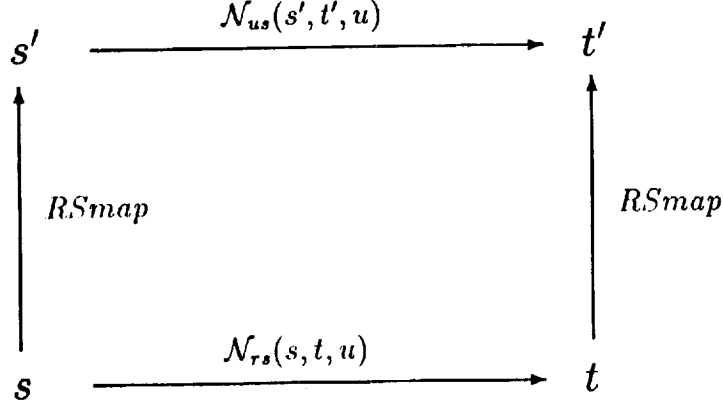


Figure 5: Mappings in the RS to US proof.

Proofs for the two main theorems are supported by a handful of lemmas. The most important is a state invariant that relates values of various state components to their corresponding consensus values.

```
state_invariant: function[RSstate_prop → bool] =
  (λ rs_prop : (∀ t : reachable(t) ⊃ rs_prop(t)))
```

```
state_rec_inv: Lemma state_invariant(state_recovery)
```

```
control_recovery: function[RSstate → bool] =
  (λ s : (∀ p : (s(p)).healthy > 1 ⊃ f_k((s(p)).proc_state) = f_k(maj(s))))
```

```
cell_recovery: function[RSstate → bool] =
  (λ s : (∀ p, c :
    rec(c, f_t((s(p)).proc_state), (s(p)).healthy)
    ⊃ f_t((s(p)).proc_state, c) = f_t(maj(s), c)))
```

```
state_recovery: function[RSstate → bool] =
  (λ s : maj_exists(s) ∧ control_recovery(s) ∧ cell_recovery(s))
```

The invariant `state_recovery` is shown to hold for all reachable states. The control recovery condition of this invariant asserts that if a processor p has been nonfaulty for at least one frame, then the control state, as extracted by f_k , is equal to the consensus value. Similarly, the cell recovery condition asserts that if cell c is due to be recovered, as indicated by the predicate `rec`, then cell state c , as extracted by f_t , is equal to the consensus value. Proving the invariant requires invoking the axioms presented in section 3.5.

Lemmas showing that a majority among RS state values continues to exist after every state transition are also proved in support of the invariant. One such lemma is also central to the proof of `frame_commutes`.

⁷Note that functions defined with “==”, such as in `rs_measure`, are semantically equivalent to those defined with “=”; the only difference is automatic expansion of “==” functions during theorem proving.

rec_maj_f.c: Lemma

$$\text{maj_working}(s) \wedge \text{state_recovery}(s) \wedge \mathcal{N}_{r,s}(s, t, u) \supset \text{maj}(t) = f_c(u, \text{maj}(s))$$

With a majority of working processors and `state_recovery` holding in current state s , this lemma concludes that `maj` applied to the next state t equals the computation step f_c applied to `maj` of s . From this lemma it is clear how RS states and their images under `maj` will correspond to the desired US states.

With the `state_recovery` invariant established, most of the work needed to prove the main theorem `frame_commutates` is in hand. One additional lemma is useful to bridge the gap between the two.

$$\begin{aligned} \text{working_majority: function[RSstate} \rightarrow \text{bool]} = \\ (\lambda s : (\forall p : p \in \text{working_set}(s) \supset (s(p)).\text{proc_state} = \text{maj}(s))) \end{aligned}$$

$$\text{consensus_prop: Lemma state_recovery}(s) \supset \text{working_majority}(s)$$

The lemma `consensus_prop` allows us to draw a key inference from the `state_recovery` invariant, which is expressed by the predicate `working_majority`. This predicate asserts that for all processors p that belong to the *working set*, i.e., for all *working* processors, p 's value of `Pstate` is equal to the majority value.

The proof of `frame_commutates` now follows from `rec_maj_f.c` and `consensus_prop` and assorted definitions. The proof of `initial_maps` follows from definitions and the lemma `initial_maj_cond`, which states that an initial state satisfies the majority condition.

$$\text{initial_maj_cond: Lemma initial_rs}(s) \supset \text{maj_condition}(\text{working_set}(s))$$

This completes the proof that the RS machine implements the US machine.

Note that our proof is in terms of a generic model of fault-tolerant computation and depends on the validity of the axioms of section 3.5. For some choices of definitions for the uninterpreted functions, there will be substantial work required to establish those axioms as theorems. For example, the Minimal Voting scheme presented in our Phase 1 report [1] requires a nontrivial proof to establish that full recovery is achieved. Such details have been omitted here. Nevertheless, the value of our revised approach is in its generality. The results can now be made to apply to a wide variety of frame-based, fault-tolerant architectures.

5 DS Specification

In the Distributed Synchronous layer we focus on two things: expanding the state to include “mailboxes” for interprocessor communication and dividing a frame transition into four sequential subtransitions. The state must also be expanded to include an indicator of which phase of a frame is currently being processed. This is done as follows.

The structure of the mailbox for a four-processor system is shown in figure 6. Each processor contains a mailbox with one slot dedicated to each other processor in the system. Each slot is large enough to contain the largest amount of data to be broadcast during one frame. The n th slot of processor n serves as the outgoing mailbox.

The local state for each processor can now be defined:

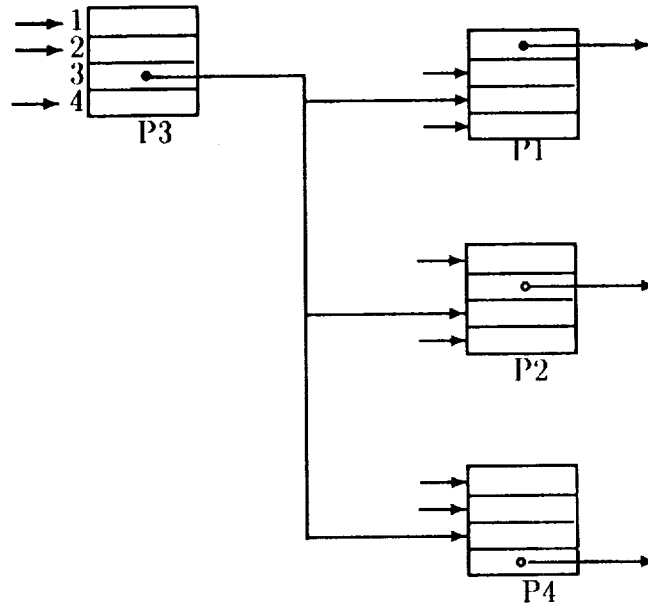


Figure 6: Structure of Mailboxes in a four-processor system

```

ds_proc_state: Type = Record healthy : nat,
                    proc_state : Pstate,
                    mailbox : MBvec
                    end record

```

The vector of all processors `ds_proc_state` is named `ds_proc_array`:

```

ds_proc_array: Type = array [processors] of ds_proc_state

```

The complete DSstate is:

```

DSstate: Type = Record phase : phases,
                    proc : ds_proc_array
                    end record

```

In the DS specification, a frame is decomposed into four phases:

```

phases: Type = (compute, broadcast, vote, sync)

```

The first field of `DSstate` holds the current phase. During each phase a distinct function is performed.

1. **Computation.** The `proc_state` component of the state is updated with the results of computation using the function f_c .
2. **Broadcast.** Interprocessor communication is effected by broadcasting the `MB` values to all other processors, which are deposited in their respective mailboxes.

3. **Voting.** The received mailbox values are voted and merged with the current Pstate values to arrive at the end-of-frame state.
4. **Synchronization.** The clock synchronization function is performed. (No details of the clocks are introduced until the DA specification layer.)

The transition relation for the frame is defined in terms of a phase-transition relation \mathcal{N}_{ds} .

$$\begin{aligned} \text{frame_N_ds: function[DSstate, DSstate, inputs} \rightarrow \text{bool]} = \\ (\lambda s, t, u : (\exists x, y, z : \\ \mathcal{N}_{ds}(s, x, u) \wedge \mathcal{N}_{ds}(x, y, u) \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))) \end{aligned}$$

Note how the intermediate states are defined using existential quantifiers and that the output state of a phase transition becomes the input of the next phase transition. The net result of performing these four phase transitions will be shown to accomplish the same thing as the single transition of the RS specification.

The phase-transition relation is defined as follows:

$$\begin{aligned} \mathcal{N}_{ds}: \text{function[DSstate, DSstate, inputs} \rightarrow \text{bool]} = \\ (\lambda s, t, u : \text{maj_working}(t) \\ \wedge t.\text{phase} = \text{next_phase}(s.\text{phase}) \\ \wedge (\forall i : \\ \text{if } s.\text{phase} = \text{sync} \\ \text{then } \mathcal{N}_{ds}^s(s, t, i) \\ \text{else } t.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \\ \wedge (s.\text{phase} = \text{compute} \supset \mathcal{N}_{ds}^c(s, t, u, i)) \\ \wedge (s.\text{phase} = \text{broadcast} \supset \mathcal{N}_{ds}^b(s, t, i)) \\ \wedge (s.\text{phase} = \text{vote} \supset \mathcal{N}_{ds}^v(s, t, i)) \\ \text{end if})) \end{aligned}$$

Notice that the phase-transition relation only holds when the next state t has a majority of working processors. This corresponds to the analogous condition in \mathcal{N}_{rs} presented in section 3.3, where it appears as one conjunct of the `allowable_faults` relation. Hence, all reachable states in the DS specification must have a majority of working processors.

The phase field of the state is advanced by the function `next_phase`. The phase-transition relation is defined in terms of four sub-relations: \mathcal{N}_{ds}^c , \mathcal{N}_{ds}^b , \mathcal{N}_{ds}^v , and \mathcal{N}_{ds}^s , which correspond to the `compute`, `broadcast`, `vote` and `sync` phases, respectively. The quantifier $\forall i$ invokes the sub-relations for all of the processors of the system. Note that the statement $t.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$ after the `else` requires that the value of `healthy` remain constant throughout a frame. Thus, if a processor is faulty anywhere in a frame it is considered to be faulty throughout; the value of `healthy` may only change at the frame boundaries, i.e., at the `sync` to `compute` transitions. Similarly, full recovery of state information does not occur until the end of a frame. This is consistent with the previous work [1].

Table 1 provides a summary of the functions that are performed during each phase on *nonfaulty* processors. In the table s_i is an abbreviation for $s.\text{proc}(i)$.

The \mathcal{N}_{ds}^c sub-relation defines the behavior of a single processor during the `compute` phase:

Phase	Held constant	Modified
<i>compute</i>	healthy	$t_i.\text{proc_state} = f_c(u, s_i.\text{proc_state})$ $t_i.\text{mailbox}(i) = f_s(f_c(u, s_i.\text{proc_state}))$
<i>broadcast</i>	proc_state healthy	$(\forall p : t_i.\text{mailbox}(p) = s_p.\text{mailbox}(p))$
<i>vote</i>	mailbox healthy	$t_i.\text{proc_state} = f_v(s_i.\text{proc_state}, s_i.\text{mailbox})$
<i>sync</i>	proc_state	$t_i.\text{healthy} = 1 + s_i.\text{healthy}$

Table 1: Summary of activities during various phases

$$\begin{aligned}
\mathcal{N}_{ds}^c: & \text{function}[\text{DSstate}, \text{DSstate}, \text{inputs}, \text{processors} \rightarrow \text{bool}] = \\
& (\lambda s, t, u, i : \\
& \quad s.\text{proc}(i).\text{healthy} > 0 \\
& \quad \supset t.\text{proc}(i).\text{proc_state} = f_c(u, s.\text{proc}(i).\text{proc_state}) \\
& \quad \wedge t.\text{proc}(i).\text{mailbox}(i) = f_s(f_c(u, s.\text{proc}(i).\text{proc_state})))
\end{aligned}$$

During this phase, the `proc_state` field is updated with the results of the computation:

$$f_c(u, s.\text{proc}(i).\text{proc_state})$$

Also, the mailbox is loaded with the subset of the results to be broadcast as defined by the function f_s . Recall that a processor's own mailbox slot acts as the place to post outgoing data for broadcast to other processors.

The \mathcal{N}_{ds}^b sub-relation defines the behavior of a single processor during the broadcast phase:

$$\begin{aligned}
\mathcal{N}_{ds}^b: & \text{function}[\text{DSstate}, \text{DSstate}, \text{processors} \rightarrow \text{bool}] = \\
& (\lambda s, t, i : s.\text{proc}(i).\text{healthy} > 0 \\
& \quad \supset t.\text{proc}(i).\text{proc_state} = s.\text{proc}(i).\text{proc_state} \\
& \quad \wedge \text{broadcast_received}(s, t, i))
\end{aligned}$$

During this phase the `proc_state` field remains unchanged and the `broadcast_received` relation holds:

$$\begin{aligned}
\text{broadcast_received:} & \text{function}[\text{DSstate}, \text{DSstate}, \text{processors} \rightarrow \text{bool}] = \\
& (\lambda s, t, q : (\forall p : \\
& \quad s.\text{proc}(p).\text{healthy} > 0 \\
& \quad \supset t.\text{proc}(q).\text{mailbox}(p) = s.\text{proc}(p).\text{mailbox}(p)))
\end{aligned}$$

This states that each nonfaulty processor q receives the values sent by other nonfaulty processors. If the sending processor p is faulty, then the consequent of the relation need not hold and the value found in p 's slot of q 's mailbox is indeterminate. If the receiving processor q is faulty, the `broadcast_received` relation is not required to hold in \mathcal{N}_{ds}^b . In this situation, all of q 's mailbox values are unspecified.

The \mathcal{N}_{ds}^v sub-relation defines the behavior of a single processor during the vote phase:

$$\begin{aligned}
\mathcal{N}_{ds}^v: \text{function}[DSstate, DSstate, \text{processors} \rightarrow \text{bool}] = \\
& (\lambda s, t, i : s.\text{proc}(i).\text{healthy} > 0 \\
& \quad \supset t.\text{proc}(i).\text{mailbox} = s.\text{proc}(i).\text{mailbox} \\
& \quad \quad \wedge t.\text{proc}(i).\text{proc_state} \\
& \quad \quad = f_v(s.\text{proc}(i).\text{proc_state}, s.\text{proc}(i).\text{mailbox}))
\end{aligned}$$

During this phase the `mailbox` field remains unchanged and the local processor state is updated with the result of voting the values broadcast by the other processors. The vote function is named f_v .

The \mathcal{N}_{ds}^s sub-relation defines the behavior of a single processor during the sync phase:

$$\begin{aligned}
\mathcal{N}_{ds}^s: \text{function}[DSstate, DSstate, \text{processors} \rightarrow \text{bool}] = \\
& (\lambda s, t, i : (s.\text{proc}(i).\text{healthy} > 0 \\
& \quad \supset t.\text{proc}(i).\text{proc_state} = s.\text{proc}(i).\text{proc_state}) \\
& \quad \wedge (t.\text{proc}(i).\text{healthy} > 0 \\
& \quad \quad \supset t.\text{proc}(i).\text{healthy} = 1 + s.\text{proc}(i).\text{healthy}))
\end{aligned}$$

During the `sync` phase, the computation state of a nonfaulty processor remains unchanged. At the end of the `sync` phase, the current frame ends, so the value of `healthy` is incremented by one if the processor is to be nonfaulty in the next frame. This is the same condition appearing in the relation `allowable_faults` of section 3.3. Any processor assumed to be faulty in the next frame will have its `healthy` field set to zero. A limit on how many processors can be faulty simultaneously is imposed by the predicate `maj_working`. Therefore, not every possible assignment of values to the `healthy` fields is admissible; each assignment must satisfy the Maximum Fault Assumption.

The predicate `initial_ds` puts forth the conditions for a valid initial state. The initial phase is set to `compute` and each element of the DS state array has its `healthy` field equal to `recovery_period` and its `proc_state` field equal to `initial_proc_state`.

$$\begin{aligned}
\text{initial_ds}: \text{function}[DSstate \rightarrow \text{bool}] = \\
& (\lambda s : s.\text{phase} = \text{compute} \\
& \quad \wedge (\forall i : s.\text{proc}(i).\text{healthy} = \text{recovery_period} \\
& \quad \quad \wedge s.\text{proc}(i).\text{proc_state} = \text{initial_proc_state}))
\end{aligned}$$

As before, the constant `recovery_period` is the number of frames required to fully recover a processor's state after experiencing a transient fault. By initializing the `healthy` fields to this value, we are starting the system with all processors *working*. Note that the `mailbox` fields are *not* initialized; any `mailbox` values can appear in a valid initial `DSstate`.

6 DS to RS Proof

The DS specification performs the functionality of the RS specification in four sequential steps. Thus, we must show that the “frame” transition function, `frame_N_ds`,

$$\text{frame_N_ds}(s, t, u) = (\exists x, y, z : \mathcal{N}_{ds}(s, x, u) \wedge \mathcal{N}_{ds}(x, y, u) \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))$$

accomplishes the same function as a single transition of the RS level transition function $\mathcal{N}_{rs}(s, t, u)$ under an appropriate mapping function.

6.1 DS to RS Mapping

The DS to RS mapping function, $DSmap$, is defined as:

$$DSmap: \text{function}[DSstate \rightarrow RSstate] = (\lambda ds : ss_update(ds, nrep))$$

where ss_update is given by:

```

ss_update: Recursive function[DSstate, nat  $\rightarrow$  RSstate] =
  ( $\lambda ds, p$  : if ( $p = 0$ )  $\vee$  ( $p > nrep$ )
    then rs0
    else ss_update( $ds, p - 1$ )
    with [( $p$ ) := rsproc0
          with [(healthy) :=  $ds.proc(p).healthy$ ,
                (proc_state) :=  $ds.proc(p).proc\_state$ ]]
    end if) by ssu_measure
  
```

This mapping copies the `healthy` and `proc_state` fields for each processor as illustrated in figure 7. To establish that DS implements RS, the commutativity diagram of figure 8 must

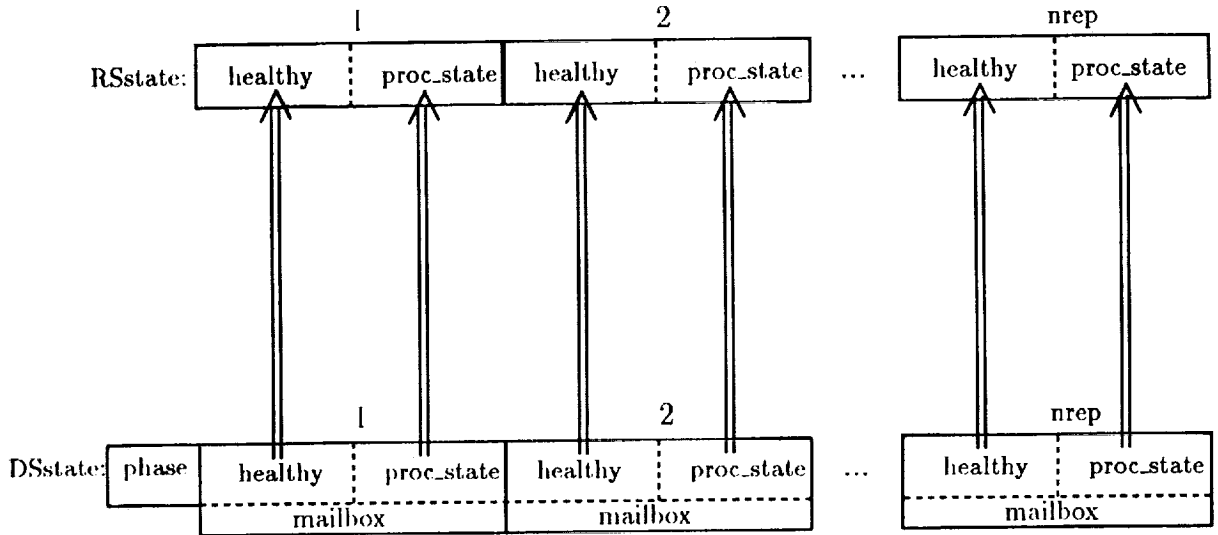


Figure 7: Mapping DS to RS: the $DSmap$ function

be shown to commute. To establish that the diagram commutes, the following formula must be proved.

frame_commutes: Theorem

$$s.phase = compute \wedge \text{frame_N_ds}(s, t, u) \supset \mathcal{N}_{rs}(DSmap(s), DSmap(t), u)$$

Note that to make the correct correspondence, we must consider only DS states found at the beginning of each frame, namely those whose phase is `compute`. Refer to figure 4 on page 12 for a visual interpretation of this theorem.

It is also necessary to show that the initial states are mapped properly:

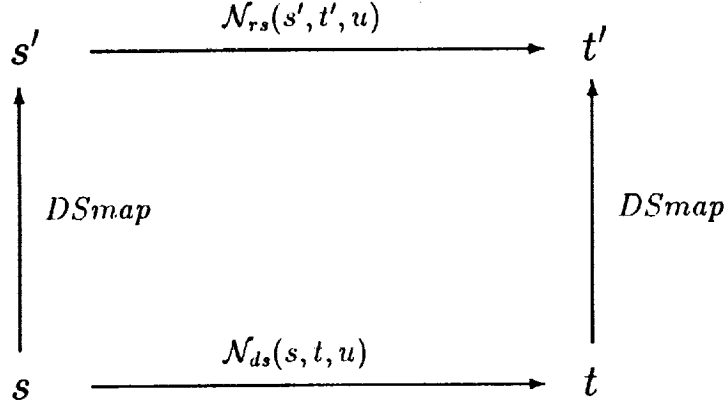


Figure 8: Commutative Diagram for DS to RS Proof

initial_maps: Theorem $\text{initial_ds}(s) \supset \text{initial_rs}(\text{DSmap}(s))$

Several basic lemmas follow from the definition of the mapping function:

map_1: Lemma $\text{DSmap}(s)(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$

map_2: Lemma $\text{DSmap}(s)(i).\text{proc_state} = s.\text{proc}(i).\text{proc_state}$

map_3: Lemma $\text{allowable_faults}(s, t) \supset \text{RS.allowable_faults}(\text{DSmap}(s), \text{DSmap}(t))$

map_4: Lemma $\text{RS.good_values_sent}(\text{DSmap}(s), u, w) = \text{good_values_sent}(s, u, w)$

map_5: Lemma $\text{RS.voted_final_state}(\text{DSmap}(s), \text{DSmap}(t), u, h, i) = \text{voted_final_state}(s, t, u, h, i)$

map_7: Lemma $\text{RS.maj_working}(\text{DSmap}(s)) = \text{DS.maj_working}(s)$

6.2 The Proof

The proof of the `frame_commutates` theorem involves the expansion of the `frame_N_ds` relation and showing that the resulting formula logically implies $\mathcal{N}_{rs}(\text{DSmap}(s), \text{DSmap}(t), u)$. We begin with the definition of `frame_N_ds`:

$$\text{frame_N_ds}(s, t, u) = (\exists x, y, z : \mathcal{N}_{ds}(s, x, u) \wedge \mathcal{N}_{ds}(x, y, u) \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))$$

Since $s.\text{phase} = \text{compute}$, $\mathcal{N}_{ds}(s, x, u)$ can be rewritten as:

$$\begin{aligned} \mathcal{N}_{ds}(s, x, u) = & \text{maj_working}(x) \wedge x.\text{phase} = \text{broadcast} \\ & \wedge (\forall i : x.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \wedge \mathcal{N}_{ds}^c(s, x, u, i)) \end{aligned}$$

Substituting for $\mathcal{N}_{ds}(s, x, u)$ we obtain

$$\begin{aligned}
s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u) \\
\supset (\exists x, y, z : \text{maj_working}(x) \\
\wedge (\forall i : x.\text{phase} = \text{broadcast} \\
\wedge x.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \wedge \mathcal{N}_{ds}^c(s, x, u, i)) \\
\wedge \mathcal{N}_{ds}(x, y, u) \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))
\end{aligned}$$

Next, expand \mathcal{N}_{ds}^c , the \mathcal{N}_{ds} term for the broadcast phase, and combine universal quantifiers:

$$\begin{aligned}
s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u) \\
\supset (\exists x, y, z : \text{maj_working}(x) \wedge \text{maj_working}(y) \\
\wedge (\forall i : x.\text{phase} = \text{broadcast} \\
\wedge x.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \\
\wedge (s.\text{proc}(i).\text{healthy} > 0 \\
\supset x.\text{proc}(i).\text{proc_state} = f_c(u, s.\text{proc}(i).\text{proc_state})) \\
\wedge y.\text{phase} = \text{vote} \\
\wedge y.\text{proc}(i).\text{healthy} = x.\text{proc}(i).\text{healthy} \\
\wedge (x.\text{proc}(i).\text{healthy} > 0 \\
\supset (y.\text{proc}(i).\text{proc_state} = x.\text{proc}(i).\text{proc_state} \\
\wedge (\forall j : x.\text{proc}(j).\text{healthy} > 0 \\
\supset y.\text{proc}(i).\text{mailbox}(j) = f_s(x.\text{proc}(j).\text{proc_state})))))) \\
\wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))
\end{aligned}$$

Simplifying to eliminate x yields:

$$\begin{aligned}
s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u) \\
\supset (\exists y, z : \text{maj_working}(y) \\
\wedge (\forall i : y.\text{phase} = \text{vote} \\
\wedge y.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \\
\wedge (s.\text{proc}(i).\text{healthy} > 0 \\
\supset (y.\text{proc}(i).\text{proc_state} = f_c(u, s.\text{proc}(i).\text{proc_state}) \\
\wedge (\forall j : s.\text{proc}(j).\text{healthy} > 0 \\
\supset y.\text{proc}(i).\text{mailbox}(j) = f_s((y.\text{proc}(j)).\text{proc_state})))))) \\
\wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))
\end{aligned}$$

Expanding the \mathcal{N}_{ds} term for the third phase and simplifying produces:

$$\begin{aligned}
s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u) \\
\supset (\exists z : \text{maj_working}(z) \\
\wedge (\forall i : z.\text{phase} = \text{sync} \\
\wedge z.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \\
\wedge (s.\text{proc}(i).\text{healthy} > 0 \\
\supset z.\text{proc}(i).\text{proc_state} = f_v(f_c(u, s.\text{proc}(i).\text{proc_state}), z.\text{proc}(i).\text{mailbox}) \\
\wedge (\forall j : s.\text{proc}(j).\text{healthy} > 0 \\
\supset z.\text{proc}(i).\text{mailbox}(j) = f_s(f_c(u, (s.\text{proc}(j)).\text{proc_state})))))) \\
\wedge \mathcal{N}_{ds}(z, t, u))
\end{aligned}$$

Expanding the fourth phase \mathcal{N}_{ds} term and simplifying gives:

$$\begin{aligned}
s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u) \\
\supset (\exists z : \text{maj_working}(t) \\
\wedge (\forall i : t.\text{phase} = \text{compute} \\
\wedge (s.\text{proc}(i).\text{healthy} > 0 \\
\supset t.\text{proc}(i).\text{proc_state} = f_v(f_c(u, s.\text{proc}(i).\text{proc_state}), z.\text{proc}(i).\text{mailbox}) \\
\wedge (\forall j : s.\text{proc}(j).\text{healthy} > 0 \\
\supset z.\text{proc}(i).\text{mailbox}(j) = f_s(f_c(u, (s.\text{proc}(j)).\text{proc_state})))))) \\
\wedge (t.\text{proc}(i).\text{healthy} > 0 \\
\supset t.\text{proc}(i).\text{healthy} = 1 + s.\text{proc}(i).\text{healthy}))
\end{aligned}$$

Letting $h(i) = z.\text{proc}(i).\text{mailbox}$,

$$\begin{aligned}
s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u) \\
\supset \text{maj_working}(t) \\
\wedge (\exists h : (\forall i : t.\text{phase} = \text{compute} \\
\wedge (t.\text{proc}(i).\text{healthy} > 0 \\
\supset t.\text{proc}(i).\text{healthy} = 1 + s.\text{proc}(i).\text{healthy}) \\
\wedge (s.\text{proc}(i).\text{healthy} > 0 \\
\supset t.\text{proc}(i).\text{proc_state} = f_v(f_c(u, s.\text{proc}(i).\text{proc_state}), h(i)) \\
\wedge (\forall j : s.\text{proc}(j).\text{healthy} > 0 \\
\supset h(i)(j) = f_s(f_c(u, (s.\text{proc}(j)).\text{proc_state}))))))
\end{aligned}$$

This must be shown to logically imply $\mathcal{N}_{rs}(\text{DSmap}(s), \text{DSmap}(t), u)$, which can be rewritten as:

$$\begin{aligned}
(\exists h : (\forall i : s.\text{proc}(i).\text{healthy} > 0 \\
\supset (\forall j : s.\text{proc}(j).\text{healthy} > 0 \supset h(i)(j) = f_s(f_c(u, s.\text{proc}(j).\text{proc_state})))) \\
\wedge t.\text{proc}(i).\text{proc_state} = f_v(f_c(u, s.\text{proc}(i).\text{proc_state}), h(i)))) \\
\wedge \text{allowable_faults}(s, t)
\end{aligned}$$

The first conjunct can be seen to follow by inspection. By expanding `allowable_faults`,

$$\begin{aligned}
\text{allowable_faults: function[RSstate, RSstate} \rightarrow \text{bool]} = \\
(\lambda s, t : \text{maj_working}(t) \\
\wedge (\forall i : t(i).\text{healthy} > 0 \supset (t(i)).\text{healthy} = 1 + s(i).\text{healthy}))
\end{aligned}$$

the second conjunct can be seen to follow as well. Q.E.D.

7 DA Specification

The DA specification performs the same functions as the DS specification; however, explicit consideration is given to the timing of the system. Every processor of the system has its own clock and consequently task executions on one processor take place at different times than on other processors. Nevertheless, the model at this level explicitly takes advantage of the fact that the clocks of the system are synchronized to within a bounded skew δ . Therefore, it is necessary to give an overview of clock synchronization theory before elaborating the DA specification.

7.1 Clock Synchronization Theory

In this section we will discuss the synchronization theory upon which the DA specification depends. Although the RCP architecture does not depend upon any particular clock synchronization algorithm, we have used the specification for the interactive consistency algorithm (ICA) [9, 8] since EHDM specifications for ICA already exist.

In this section we show the essential aspects of this theory. The formal definition of a clock is fundamental. A clock can be modeled as a function from real time t to clock time T : $C(t) = T$ or as a function from clock time to real time: $c(T) = t$. Since the ICA theory was expressed in terms of the latter, we will also be modeling clocks as functions from clock time to real time. We must be careful to distinguish between an uncorrected clock and a clock which is being resynchronized periodically. We will use the notation $c(T)$ for an uncorrected clock and $rt^{(i)}(T)$ to represent a synchronized clock during its i th frame.⁸

Good clocks have different drift rates with respect to perfect time. Nevertheless, this drift rate is bounded. Thus, we can define a good clock as one whose drift rate is strictly bounded by $\rho/2$. A clock is "good", (i.e. a predicate $\text{good_clock}(T_0, T_n)$ is true), between clock times T_0 and T_n iff:

$$\begin{aligned} & (\forall T_1, T_2 : T_0 \leq T_1 \leq T_n \wedge T_0 \leq T_2 \leq T_n \\ & \quad \supset |c_p(T_1) - c_p(T_2) - (T_1 - T_2)| \leq \frac{\rho}{2} * |T_1 - T_2|) \end{aligned}$$

The synchronization algorithm is executed once every frame of duration `frame_time`. The notation $T^{(i)}$ is used to represent the start of the i th frame, i.e., $(T^0 + i * \text{frame_time})$. The notation $T \in R^{(i)}$ means that T falls in the i th frame, i.e.,

$$(\exists \Pi : 0 \leq \Pi \leq \text{frame_time} \wedge T = T^{(i)} + \Pi)$$

During the i th frame the synchronized clock on processor p , rt_p , is defined by:

$$rt_p(i, T) = c_p(T + \text{Corr}_p^{(i)})$$

where Corr is the cumulative sum of the corrections that have been made to the (logical) clock. It is defined by :

$$\begin{aligned} \text{Corr}_p^{(i)} = & \text{if } i > 0 \text{ then } \text{Corr}_p^{(i-1)} + \Delta_p^{(i-1)} \\ & \text{else } \text{initial_Corr}(p) \\ & \text{end if} \end{aligned}$$

where $\text{initial_Corr}(p)$ is conveniently equated to zero (i.e. $\text{Corr}_p^{(0)} = 0$). The function $\Delta_p^{(i-1)}$ is the correction factor for the current frame as computed by the clock synchronization algorithm.

We now define what is meant by a clock being nonfaulty in the current frame. The predicate nonfaulty_clock is defined as follows:

$$\text{A1: Lemma } \text{nonfaulty_clock}(p, i) = \text{goodclock}(p, T^{(0)} + \text{Corr}_p^{(0)}, T^{(i+1)} + \text{Corr}_p^{(i)})$$

⁸This differs from the notation, $c^{(i)}(T)$, used in [8].

Note that in order for a clock to be non-faulty in the current frame it is necessary that it has been working continuously from time 0.⁹

The clock synchronization theory provides two important properties about the clock synchronization algorithm, namely that the skew between good clocks is bounded and that the correction to a good clock is always bounded. The maximum skew is denoted by δ and the maximum correction is denoted by Σ . More formally,

Clock Synchronization Conditions: For all nonfaulty clocks p and q :

$$S1: \forall T \in R^{(i)} : |rt_p^{(i)}(T) - rt_q^{(i)}(T)| < \delta$$

$$S2: |\text{Corr}_p^{(i+1)} - \text{Corr}_p^{(i)}| < \Sigma$$

The value of δ is determined by several key parameters of the synchronization system: $\rho, \epsilon, \delta_0, m, \text{nrep}$ listed in table 2. The formal definition of ρ has already been given. The

parameter	meaning
ρ	upper bound on drift rate of a good clock
ϵ	upper bound on error in reading another processor's clock
δ_0	upper bound on initial skew
m	maximum number of faulty clocks tolerated
nrep	number of clocks in system

Table 2: Meaning of Synchronization Parameters

parameter ϵ is a bound on the error in reading another processor's clock. The synchronization algorithm requires that every processor in the system obtain an estimate of its skew relative to every other clock in the system. The notation $\Delta_{qp}^{(i)}$ is used to represent the skew between clocks q and p during the i th frame as perceived by p . Thus, the real time at which p 's clock reads $T_0 + \Delta_{qp}^{(i)}$ should be very close to the real time that q 's clock reads T_0 . This is constrained by an axiom to be less than ϵ :

$$\begin{aligned} \text{Axiom} \text{ If conditions S1 and S2 hold throughout the } i\text{th frame, then} \\ \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \\ \supset |\Delta_{qp}^{(i)}| \leq \text{sync_time} \\ \wedge (\exists T_0 : T_0 \in S^{(i)} \wedge |rt_p^{(i)}(T_0 + \Delta_{qp}^{(i)}) - rt_q^{(i)}(T_0)| < \epsilon) \end{aligned}$$

The amount of time reserved for executing the clock synchronization algorithm is denoted by the constant `sync_time`.

The third parameter, δ_0 , is constrained as follows:

$$A0: \text{Axiom } |rt_p^{(0)}(0) - rt_q^{(0)}(0)| < \delta_0$$

⁹This is a limitation not of the operating system, but of existing, mechanically verified fault-tolerant clock synchronization theory. Future work will concentrate on how to make clock synchronization robust in the presence of transient faults.

Thus, δ_0 bounds the initial clock skew.

The property that the ICA clock synchronization algorithm meets the two synchronization conditions S1 and S2 was proved in [8]. These were named **Theorem_1** and **Theorem_2**: formally as:

Theorem_1: Theorem

$$\begin{aligned} \text{S1A}(i) \supset (\forall p, q : (\forall T : \\ \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T \in R^{(i)} \\ \supset |rt_p^{(i)}(T) - rt_q^{(i)}(T)| \leq \delta) \end{aligned}$$

Theorem_2: Theorem $|\text{Corr}_p^{(i+1)} - \text{Corr}_p^{(i)}| < \Sigma$

where the premise for **Theorem_1**, **S1A**, is defined by:

$$(\lambda i : (\forall r : (m + 1 \leq r \text{ and } r \leq n) \supset \text{nonfaulty_clock}(r, i)))$$

and where m is equal to the maximum number of faulty processors.

We have used the following equivalent but more convenient premise: **S1A** : function[period \rightarrow bool] == $(\lambda i : \text{enough_clocks}(i))$.¹⁰ where

$$\begin{aligned} \text{enough_clocks: function[period} \rightarrow \text{bool]} = \\ (\lambda i : 3 * \text{num_good_clocks}(i, \text{nrep}) > 2 * \text{nrep}) \end{aligned}$$

and

$$\begin{aligned} \text{num_good_clocks: Recursive function[period, nat} \rightarrow \text{nat]} = \\ (\lambda i, k : \text{if } k = 0 \vee k > \text{nrep} \\ \text{then } 0 \\ \text{elseif nonfaulty_clock}(k, i) \\ \text{then } 1 + \text{num_good_clocks}(i, k - 1) \\ \text{else num_good_clocks}(i, k - 1) \\ \text{end if) by num_measure} \end{aligned}$$

The theorems proved in [8] also depend upon the following axioms not mentioned above.

$$\text{A2_aux: Axiom } \Delta_{pp}^{(i)} = 0$$

$$\text{C0: Axiom } m < \text{nrep} \wedge m \leq \text{nrep} - \text{num_good_clocks}(i, \text{nrep})$$

$$\text{C1: Axiom frame_time} \geq 3 * \text{sync_time}$$

$$\text{C2: Axiom sync_time} \geq \Sigma$$

$$\text{C3: Axiom } \Sigma \geq \Delta$$

$$\text{C4: Axiom } \Delta \geq \delta + \epsilon + \frac{\rho}{2} * \text{sync_time}$$

$$\text{C5: Axiom } \delta \geq \delta_0 + \rho * \text{frame_time}$$

$$\begin{aligned} \text{C6: Axiom } \delta \geq 2 * (\epsilon + \rho * \text{sync_time}) + 2 * m * \Delta / (\text{nrep} - m) \\ + \text{nrep} * \rho * \text{frame_time} / (\text{nrep} - m) + \rho * \Delta \\ + \text{nrep} * \rho * \Sigma / (\text{nrep} - m) \end{aligned}$$

¹⁰Note that this form also subsumes axiom C0 below.

With the S1A premise expanded, the main synchronization theorem becomes:

sync_thm: Theorem enough_clocks(i)
 $\supset (\forall p, q : (\forall T : T \in R^{(i)} \wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i)$
 $\supset |rt_p^{(i)}(T) - rt_q^{(i)}(T)| \leq \delta))$

The proof that DA implements DS depends crucially upon this theorem.

7.2 The DA Formalization

Now that a clock synchronization theory is at our disposal, the DA model can be specified. Two new fields are added to the state vector associated with each processor: `lclock` and `cum_delta`:

```
da_proc_state: Type = Record healthy : nat,
                          proc_state : Pstate,
                          mailbox : MBvec,
                          lclock : logical_clocktime,
                          cum_delta : number
                          end record
```

The complete `DAstate` is:

```
DAstate: Type = Record phase : phases,
                      sync_period : nat,
                      proc : da_proc_array
                      end record
```

where `da_proc_state` is defined by:

```
da_proc_array: Type = array [processors] of da_proc_state
```

The `sync_period` field holds the current frame of the system. Note this does not represent the frame counter on any particular processor, but rather the ideal, unbounded frame counter.

The `lclock` field of a `DAstate` stores the current value of the processor's local clock. The real-time corresponding to this clock time can be found through use of the auxiliary function `da_rt`.

```
da_rt: function[DAstate, processors, logical_clocktime → realtime] =
  (λ da, p, T : c_p(T + da.proc(p).cum_delta)
```

This function corresponds to the `rt` function of the clock synchronization theory. Thus, `da_rt(s, p, T)` represents processor `p`'s synchronized clock. Given a clock time `T` in the current frame (`s.sync_period`), `da_rt` returns the real-time that processor `p`'s clock reads `T`. The current value of the cumulative correction is stored in the field `cum_delta`.

Every frame the clock synchronization algorithm is executed, and $\Delta_p^{(i)}$ is added to `cum_delta`. Note that this corresponds to the `Corr` function of the clock synchronization theory. The relationship between `c_p`, `da_rt`, and `cum_delta` is illustrated in figure 9.

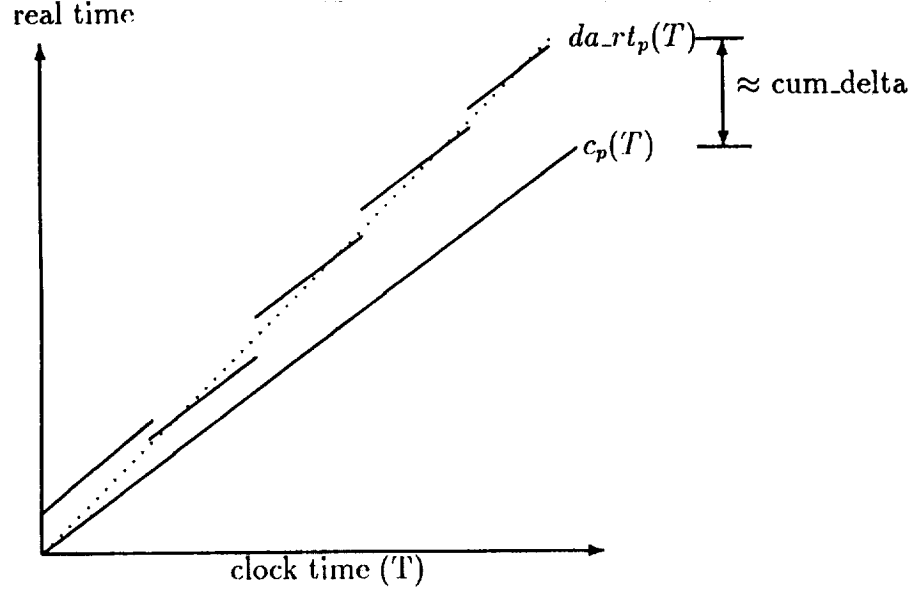


Figure 9: Relationship between c_p and da_rt

Since the original ICA clock theory was not cast into the state-machine framework used in this work, it is necessary to show that the da_rt function is equivalent to the rt function of the clock synchronization theory. The first step is to equate the period of the clock synchronization with the length of a frame in the operating system. Since the length of the period in the clock theory is a parameter of the theory, this is accomplished by setting it equal to `frame_length`. Similarly, the execution time of the synchronization algorithm is a parameter of the clock theory which is set equal to `sync_period`.¹¹ The clock synchronization theory also requires that a constraint be placed on the duration of the `sync` phase:

AXIOM: $\text{duration}(\text{sync}) \geq \text{sync_period}$

The next step is to equate the clocks of the state-machine with the clocks in the `sync` theory. This is done by proving the following lemma:

da_rt_lem: Lemma $\text{reachable}(da) \wedge \text{nonfaulty_clock}(p, da.\text{sync_period})$
 $\supset da_rt(da, p, T) = rt_p^{(da.\text{sync_period})}(T)$

This lemma follows from the fact that in every period (during the `sync` phase) the `cum_delta` field is incremented by Δ_i :

$$t.\text{proc}(i).\text{cum_delta} = s.\text{proc}(i).\text{cum_delta} + \Delta_i^{s.\text{sync_period}}$$

The algorithm that is specified in the clock theory uses Δ_i as its correction factor each frame. The exact same correction factor is used in the DA model. Thus, the RCP system executes

¹¹These are named R and S in [9, 8]. However, these names conflicted with their use in [1].

the same algorithm as specified in the clock theory, and `cum_delta` will always be equal to `Corr`. Thus, $rt_p = da_rt_p$.

The specification of time-critical behavior in the DA model is accomplished using the `da_rt` function. For example, the `broadcast_received` function is expressed in terms of `da_rt`:

```

broadcast_received: function[DAstate, DAstate, processors → bool] =
  ( λ s, t, q : ( ∀ p :
    (s.proc(p).healthy > 0
      ∧ da_rt(s, p, s.proc(p).lclock) + max_comm_delay
        ≤ da_rt(t, q, t.proc(q).lclock)
      ⊃ t.proc(q).mailbox(p) = s.proc(p).mailbox(p)
    )
  )

```

Thus, the data in the incoming bin p on processor q is only defined to be equal to the value broadcast by p (i.e. $s.proc(p).mailbox(p)$) when the real time on the receiving end (i.e. $da_rt(t, q, t.proc(q).lclock)$) is greater than $da_rt(s, p, s.proc(p).lclock)$ plus `max_comm_delay`. This specification anticipates the design of a communications system that can deliver a message in a bounded amount of time, in particular within `max_comm_delay` units of time.

In the DA level there is no single transition that covers the entire frame. There is only a transition relation for a phase. The \mathcal{N}_{da} relation is:

```

 $\mathcal{N}_{da}$ : function[DAstate, DAstate, inputs → bool] =
  ( λ s, t, u : enough_hardware(t) ∧ t.phase = next_phase(s.phase)
    ∧ ( ∀ i : if s.phase = sync
      then  $\mathcal{N}_{da}^s(s, t, i)$ 
      else t.proc(i).healthy = s.proc(i).healthy
        ∧ t.proc(i).cum_delta = s.proc(i).cum_delta
        ∧ t.sync_period = s.sync_period
        ∧ (nonfaulty_clock(i, s.sync_period)
          ⊃ clock_advanced(s.proc(i).lclock, t.proc(i).lclock, duration(s.phase)))
        ∧ (s.phase = compute ⊃  $\mathcal{N}_{da}^c(s, t, u, i)$ )
        ∧ (s.phase = broadcast ⊃  $\mathcal{N}_{da}^b(s, t, i)$ )
        ∧ (s.phase = vote ⊃  $\mathcal{N}_{da}^v(s, t, i)$ )
      end if)
  )

```

Note that the transition to a new state is only valid when the `enough_hardware` function holds in the next state. This function is defined as follows:

```

enough_hardware: function[DAstate → bool] =
  ( λ t : maj_working(t) ∧ enough_clocks(t.sync_period)

```

`maj_working` is defined identically in RS, DS, and DA. Its definition is presented in section 3.3. The definition of `enough_clocks` appears in section 7.1.

As in the DS level, the state transition relation \mathcal{N}_{da} is defined in terms of four sub-relations, each of which applies to a particular phase type. These are called \mathcal{N}_{da}^c , \mathcal{N}_{da}^b , \mathcal{N}_{da}^v and \mathcal{N}_{da}^s .

The \mathcal{N}_{da}^c sub-relation is:

$$\begin{aligned} \mathcal{N}_{da}^c: & \text{function}[D\text{Astate}, D\text{Astate}, \text{inputs}, \text{processors} \rightarrow \text{bool}] = \\ & (\lambda s, t, u, i : \\ & \quad s.\text{proc}(i).\text{healthy} > 0 \\ & \quad \supset t.\text{proc}(i).\text{proc_state} = f_c(u, s.\text{proc}(i).\text{proc_state}) \\ & \quad \wedge t.\text{proc}(i).\text{mailbox}(i) = f_s(f_c(u, s.\text{proc}(i).\text{proc_state})) \end{aligned}$$

Just as in the corresponding DS relation, the `proc_state` field is updated with the results of the computation, $f_c(u, s.\text{proc}(i).\text{proc_state})$. Also, the mailbox is loaded with the subset of the results to be broadcast as defined by the function f_s . Unlike the DS model, the local clock time is changed in the new state. This is accomplished by the predicate `clock_advanced`, which is not based on a simple incrementation operation because the number of clock cycles consumed by an instruction stream will exhibit a small amount of variation on real processors. The function `clock_advanced` accounts for this variability, meaning the start of the next phase is not deterministically related to the start time of the current phase.

ν : number

$$\begin{aligned} \text{clock_advanced}: & \text{function}[\text{logical_clocktime}, \text{logical_clocktime}, \text{number} \rightarrow \text{bool}] = \\ & (\lambda X, Y, D : X + D * (1 - \nu) \leq Y \wedge Y \leq X + D * (1 + \nu)) \end{aligned}$$

where ν represents the maximum rate at which one processor's execution time over a phase can vary from the *nominal* amount given by the `duration` function. ν is intended to be a nonnegative fractional value, $0 \leq \nu < 1$. The *nominal* amount of time spent in each phase is specified by a function named `duration`:

$$\text{duration}: \text{function}[\text{phases} \rightarrow \text{logical_clocktime}]$$

However, the actual amount of clock time spent in a phase is not fixed, but can vary within limits. For example, the actual duration of the `compute` phase can be anything from $(1 - \nu) * \text{duration}(\text{compute})$ to $(1 + \nu) * \text{duration}(\text{compute})$. The value of ν is a parameter of the specification and can be set to any desired value. However, there are some constraints on the implementation that are expressed in terms of ν :

broadcast_duration: Axiom

$$\begin{aligned} & \text{duration}(\text{broadcast}) * (1 - \frac{\delta}{2}) - 2 * \nu * \text{duration}(\text{compute}) - \nu * \text{duration}(\text{broadcast}) - \\ \delta \geq & \text{max_comm_delay} \end{aligned}$$

broadcast_duration2: Axiom

$$\text{duration}(\text{broadcast}) - 2 * \nu * \text{duration}(\text{compute}) - \nu * \text{duration}(\text{broadcast}) \geq 0$$

pos_durations: Axiom

$$\begin{aligned} 0 \leq & (1 - \nu) * \text{duration}(\text{compute}) \wedge 0 \leq (1 - \nu) * \text{duration}(\text{broadcast}) \\ \wedge 0 \leq & (1 - \nu) * \text{duration}(\text{vote}) \wedge 0 \leq (1 - \nu) * \text{duration}(\text{sync}) \end{aligned}$$

all_durations: Axiom

$$\begin{aligned} (1 + \nu) * \text{duration}(\text{compute}) + (1 + \nu) * \text{duration}(\text{broadcast}) \\ \leq \text{frame_time} \end{aligned}$$

The constants ρ and δ are drawn from the clock synchronization theory, as explained in section 7.1.

There may be many possible causes of the variation in execution times on different processors. The asynchronous interface between a processor and its memory can lead to different execution times between two processors even when they execute exactly the same instructions on exactly the same data. Another possible cause of different execution times could be the use of different schedules on different processors.

The \mathcal{N}_{da}^b sub-relation is:

$$\begin{aligned} \mathcal{N}_{da}^b: \text{function}[\text{DAstate}, \text{DAstate}, \text{processors} \rightarrow \text{bool}] = \\ (\lambda s, t, i : s.\text{proc}(i).\text{healthy} > 0 \\ \supset t.\text{proc}(i).\text{proc_state} = s.\text{proc}(i).\text{proc_state} \\ \wedge \text{broadcast_received}(s, t, i)) \end{aligned}$$

As in the corresponding DS relation, the `proc_state` field remains unchanged and the `broadcast_received` relation must hold. When it holds, all the nonfaulty processors receive the values sent by other nonfaulty processors. However, this is now contingent upon certain constraints on the times that things happen.

The \mathcal{N}_{da}^v sub-relation is:

$$\begin{aligned} \mathcal{N}_{da}^v: \text{function}[\text{DAstate}, \text{DAstate}, \text{processors} \rightarrow \text{bool}] = \\ (\lambda s, t, i : s.\text{proc}(i).\text{healthy} > 0 \\ \supset t.\text{proc}(i).\text{mailbox} = s.\text{proc}(i).\text{mailbox} \\ \wedge t.\text{proc}(i).\text{proc_state} = f_v(s.\text{proc}(i).\text{proc_state}, s.\text{proc}(i).\text{mailbox})) \end{aligned}$$

As before, the `mailbox` field remains unchanged and the local processor state is updated with the result of voting the values broadcast by the other processors.

The \mathcal{N}_{da}^s sub-relation is:

$$\begin{aligned} \mathcal{N}_{da}^s: \text{function}[\text{DAstate}, \text{DAstate}, \text{processors} \rightarrow \text{bool}] = \\ (\lambda s, t, i : (s.\text{proc}(i).\text{healthy} > 0 \\ \supset t.\text{proc}(i).\text{proc_state} = s.\text{proc}(i).\text{proc_state}) \\ \wedge (t.\text{proc}(i).\text{healthy} > 0 \\ \supset t.\text{proc}(i).\text{healthy} = 1 + s.\text{proc}(i).\text{healthy} \\ \wedge \text{nonfaulty_clock}(i, t.\text{sync_period})) \\ \wedge t.\text{sync_period} = 1 + s.\text{sync_period} \\ \wedge (\text{nonfaulty_clock}(i, s.\text{sync_period}) \\ \supset t.\text{proc}(i).\text{lclock} = (1 + s.\text{sync_period}) * \text{frame_time} \\ \wedge t.\text{proc}(i).\text{cum_delta} = s.\text{proc}(i).\text{cum_delta} + \Delta_i^{s.\text{sync_period}})) \end{aligned}$$

During the `sync` phase, the processor state remains unchanged. As in the DS specification, the `healthy` field is incremented by one. Unlike the DS model, the local clock time is changed in the new state. For this sub-relation, the clock is not advanced in accordance with the function `clock_advanced`, because this phase is terminated by a clock interrupt. At a predetermined local clock time, the clock interrupt fires and the next frame is initiated. The specification requires that the interrupts fire at clock times that are integral multiples of the frame length, `frame_time`.

In addition to requirements conditioned on having a nonfaulty processor, the DA specifications are concerned with having a nonfaulty clock as well. It is assumed that the clock is an independent piece of hardware whose faults can be isolated from those of the corresponding processor. Although some implementations of a fault-tolerant architecture such as RCP could execute part of the clock synchronization function in software, thereby making clock faults and processor faults mutually dependent, we assume that RCP implementations will have a dedicated hardware clock synchronization function. This means that a clock can continue to function properly during a transient fault period on its adjoining processor. The converse is not true, however. Since the software executing on a processor depends on the clock to properly schedule events, a nonfaulty processor having a faulty clock may produce errors. Therefore, a one-way fault dependency exists.

		Processor		
Clock	Function	Faulty	Recovering	Working
Faulty	Voting	N	N	N
	Clock sync	N	N	N
Nonfaulty	Voting	N	N	Y
	Clock sync	Y	Y	Y

Figure 10: Relationship of clock and processor faults.

Figure 10 summarizes the interaction between clock faults and processor faults. It shows for each combination of fault mode whether a processor can make a sound contribution to voting the state variables and whether a clock can properly contribute to clock synchronization. These conditions have been encoded in the various DA specifications. In particular, the relation \mathcal{N}_{da}^s shown above requires that for a processor to be nonfaulty in the next frame it must have a nonfaulty clock through the end of that frame. Recall that the definition of nonfaulty clock requires that it be continuously nonfaulty from time zero.¹²

The predicate `initial_da` puts forth the conditions for a valid initial state. The initial phase is set to `compute` and the initial sync period is set to zero. Each element of the DA state array has its `healthy` field equal to `recovery_period` and its `proc_state` field equal to `initial_proc_state`.

```

initial_da: function[DAstate → bool] =
  ( λ s : s.phase = compute ∧ s.sync_period = 0
    ∧ ( ∀ i : s.proc(i).healthy = recovery_period
      ∧ s.proc(i).proc_state = initial_proc_state
      ∧ s.proc(i).cum_delta = 0
      ∧ s.proc(i).lclock = 0 ∧ nonfaulty_clock(i, 0)))

```

As before, the constant `recovery_period` is the number of frames required to fully recover a processor's state after experiencing a transient fault. By initializing the `healthy` fields to this

¹²This does not represent a deficiency in the design of the DA model but rather is a limitation imposed by the existing, mechanically verified clock synchronization algorithm. Future work will concentrate on liberating the clock synchronization property from this restriction.

value, we are starting the system with all processors *working*. Note that the mailbox fields are *not* initialized; any mailbox values can appear in a valid initial DAstate.

8 DA to DS Proof

8.1 DA to DS Mapping

The DA to DS mapping function, DAmapping, is defined as:

```
DAmapping: function[DAstate → DSstate] =
  (λ da : ss_update(da, nrep) with [(phase) := da.phase])
```

where `ss_update` is given by:

```
ss_update: Recursive function[DAstate, nat → DSstate] =
  (λ da, k : if (k = 0) ∨ (k > nrep)
    then ds0
    else ss_update(da, k - 1)
    with [(proc)(k) := dsproc0
          with [(healthy) := da.proc(k).healthy,
                (proc_state) := da.proc(k).proc_state,
                (mailbox) := da.proc(k).mailbox]]
    end if) by da_measure
```

Thus, the `lclock`, `cum_delta`, and `sync_period` fields are not mapped (i.e., are abstracted away) and all of the other fields are mapped identically. To establish that DA implements DS, the commutativity diagram of figure 11 must be shown to commute. To establish that the

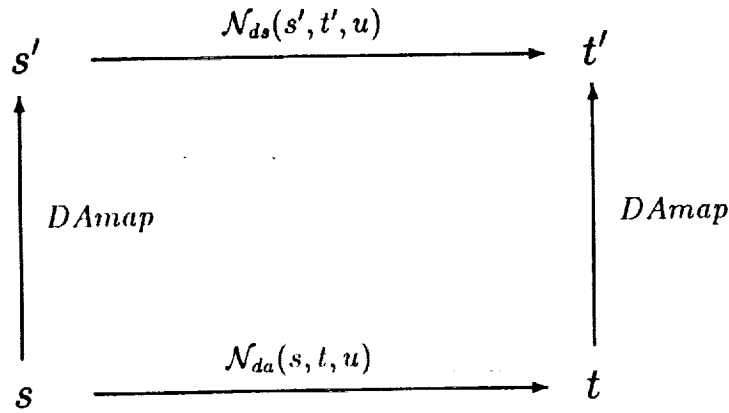


Figure 11: Commutative Diagram for DA to DS Proof

diagram commutes, the following formulas must be proved:

phase_commutates: Theorem $\text{reachable}(s) \wedge \mathcal{N}_{da}(s, t, u) \supset \mathcal{N}_{ds}(\text{DAmapping}(s), \text{DAmapping}(t), u)$

initial_maps: Theorem $\text{initial}_{da}(s) \supset \text{initial}_{ds}(\text{DAmapping}(s))$

The lemmas below directly follow from the definition of the mapping.

- map_1: Lemma $\text{DAm}ap(s).\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
- map_2: Lemma $\text{DAm}ap(s).\text{proc}(i).\text{proc_state} = s.\text{proc}(i).\text{proc_state}$
- map_3: Lemma $\text{DAm}ap(s).\text{phase} = s.\text{phase}$
- map_4: Lemma $\text{DAm}ap(s).\text{proc}(i).\text{mailbox} = s.\text{proc}(i).\text{mailbox}$
- map_7: Lemma $\text{DS.maj_working}(\text{DAm}ap(s)) = \text{DA.maj_working}(s)$

8.2 The Proof

The `phase_commut`s theorem must be shown to hold for all four phases. Thus, the proof is decomposed into four separate cases, each of which is handled by a lemma of the form:

$$\text{phase_com_}\mathcal{X}\text{: Lemma } s.\text{phase} = \mathcal{X} \wedge \mathcal{N}_{da}(s, t, u) \supset \mathcal{N}_{ds}(\text{DAm}ap(s), \text{DAm}ap(t), u)$$

where \mathcal{X} is any one of `{compute, broadcast, vote, sync}`. The proof of this theorem requires the expansion of the \mathcal{N}_{da} relation and showing that the resulting formula logically implies $\mathcal{N}_{ds}(\text{DAm}ap(s), \text{DAm}ap(t), u)$.

8.2.1 Decomposition Scheme

The proof of each lemma `phase_com_` \mathcal{X} is facilitated by using a common, general scheme for each phase that further decomposes the proof by means of four subordinate lemmas. The general form of these lemmas is as follows:

$$\text{Lemma 1: } s.\text{phase} = \mathcal{X} \wedge \mathcal{N}_{da}(s, t, u) \supset (\forall i : \mathcal{N}_{da}^{\mathcal{X}}(s, t, i))$$

$$\text{Lemma 2: } s.\text{phase} = \mathcal{X} \wedge \mathcal{N}_{da}^{\mathcal{X}}(s, t, i) \supset \mathcal{N}_{ds}^{\mathcal{X}}(\text{DAm}ap(s), \text{DAm}ap(t), i)$$

$$\text{Lemma 3: } s.\text{phase} = \mathcal{X} \wedge \text{DS.maj_working}(tt) \wedge (\forall i : \mathcal{N}_{ds}^{\mathcal{X}}(ss, tt, i)) \supset \mathcal{N}_{ds}(ss, tt, u)$$

$$\text{Lemma 4: } s.\text{phase} = \mathcal{X} \wedge \mathcal{N}_{da}(s, t, u) \supset \text{DS.maj_working}(\text{DAm}ap(t))$$

A few differences exist among the lemmas for the four phases, but they adhere to this scheme fairly closely. The `phase_com_` \mathcal{X} lemma follows by chaining the four lemmas together:

$$\mathcal{N}_{da}(s, t, u) \supset (\forall i : \mathcal{N}_{da}^{\mathcal{X}}(s, t, i)) \supset (\forall i : \mathcal{N}_{ds}^{\mathcal{X}}(\text{DAm}ap(s), \text{DAm}ap(t), i)) \supset \mathcal{N}_{ds}(\text{DAm}ap(s), \text{DAm}ap(t), u)$$

In three of the four cases above, proofs for the lemmas are elementary. The proof of Lemma 1 follows directly from the definition of \mathcal{N}_{da} . Lemma 3 follows directly from the definition of \mathcal{N}_{ds} . Lemma 4 follows from the definition of \mathcal{N}_{da} , `enough_hardware` and the basic mapping lemmas.

Furthermore, in three of the four phases, the proof of Lemma 2 is straightforward. For all but the `broadcast` phase, Lemma 2 follows from the definition of $\mathcal{N}_{ds}^{\mathcal{X}}$, $\mathcal{N}_{da}^{\mathcal{X}}$, and the basic mapping lemmas.

However, in the broadcast phase, Lemma 2 from the scheme above, which is named `com_broadcast_2`, is a much deeper theorem. The broadcast phase is where the effects of asynchrony are felt; we must show that interprocessor communications are properly received in the presence of asynchronously operating processors. Without clock synchronization we would be unable to assert that broadcast data is received. Hence the need to invoke clock synchronization theory and its attendant reasoning over inequalities of time.

8.2.2 Proof of `com_broadcast_2`

The lemma `com_broadcast_2` is the most difficult of the four lemmas for the broadcast phase. It follows from the definition of \mathcal{N}_{ds}^b , \mathcal{N}_{da}^b , the basic mapping lemmas and a fairly difficult lemma, `com_broadcast_5`:

`com_broadcast_5`: Lemma

$$\begin{aligned} & \text{reachable}(s) \wedge \mathcal{N}_{da}(s, t, u) \wedge s.\text{phase} = \text{broadcast} \\ & \wedge s.\text{proc}(i).\text{healthy} > 0 \wedge \text{broadcast_received}(s, t, i) \\ & \supset \text{broadcast_received}(\text{DMap}(s), \text{DMap}(t), i) \end{aligned}$$

This lemma deals with the main difference between the DA level and the DS level—the timing constraint on the function `broadcast_received`:

$$\begin{aligned} \text{broadcast_received: function}[\text{DAstate}, \text{DAstate}, \text{processors} \rightarrow \text{bool}] = \\ & (\lambda s, t, q : (\forall p : \\ & \quad (s.\text{proc}(p).\text{healthy} > 0 \\ & \quad \wedge \text{da_rt}(s, p, (s.\text{proc}(p).\text{lclock}) + \text{max_comm_delay} \leq \text{da_rt}(t, q, t.\text{proc}(q).\text{lclock}) \\ & \quad \supset t.\text{proc}(q).\text{mailbox}(p) = s.\text{proc}(p).\text{mailbox}(p) \end{aligned}$$

The timing constraint

$$\text{da_rt}(s, p, s.\text{proc}(p).\text{lclock}) + \text{max_comm_delay} \leq \text{da_rt}(t, q, t.\text{proc}(q).\text{lclock})$$

must be discharged in order to show that the DA level implements the DS level. The following lemma is instrumental to this goal.

$$\begin{aligned} \text{ELT: Lemma } T_2 \geq T_1 + \text{bb} \wedge (T_1 \geq T^0) \wedge (\text{bb} \geq T^0) \wedge T_2 \in R^{(\text{sp})} \wedge T_1 \in R^{(\text{sp})} \\ & \wedge \text{nonfaulty_clock}(p, \text{sp}) \wedge \text{nonfaulty_clock}(q, \text{sp}) \wedge \text{enough_clocks}(\text{sp}) \\ & \supset \text{rt}_p^{(\text{sp})}(T_2) \geq \text{rt}_q^{(\text{sp})}(T_1) + (1 - \frac{\rho}{2}) * |\text{bb}| - \delta \end{aligned}$$

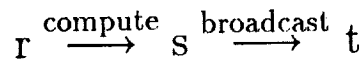
This lemma establishes an important property of timed events in the presence of a fault-tolerant clock synchronization algorithm and is proved in the next subsection. Suppose that on processor q an event occurs at T_1 according to its own clock and another event occurs on processor p at time T_2 according to its own clock. Then, assuming that the clock times fall within the current frame and the clocks are working and the system still is safe (i.e. more than two thirds of the clocks are non-faulty), then the following is true about the real times of the events:

$$\text{rt}_p^{(\text{sp})}(T_2) \geq \text{rt}_q^{(\text{sp})}(T_1) + (1 - \frac{\rho}{2}) * |\text{bb}| - \delta$$

where $bb = T_2 - T_1$, $T_1 = s.\text{proc}(p).\text{lclock}$ and $T_2 = t.\text{proc}(q).\text{lclock}$.

If we apply this lemma to the broadcast phase, letting T_1 be the time that the sender loads his outgoing mailbox bin and T_2 is the earliest time that the receivers can read their mailboxes (i.e. at the start of the `vote` phase), we know that these events are separated in time by more than $(1 - \frac{\epsilon}{2}) * |bb| - \delta$.

In this case bb is approximately equal to `duration(broadcast)`. However, since there may be some variations in the time spent in the compute and broadcast phases on different processors (i.e. they can drift from the nominal value at a rate less than ν), the analysis is a little tricky. First consider the situation where processor q is sending a message to processor p during its broadcast phase. Let r be the state at the start of the compute phase, s be the state at the start of the broadcast phase and t be the state at the start of the vote phase:



Then, let

- R_q = the clock time at the start of the compute phase on processor q
- S_q = the clock time at the start of the broadcast phase on processor q
- T_q = the clock time at the start of the vote phase on processor q
- R_p = the clock time at the start of the compute phase on processor p
- S_p = the clock time at the start of the broadcast phase on processor p
- T_p = the clock time at the start of the vote phase on processor p

This is illustrated in figure 12. By the definition of `clock_advanced`, the following can be

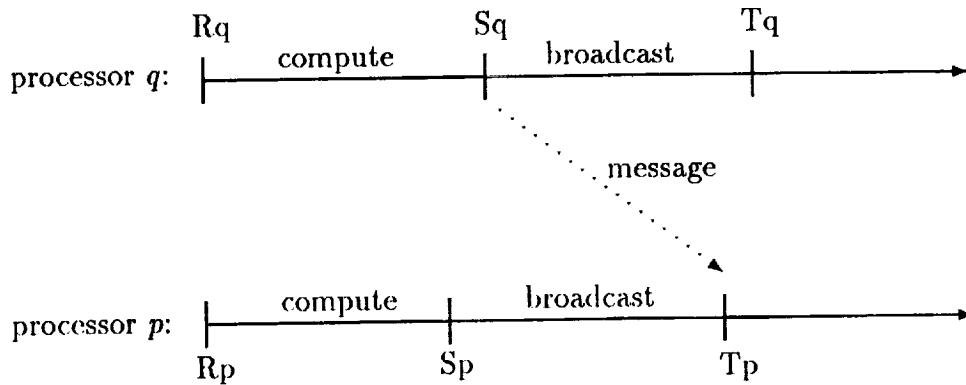


Figure 12: Relationship between phase times on different processors

established:

$$\begin{aligned}
 & (\exists \text{pdurc, pdurb, qdurc, qdurb} : \\
 & \quad \text{near}(\text{pdurc, compute}) \wedge \text{near}(\text{pdurb, broadcast}) \\
 & \quad \wedge \text{near}(\text{qdurc, compute}) \wedge \text{near}(\text{qdurb, broadcast}) \\
 & \quad \wedge R_p = R_q
 \end{aligned}$$

$$\begin{aligned} \wedge Sq &= Rq + qdurc \\ \wedge Tp &= Sq - qdurc + pdurc + pdurb)) \end{aligned}$$

where $\text{near}(\text{dur}, \text{ph})$ is given by

$$\text{near}(\text{dur}, \text{ph}) = (1 - \nu) * \text{duration}(\text{ph}) \leq \text{dur} \leq (1 + \nu) * \text{duration}(\text{ph}))$$

This result depends upon a critical invariant of the system:

$$\begin{aligned} (\forall p, q : s.\text{phase} = \text{compute} \wedge \\ \text{nonfaulty_clock}(p, s.\text{sync_period}) \wedge \text{nonfaulty_clock}(q, s.\text{sync_period}) \\ \supset s.\text{proc}(p).\text{lclock} = s.\text{proc}(q).\text{lclock}) \end{aligned}$$

given that the state s is $\text{reachable}(s)$. This invariant exists in the system because of the use of an interrupt timer to initiate the start of a frame on each of the processors at the pre-determined times $i * \text{frame_time}$. Using the definition of $R^{(i)}$ and the axioms pos_durations and all_durations , we obtain:

$$\begin{aligned} \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \\ \supset Sq \in R^{(i)} \wedge Tp \in R^{(i)} \\ \wedge Tp \geq Sq + \text{duration}(\text{broadcast}) \\ - 2 * \nu * \text{duration}(\text{compute}) - \nu * \text{duration}(\text{broadcast}) \end{aligned}$$

where i is the current synchronization period (i.e. $i = r.\text{sync_period} = s.\text{sync_period} = t.\text{sync_period}$). We now have a relationship between the clock time that the message was sent and the clock time that it was received in a form appropriate for application of the ELT theorem. In other words, $T_2 = Tp$, $T_1 = Sq$ and $\text{bb} = \text{pdurc} - \text{qdurc} + \text{pdurb}$. Thus, we can convert the relationship between the events expressed in clock times to a relationship between the *real* times of these events:

$$rt_p^{(i)}(Tp) \geq rt_q^{(i)}(Sq) + (1 - \frac{\epsilon}{2}) * |\text{duration}(\text{broadcast}) - \text{Epsi}| - \delta$$

where $\text{Epsi} = 2 * \nu * \text{duration}(\text{compute}) + \nu * \text{duration}(\text{broadcast})$. Using the $\text{broadcast_duration}$ implementation axiom:

$$\begin{aligned} \text{broadcast_duration: Axiom} \\ \text{duration}(\text{broadcast}) * (1 - \frac{\epsilon}{2}) - 2 * \nu * \text{duration}(\text{compute}) \\ - \nu * \text{duration}(\text{broadcast}) - \delta \geq \text{max_comm_delay} \end{aligned}$$

we have:

$$rt_p^{(i)}(Tp) \geq rt_q^{(i)}(Sq) + \text{max_comm_delay}$$

Using the da_rt_lem lemma:

$$\text{da_rt}(t, q, Tp) \geq \text{da_rt}(s, p, Sq) + \text{max_comm_delay}$$

This will discharge the premise of $\text{broadcast_received}$. Thus,

com_broadcast_5: Lemma

$$\begin{aligned} & \text{reachable}(s) \wedge \mathcal{N}_{da}(s, t, u) \wedge s.\text{phase} = \text{broadcast} \\ & \wedge s.\text{proc}(p).\text{healthy} > 0 \wedge \text{broadcast_received}(s, t, p) \\ & \supset \text{broadcast_received}(\text{DAMap}(s), \text{DAMap}(t), p) \end{aligned}$$

Of course there are several technicalities such as the $\text{reachable}(\text{da})$ premise that must be discharged in order to apply the da_rt_lem lemma and the other state invariants and establishing that $s.\text{proc}(p).\text{healthy} > 0 \supset \text{nonfaulty_clock}(p, s.\text{sync_period})$.

Proof of ELT Lemma: In this section we prove,

$$\begin{aligned} \text{Lemma 1 (earliest_later_time Lemma)} \quad & T_2 = T_1 + \text{BB} \\ & \wedge (T_1 \geq T^0) \wedge (\text{BB} \geq T^0) \wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \\ & \wedge \text{enough_clocks}(i) \wedge T_2 \in R^{(i)} \wedge T_1 \in R^{(i)} \\ & \supset \text{rt}_p^{(i)}(T_2) \geq \text{rt}_q^{(i)}(T_1) + (1 - \frac{\rho}{2}) * |\text{BB}| - \delta \end{aligned}$$

from which the ELT lemma immediately follows.

Proof. This lemma depends primarily upon the definition of a good clock and the synchronization theorem (i.e. sync_thm). The good clock definition yields:

$$\begin{aligned} & \text{goodclock}(q, T^0, T_1 + \text{BB}) \wedge (T_1 \geq T^0) \wedge (\text{BB} \geq T^0) \\ & \supset (1 - \frac{\rho}{2}) * |\text{BB}| \leq c_q(T_1 + \text{BB}) - c_q(T_1) \\ & \quad \wedge c_q(T_1 + \text{BB}) - c_q(T_1) \leq (1 + \frac{\rho}{2}) * |\text{BB}| \end{aligned}$$

Note that the definition of a good clock is defined in terms of the uncorrected clocks, $c_p(T)$. Using the definition of rt , we can rewrite the first formula as:

$$\begin{aligned} \text{Lemma goodclock}(q, T^0, T_1 + \text{Corr}_q^{(i)} + \text{BB}) \\ & \wedge (T_1 \geq T^0) \wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0) \wedge (\text{BB} \geq T^0) \\ & \supset (1 - \frac{\rho}{2}) * |\text{BB}| \leq \text{rt}_q^{(i)}(T_1 + \text{BB}) - \text{rt}_q^{(i)}(T_1) \\ & \quad \wedge \text{rt}_q^{(i)}(T_1 + \text{BB}) - \text{rt}_q^{(i)}(T_1) \leq (1 + \frac{\rho}{2}) * |\text{BB}| \end{aligned}$$

and obtain a formula in terms of the function rt .

The sync_thm theorem gives us:

$$\begin{aligned} & \text{enough_clocks}(i) \wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T \in R^{(i)} \\ & \supset -\delta \leq \text{rt}_p^{(i)}(T) - \text{rt}_q^{(i)}(T) \leq \delta \end{aligned}$$

Combining the previous two formulas and substituting T_2 for T in sync_thm , we obtain:

$$\begin{aligned} & T_2 = T_1 + \text{BB} \wedge (T_1 \geq T^0) \wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0) \wedge (\text{BB} \geq T^0) \wedge T_2 \in R^{(i)} \\ & \quad \wedge \text{enough_clocks}(i) \wedge \text{goodclock}(q, T^0, T_1 + \text{Corr}_q^{(i)} + \text{BB}) \wedge \text{nonfaulty_clock}(p, i) \wedge \\ & \quad \text{nonfaulty_clock}(q, i) \\ & \supset \text{rt}_p^{(i)}(T_2) \geq \text{rt}_q^{(i)}(T_1) + (1 - \frac{\rho}{2}) * |\text{BB}| - \delta \end{aligned}$$

From the definition of nonfaulty and goodclock , we have:

$$\begin{aligned} & T_1 + \text{BB} \leq T^{(i+1)} \wedge \text{nonfaulty_clock}(q, i) \\ & \supset \text{goodclock}(q, T^0, T_1 + \text{Corr}_q^{(i)} + \text{BB}) \end{aligned}$$

Using these last two results we have:

$$\begin{aligned} T_2 = T_1 + \text{BB} \wedge T_2 \leq T^{(i+1)} \wedge (T_1 \geq T^0) \wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0) \wedge (\text{BB} \geq T^0) \\ \wedge \text{enough_clocks}(i) \wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T_2 \in R^{(i)} \\ \supset \text{rt}_p^{(i)}(T_2) \geq \text{rt}_q^{(i)}(T_1) + (1 - \frac{\rho}{2}) * |\text{BB}| - \delta \end{aligned}$$

Then from the definition of $R^{(i)}$, $T^{(i)}$ and the fact that $\text{Corr}_q^{(0)} = 0$, we have

$$\begin{aligned} \text{ft11: Lemma } T_2 = T_1 + \text{BB} \wedge (T_1 \geq T^0) \wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0) \wedge (\text{BB} \geq T^0) \\ \wedge \text{enough_clocks}(i) \wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T_2 \in R^{(i)} \\ \supset \text{rt}_p^{(i)}(T_2) \geq \text{rt}_q^{(i)}(T_1) + (1 - \frac{\rho}{2}) * |\text{BB}| - \delta \end{aligned}$$

Using the `adj_always_pos` theorem from [8], we obtain

$$\text{ft12: Lemma } T_1 \in R^{(i)} \supset (T_1 + \text{Corr}_q^{(i)} \geq T^0)$$

The key lemma follows immediately from the last two formulas, (ft11 and ft12).

9 Implementation Considerations

Although many RCP design decisions have yet to be made, there are a number of implementation issues that need to be considered early. Some of these have emerged as consequences of the formalization effort completed in Phase 2. Others are the result of preliminary investigations into the needs of implementations that can satisfy the RCP specifications. Following is a discussion of these issues and available options.

9.1 Restrictions Imposed by the DA Model

Recall that the DA extended state machine model described in section 2.4 recognized four different classes of state transition: L, B, R, C. Although each is used for a different phase of the frame, the transition types were introduced because operation restrictions must be imposed on implementations to correctly realize the DA specifications. Failure to satisfy these restrictions can render an implementation at odds with the underlying execution model, where shared data objects are subject to the problems of concurrency. The set of constraints on the DA model's implementation concerns possible concurrent accesses to the mailboxes.

While a broadcast send operation is in progress, the receivers' mailbox values are undefined. If the operation is allowed sufficient time to complete, the mailbox values will match the original values sent. If insufficient time is allowed, or a broadcast operation is begun immediately following the current one, the final mailbox value cannot be assured. Furthermore, we make the additional restriction that all other uses of the mailbox be limited to read-only accesses. This provides a simple sufficient condition for noninterfering use of the mailboxes, thereby avoiding more complex mutual exclusion restrictions.

Operation Restrictions. Let s and t be successive DA states, i be the processor with the earliest value of $c_i(s(i).lclock)$, and j be the processor with the latest

value of $c_j(t(j).\text{lclock})$. If s corresponds to a broadcast (B) operation, all processors must have completed the previous operation of type R by time $c_i(s(i).\text{lclock})$, and the next operation of type B can begin no earlier than time $c_j(t(j).\text{lclock})$. No processor may write to its mailbox during an operation of type B or R.

By introducing a prescribed discipline on the use of mailboxes, we ensure that the axiom describing the net effect of broadcast communication can be legitimately used in the DA proof. Although the restrictions are expressed in terms of real time inequalities over all processors' clocks, it is possible to derive sufficient conditions that satisfy the restrictions and can be established from local processor specifications only, assuming a clock synchronization mechanism is in place.

9.2 Processor Scheduling

The DA model of the RCP deals with the timing and coordination of the replicated processors in a fairly complete manner. The model defines in detail the functionality of the system with regard to the activities that are necessary to ensure its fault-masking and transient recovery capability. Nevertheless, the delineation of the task execution process on each local processor has not been elaborated in any more detail than in the US model. This was done deliberately in order to obtain as general a specification as possible. Thus, the 4-level hierarchy presented in this paper could be further refined into a set of entirely different kinds of implementations. They could differ drastically in the types of task scheduling that are utilized as well as the type of hardware or software used.

Nevertheless, one aspect of scheduling needs to be carefully controlled, namely the basic frame structure. The RCP specifications were developed with a very crisp execution model in mind regarding the basic timing of a frame and its major parts. We assume the existence of one or more nonmaskable hardware interrupts, triggered by the clock subsystem, that are used to effect the transition from one frame to the next and one major phase to the next. As a minimum, the following transitions must be triggered by timer interrupts or an equally strong hardware mechanism.

- **Start of frame.** The last portion of a frame is reserved for clock synchronization activities. This includes not only executing the clock synchronization functions, but also reserving some dead time to be sacrificed when clock adjustments cause local clock time discontinuities. An interrupt is set to fire at the proper value of clock time so that all processors begin the new frame with the same local clock reading.
- **Beginning of vote phase.** After waiting for the completion of broadcast communication from other processors, the vote phase is begun to selectively restore portions of the computation state. Also needing to be recovered are any control state variables used by the operating system. If a transient fault occurs, recovery cannot begin until the control state is first restored through voting. However, a processor operating after a transient fault may be executing with a corrupted memory state. The only way to ensure that corrupted memory does not prevent the eventual recovery of control state information is to force the vote to happen through a nonmaskable interrupt.

The use of timer interrupts are highly desirable in other situations, but those listed above are considered essential.

Scheduling of applications tasks is an area where the implementation retains some flexibility owing to our use of a general fault-tolerant computing model in the US and RS specifications. Often it is considered desirable to achieve some type of schedule diversity across processors as a means of gaining more transient fault immunity. A limited way of accomplishing this is available under the current RCP design. Since the specifications only state what must be true after all tasks have been executed within a frame, it is possible to juggle the order of tasks within each frame to implement diversity. For example, if N tasks are scheduled in a particular frame, each processor may execute them in a different order up to the limits of data dependency among tasks. It is also possible to introduce different spreads of slack time, dummy tasks, etc. to achieve similar effects.

9.3 Hardware Protection Features

Correct recovery of state information after a transient fault has been formalized in the RS to US proof. Transient recovery of state information occurs gradually, one cell at a time. Consequently, depending on the voting pattern used, some tasks will be executing in the presence of erroneous state information. Implicit in the RS specifications is that computation of task outputs is not subject to interference by other tasks executing with erroneous data inputs. In the specifications, this is due simply to the use of a functional representation of the effects of task execution.

Nonetheless, in a real processor a program in execution can interfere with another unless hardware protection mechanisms are in place. To see why this is so, suppose, for instance, that task T_1 is followed by task T_2 in a particular frame and neither's output is voted during that frame. Suppose further that in the transient fault recovery scheme, T_2 's inputs come from recently voted cells while T_1 's do not. Thus, we expect T_2 's cell to be recovered after this frame. After a transient fault, T_1 may be executing instructions on erroneous data, possibly overwriting recovered information such as that required by T_2 . This would invalidate our assumption that T_2 's state is recovered at the end of the frame.

In a similar manner, interference can be caused in the time domain as well as the data domain. In the example above, if T_1 's erroneous input causes it to run longer than its upper execution time bound, T_2 may not get to execute in this frame. Again, this would result in our assumptions about T_2 's output being invalid. Therefore, hardware protection features are required to prevent both kinds of interference in a system that attempts to recover state information selectively.

There are several well-known hardware techniques for providing this type of protection.

- **Memory protection.** Hardware write protection devices are found on many modern computer architectures. What RCP requires is less than a full-blown memory management unit (MMU). All that is necessary is to be able to prevent a task in execution from writing into memory areas for which the operating system has not given explicit write permission. The ability to give a task write access to a small set of physical memory regions is sufficient. Generating hardware exceptions such as traps on illicit write attempts is desirable but not essential.

- **Watch-dog timers.** Timer interrupts or special-purpose timing logic will be required to prevent a task from consuming more than its allotted amount of execution time. When a watch-dog timer is triggered, the operating system need only dispatch the next task on the schedule. The actual hardware used to carry out this timing function needs to have adequate resolution and be distinct from the timer interrupts used to signal the end-of-frame and start-of-voting events.
- **Privileged Operating Modes.** To protect the protection mechanisms, it is usually necessary for a processor to have at least one privileged execution domain. Processors typically provide at least a user domain and a (privileged) supervisor domain to implement conventional operating system designs. In RCP, we need these features so the tasks cannot accidentally change or disable the memory write protection or watch-dog timer functions. There may be other uses for privileged mode as well.

It is important to realize that use of these features may be obviated in special cases. If sufficiently frequent voting is used, for example, it may not be necessary to provide these features as long as a task is always executing with valid data as input.

9.4 Voting Mechanisms

Exact-match voting of state information exchanged among processors is usually envisioned as applying the majority function to mailbox values. Note, however, that the voting function f_v , described in section 3.3, is unspecified and need not be based on the majority operation. Other types of voting may be used provided that the transient recovery axioms of section 3.5.2 are still true.

A desirable alternative to majority voting is *plurality* voting. If the values subject to voting are $\{a, a, b, c\}$, for example, a majority does not exist, but a plurality does, namely $\{a, a\}$. The reason this can be valuable is that during a massive transient fault that affects more than a majority of processors, the Maximum Fault Assumption no longer holds and transient fault recovery is not assured by the proofs previously described. However, the likelihood is that the affected processors will not exhibit exactly the same errors. If a minority of processors is still working, it is likely that the values produced by the replicated processors will appear something like the example $\{a, a, b, c\}$. Hence, plurality voting has a good chance of recovering the correct state in spite of the absence of a working majority.

This problem has been studied by Miner and Caldwell [26]. They showed that the substitution of plurality voting for majority voting can be used to produce identical results as long as the Maximum Fault Assumption holds:

$$\text{maj_exists}(s) \supset \text{maj}(s) = \text{plur}(s)$$

By using an implementation based on plurality voting, we enjoy the same provable behavior when the Maximum Fault Assumption holds, and we enjoy added transient fault immunity in the rare case that it is violated. All that is necessary to achieve this is to show that the choice of function for f_v meets the requirements of the transient recovery axioms.

10 Future Work

There are four main areas where further work may be profitable.

1. Development of a still more detailed specification and verification that it meets the DA specification.
2. Development of task scheduling/voting strategies that satisfy the axioms of the US model.
3. More detailed specification of the behavior of the actuator outputs.
4. Development of a detailed reliability model.

10.1 Further Refinement

Although the DA specification is a fairly detailed design of the system-wide behavior of the RCP, there is very little implementation detail about what occurs locally on each processor. The next level of the specification hierarchy, the local processor LP specification will define the data structures and algorithms to be implemented on each local processor.

At some point the design must be implemented on hardware. It is anticipated that both standard hardware such as microprocessors and memory management units will be required as well as special hardware to implement the clock synchronization and Byzantine agreement functions. In the same way that this work capitalized on the work done elsewhere in clock synchronization, the LP specification will build on the work being performed under contract to NASA Langley in hardware verification.

NASA Langley has awarded three contracts specifically devoted to formal methods (from the competitive NASA RFP 1-22-9130.0238). The selected contractors were SRI International, Computational Logic Inc., and Odyssey Research Associates. Another task-assignment contract with Boeing Military Aircraft Company (BMAC) is being used to explore formal methods as well. Through this contract BMAC is funding research at the University of California at Davis and California Polytechnic State University to assist them in the use of formal methods in aerospace applications. The efforts are roughly divided as follows:

SRI:	Clock synchronization, operating system
CLI:	Byzantine Agreement Circuits, clock synchronization
ORA:	Byzantine Agreement Circuits, applications
BMAC:	Hardware Verification, formal requirements analysis

The DA specification critically depended upon a clock synchronization property. Previous work by SRI had verified that the ICA algorithm meets this property. Ongoing work at SRI is directed at implementing a synchronization algorithm in hardware verifying it. This will lead to the verification hierarchy shown in figure 13.

Implicit in the RS, DS and DA models is the assumption that it is possible to distribute single source information such as sensor data to the redundant processors in a consistent man-

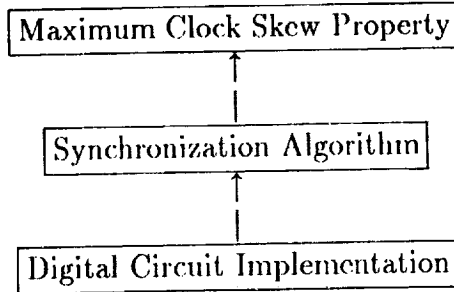


Figure 13: Clock Synchronization Hierarchy

ner even in the presence of faults. This is the classic Byzantine Generals problem [18].¹³ CLI is investigating the formal verification of such algorithms and their implementation. They have formally verified the original Pease, Shostak, and Lamport version of this algorithm using the Boyer Moore theorem prover [27]. They have also implemented this algorithm down to the register-transfer level and demonstrated that it implements the mathematical algorithm [28]. Future work will concentrate on tying this work together with their verified microprocessor, the FM8502 [29].

ORA has also been investigating the formal verification of Byzantine Generals algorithms. They have focused on the practical implementation of a Byzantine-resilient communications mechanism between Mini-Cayuga micro-processors [30]. The Mini-Cayuga is a small but formally verified microprocessor developed by ORA. It is a research prototype and has not been fabricated. This communications circuitry could serve as a foundation for the RCP architecture. It was designed assuming that the underlying processors were synchronized (say by a clock synchronization circuit). The issues involved with connecting the Byzantine communications circuit with a clock synchronization circuit and verifying the combination have not yet been explored.

Boeing Military Aircraft Company and U. C. Davis have been sponsored by NASA, Langley to apply formal methods to the design of conventional hardware devices. Formal Verification of the following circuits is currently under investigation:

- a floating-point coprocessor similar to the Intel 8087 (but smaller) [31, 32].
- a DMA controller similar to the Intel 8237A (but smaller) [33].
- microprocessors in I/O (small) [34, 35, 36].
- a memory management unit [37, 38].

¹³Fault-tolerant systems, although internally redundant, must deal with single-source information from the external world. For example, a flight control system is built around the notion of feedback from physical sensors such as accelerometers, position sensors, pressure sensors, etc. Although these can be replicated (and they usually are), the replicates do not produce identical results. In order to use bit-by-bit majority voting all of the computational replicates must operate on identical input data. Thus, the sensor values (the complete redundant suite) must be distributed to each processor in a manner that guarantees all working processors receive exactly the same value even in the presence of some faulty processors.

The team is currently investigating the verification of a composed set of verified hardware devices [39, 40, 41]

Researchers at NASA Langley have begun a new effort on a hardware clock synchronization technique that can serve as a foundation for the RCP architecture. The method, which is based on the Fault-Tolerant Midpoint algorithm [42], is aimed at a fully independent hardware implementation. The primary goals of this work are full mechanical verification, transient fault recovery, and an initialization scheme that provides recovery from large transient upsets.

10.2 Task Scheduling and Voting

The Phase 1 report described a scheduling system that was based upon a deterministic table. In the models presented in this paper, this is no longer strictly required although such an approach clearly fits within the axioms presented in the US model. However, it is conceivable that more sophisticated scheduling strategies could also be shown to conform.

10.3 Actuator Outputs

It is important not only that the replicated outputs sent to the actuators (on separate wires) are identical but that they appear within some bounded time of each other. Although this bound may not be very small, it is still incumbent upon the verification activity that a bound be mathematically established.

10.4 Development of a Detailed Reliability Model

In the Phase 1 paper, a simple reliability model of the RCP system was developed that demonstrated that the speed at which one must remove the effects of a transient fault is not very critical. In other words, flushing the effects of a transient fault over an extended period of time did not significantly decrease the reliability of the system as compared to extremely fast removal. In this model, a fault anywhere in the processor was sufficient to render the entire processor faulty. Clearly, in a fully developed RCP, there will be more than one fault-isolation containment region per processor. The most likely arrangement is to have a separate fault-containment region for the clocking system and one for the Byzantine agreement circuitry.

11 Concluding Remarks

In this paper a hierarchical specification of a reliable computing platform (RCP) has been developed. The top level specification is extremely general and should serve as a model for many fault-tolerant system designs. The successive refinements in the lower levels of abstraction introduce, first, processor replication and voting, second interprocess communication by use of dedicated mailboxes and finally, the asynchrony due to separate clocks in the system.

Although the first phase of this work was accomplished without the use of an automated theorem prover, we found the use of the EHDM system to be beneficial to this second phase of work for several reasons.

- The amount of detail in the lower level models is significantly greater than in the upper level models. It became extremely difficult to keep up with everything using pencil and paper.
- The strictness of the EHDM language (i.e. its requirement to precisely define all variables and functions, etc.) forced us to elaborate the design more carefully.
- Most of the proofs were not very deep but had to deal with large amounts of detail. Without a mechanical proof checker, it would be far too easy to overlook a flaw in the proofs.
- The proof support environment of EHDM, although overly strict in some cases, provided much assistance in assuring us that our proof chains were complete and that we had not overlooked some unproven lemmas.
- The decision procedures of EHDM for linear arithmetic and propositional calculus were valuable in that they relieved us of the need to reduce many formulas to primitive axioms of arithmetic. Especially useful was its ability to reason about inequalities.

Key features of the work completed during Phase 2 and improvements over the results of Phase 1 include the following.

- Specification of redundancy management and the transient fault recovery scheme uses a very general model of fault-tolerant computing similar to one proposed by Rushby [20, 21].
- Specification of the asynchronous layer design uses modeling techniques based on a time-extended state machine approach. This method allows us to build on previous work that formalized clock synchronization mechanisms and their properties.
- Formulation of the RCP specifications is based on a straightforward Maximum Fault Assumption that provides a clean interface to the realm of probabilistic reliability models. It is only necessary to determine the probability of having a majority of working processors and a two-thirds majority of nonfaulty clocks.
- A four-layer tier of specifications has been completely proved to the standards of rigor of the EHDM mechanical proof system. All proofs can be run on a Sun SPARCstation in less than one hour.
- Important constraints on lower level design and implementation constructs have been identified and investigated.

Based on the results obtained thus far, work will continue to a Phase 3 effort, which will concentrate on completing design formalizations and develop the techniques needed to produce verified implementations of RCP architectures.

Acknowledgements

The authors are grateful for the many helpful suggestions given by Dr. John Rushby of SRI International. His suggestions during the early phases of model formulation and decomposition lead to a significantly more manageable proof activity. The authors are also grateful to John and Sam Owre for the timely assistance given in the use of the EHDM system. This work was supported in part by NASA contract NAS1-19341.

References

- [1] Di Vito, Ben L.; Butler, Ricky W.; and Caldwell, James L., II: *Formal Design and Verification of a Reliable Computing Platform For Real-Time Control (Phase 1 Results)*. NASA Technical Memorandum 102716, Oct. 1990.
- [2] Di Vito, Ben L.; Butler, Ricky W.; and Caldwell, James L.: High Level Design Proof of a Reliable Computing Platform. In *2nd IFIP Working Conference on Dependable Computing for Critical Applications*, Tucson, AZ, Feb. 1991, pp. 124-136.
- [3] Butler, Ricky W.; Caldwell, James L.; and Di Vito, Ben L.: Design Rationale for a Formally Verified Reliable Computing Platform. In *6th Annual Conference on Computer Assurance (COMPASS 91)*, Gaithersburg, MD, June 1991.
- [4] von Henke, F. W.; Crow, J. S.; Lee, R.; Rushby, J. M.; and Whitehurst, R. A.: EHDM Verification Environment: An Overview. In *11th National Computer Security Conference*, Baltimore, Maryland, 1988.
- [5] Computer Resource Management, Inc.: Chapter 14: High Energy Radio Frequency Fields. In *Digital Systems Validation Handbook - volume II*, no. DOT/FAA/CT-88/10, pp. 14.1-14.50. FAA, Feb. 1989.
- [6] Federal Aviation Administration. *System Design Analysis*, September 7, 1982. Advisory Circular 25.1309-1.
- [7] Butler, Ricky W.; and Finelli, George B.: The Infeasibility of Experimental Quantification of Life-Critical Software Reliability. In *Proceedings of the ACM SIGSOFT '91 Conference on Software for Critical Systems*, New Orleans, Louisiana, Dec. 1991, pp. 66-76.
- [8] Rushby, John; and von Henke, Friedrich: *Formal Verification of a Fault-Tolerant Clock Synchronization Algorithm*. NASA Contractor Report 4239, June 1989.
- [9] Lamport, Leslie; and Melliar-Smith, P. M.: Synchronizing Clocks in the Presence of Faults. *Journal of the ACM*, vol. 32, no. 1, Jan. 1985, pp. 52-78.
- [10] Siewiorek, Daniel P.; and Swarz, Robert S.: *The Theory and Practice of Reliable System Design*. Digital Press, 1982.

- [11] Goldberg, Jack; et al.: *Development and Analysis of the Software Implemented Fault-Tolerance (SIFT) Computer*. NASA Contractor Report 172146, 1984.
- [12] Hopkins, Albert L., Jr.; Smith, T. Basil, III; and Lala, Jaynarayan H.: FTMP — A Highly Reliable Fault-Tolerant Multiprocessor for Aircraft. *Proceedings of the IEEE*, vol. 66, no. 10, Oct. 1978, pp. 1221-1239.
- [13] Lala, Jaynarayan H.; Alger, L. S.; Gauthier, R. J.; and Dzwonczyk, M. J.: *A Fault-Tolerant Processor to Meet Rigorous Failure Requirements*. Charles Stark Draper Lab., Inc., Technical Report CSDL-P-2705, July 1986.
- [14] Walter, C. J.; Kieckhafer, R. M.; and Finn, A. M.: MAFT: A Multicomputer Architecture for Fault-Tolerance in Real-Time Control Systems. In *IEEE Real-Time Systems Symposium*, Dec. 1985.
- [15] Kopetz, H.; Damm, A.; Koza, C.; Mulazzani, M.; Schwabl, W.; Senft, C.; and Zainlinger, R.: Distributed Fault-tolerant Real-time Systems: The Mars Approach. *IEEE Micro*, vol. 9, Feb. 1989, pp. 25-40.
- [16] Moser, Louise; Melliar-Smith, Michael; and Schwartz, Richard: *Design Verification of SIFT*. NASA Contractor Report 4097, Sept. 1987.
- [17] *Peer Review of a Formal Verification/Design Proof Methodology*. NASA Conference Publication 2377, July 1983.
- [18] Lamport, Leslie; Shostak, Robert; and Pease, Marshall: The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, July 1982, pp. 382-401.
- [19] Mancini, L. V.; and Pappalardo, G.: Towards a Theory of Replicated Processing. In *Lecture Notes in Computer Science*, vol. 331, pp. 175-192. Springer Verlag, 1988.
- [20] Rushby, John: *Formal Specification and Verification of a Fault-Masking and Transient-Recovery Model for Digital Flight-Control Systems*. NASA Contractor Report 4384, July 1991.
- [21] Rushby, John: Formal Specification and Verification of a Fault-Masking and Transient-Recovery Model for Digital Flight-Control Systems. In *Proceedings of the Symposium on Formal Techniques in Real Time and Fault Tolerant Systems*, Nijmegen, The Netherlands, Jan. 1992. To appear.
- [22] Shankar, Natarajan: *Mechanical Verification of a Schematic Byzantine Clock Synchronization Algorithm*. NASA Contractor Report 4386, July 1991.
- [23] Shankar, Natarajan: Mechanical Verification of a Schematic Byzantine Fault-Tolerant Clock Synchronization Algorithm. In *Proceedings of the Symposium on Formal Techniques in Real Time and Fault Tolerant Systems*, Nijmegen, The Netherlands, Jan. 1992. To appear.

- [24] Harel, D.; Lachover, H.; Naamad, A.; Pnueli, A.; Politi, M.; Sherman, R.; Shtull-Trauring, A.; and Trakhtenbrot, M.: STATEMATE: A Working Environment for the Development of Complex Reactive Systems. *IEEE Transactions on Software Engineering*, vol. 16, no. 4, Apr. 1990, pp. 403-414.
- [25] Clarke, E.M.; Emerson, E.A.; and Sistla, A.P.: Automatic Verification of Finite-State Concurrent Systems using Temporal Logic Specifications. *ACM Transactions on Programming Languages and Systems*, vol. 8, no. 2, Apr. 1986, pp. 244-263.
- [26] Miner, Paul S.; and Caldwell, James L.: A HOL Theory for Voting. In *NASA Formal Methods Workshop 1990*, NASA CP-10052, Nov. 1990, pp. 442-456.
- [27] Bevier, William R.; and Young, William D.: *Machine Checked Proofs of the Design and Implementation of a Fault-Tolerant Circuit*. NASA Contractor Report 182099, Nov. 1990.
- [28] Bevier, William R.; and Young, William D.: The Proof of Correctness of a Fault-Tolerant Circuit Design. In *Second IFIP Conference on Dependable Computing For Critical Applications*, Tucson, Arizona, Feb. 1991, pp. 107-114.
- [29] Hunt, Jr., Warren A.: Microprocessor Design Verification. *Journal of Automated Reasoning*, no. 4, 1989, pp. 429-260.
- [30] Srivas, Mandayam; and Bickford, Mark: *Verification of the FtCayuga Fault-Tolerant Microprocessor System (Volume 1: A Case Study in Theorem Prover-Based Verification)*. NASA Contractor Report 4381, July 1991.
- [31] Pan, Jing; Levitt, Karl; and Cohen, Gerald C.: *Toward a Formal Verification of a Floating-Point Coprocessor and its Composition with a Central Processing Unit*. NASA Contractor Report 187547, 1991.
- [32] Pan, Jing; and Levitt, Karl: Towards a Formal Specification of the IEEE Floating-Point Standard with Application to the Verification of Floating-Point Coprocessors. In *24th Asilomar Conference on Signals, Systems & Computers*, Monterrey, CA., Nov. 1990.
- [33] Kalvala, Sara; Levitt, Karl; and Cohen, Gerald C.: Design and Verification of a DMA Processor. To be published as a NASA Contractor Report, 1991.
- [34] Windley, Phil J.; Levitt, Karl; and Cohen, Gerald C.: *Formal Proof of the AVM-1 Microprocessor Using the Concept of Generic Interpreters*. NASA Contractor Report 187491, Mar. 1991.
- [35] Windley, Phil J.; Levitt, Karl; and Cohen, Gerald C.: *The Formal Verification of Generic Interpreters*. NASA Contractor Report 4403, Oct. 1991.
- [36] Windley, Phil J.: Abstract Hardware. In *ACM International Workshop on Formal Methods in VLSI Design*, Miami, FL, Jan. 1991.

- [37] Schubert, Thomas; Levitt, Karl; and Cohen, Gerald C.: *Formal Verification of a Set of Memory Management Units*. NASA Contractor Report 189566, 1992.
- [38] Schubert, Thomas; and Levitt, Karl: Verification of Memory Management Units. In *Second IFIP Conference on Dependable Computing For Critical Applications*, Tucson, Arizona, Feb. 1991, pp. 115-123.
- [39] Schubert, Thomas; Levitt, Karl; and Cohen, Gerald C.: *Towards Composition of Verified Hardware Devices*. NASA Contractor Report 187504, 1991.
- [40] Pan, Jing; Levitt, Karl; and Schubert, E. Thomas: Toward a Formal Verification of a Floating-Point Coprocessor and its Composition with a Central Processing Unit. In *ACM International Workshop on Formal Methods in VLSI Design*, Miami, FL, Jan. 1991.
- [41] Kalvala, Sara; Archer, Myla; and Levitt, Karl: A Methodology for Integrating Hardware Design and Verification. In *ACM International Workshop on Formal Methods in VLSI Design*, Miami, FL, Jan. 1991.
- [42] Welch, J. Lundelius; and Lynch, Nancy A.: A New Fault-tolerant Algorithm For Clock Synchronization. *Information and Computation*, vol. 77, no. 1, Apr. 1988, pp. 1-35.

Index

The following index identifies where each symbol or identifier is introduced in the main body of the report. Multiple entries appear for those names used in more than one module in the EHDM specifications.

$C(t)$ 33
 $T \in R^{(i)}$ 33
 $T^{(i)}$ 33
 $\Delta_p^{(i-1)}$ 33
 $\Delta_{qp}^{(i)}$ 34
 Σ 34
 δ 34
 δ_0 34
 ϵ 34
 ν 39
 ρ 33
 $c(T)$ 33
 f_k 19
 f_s 15
 f_t 19
 f_v 15
 m 34
 $rt^{(i)}(T)$ 33
 rt_p 33
 \mathcal{N}_{da}^a 38
 \mathcal{N}_{da}^b 40
 \mathcal{N}_{da}^c 38
 \mathcal{N}_{da}^s 40
 \mathcal{N}_{da}^v 40
 \mathcal{N}_{ds} 26
 \mathcal{N}_{ds}^b 27
 \mathcal{N}_{ds}^c 27
 \mathcal{N}_{ds}^s 28
 \mathcal{N}_{ds}^v 28
 \mathcal{N}_{rs} 15
 \mathcal{N}_{us} 14
A0 34
Corr 33
DA 9
DAmap 42
DAstate 36
DS 9
DSmap 29
DSstate 25
ELT 44
MB 13
MBvec 13
Pstate 13
RS 8
RSmap 22
RSstate 15
S1 34
S1A 35
S2 34
Theorem_1 35
Theorem_2 35
US 8
all_durations 39
allowable_faults 16
broadcast_duration 39
broadcast_duration2 39
broadcast_received 27
broadcast_received 38
cell 18
cell_recovered 20
cell_recovery 23
cell_state 18
clock_advanced 39
com_broadcast_2 44
com_broadcast_5 44
components_equal 20
consensus_prop 24
control_recovered 20
control_recovery 23
control_state 18
da_proc_state 36
da_proc_state 36
da_rt 36
da_rt_lem 37
dep 19
dep_agree 19

dep_recovery 20
ds_proc_array 25
ds_proc_state 25
duration 39
enough_clocks 35
enough_hardware 38
frame_N_ds 26
frame_commutes 22
frame_commutes 29
frame_time 33
full_recovery 20
good_clock 33
good_values_sent 16
initial_Corr 33
initial_da 41
initial_ds 28
initial_maj_cond 24
initial_maps 22
initial_maps 29
initial_maps 42
initial_recovery 20
initial_rs 17
initial_us 14
inputs 13
maj 22
maj_ax 22
maj_condition 16
maj_working 17
nonfaulty_clock 33
nrep 13
num_good_clocks 35
outputs 13
phase_com_l' 43
phase_commutes 42
phases 25
pos_durations 39
processors_exist_ax 13
reachable 22
rec 19
rec_maj_f_c 24
rs_proc_state 15
ss_update 29
ss_update 42
state_invariant 23
state_rec_inv 23
state_recovery 23
succ 19
succ_ax 19
sync_thm 36
sync_time 34
vote_maj 20
voted_final_state 16
working_majority 24
working_proc 17
working_set 17

Appendix A

LaTeX-printed Specification Listings

The following specifications were formatted with the assistance of the EIIDM latex-printer.

US: Module

Using generic_FT

Exporting all

Theory

```
s, t: Var Pstate
u: Var inputs
N_u: Definition function[Pstate, Pstate, inputs → bool] =
  (λ s, t, u : t = f_c(u, s))
initial_us: function[Pstate → bool] = (λ s : s = initial_proc_state)
```

End

generic_FT: **Module**

Using rcp_defs, sets[processors], cardinality[processors]

Exporting all with rcp_defs, sets[processors], cardinality[processors]

Theory

```
us, ps, X, Y: Var Pstate
p, i, j: Var processors
k, l, q: Var nat
u: Var inputs
w: Var MBvec
h: Var MBmatrix
A, B: Var set[processors]
maj_condition: function[set[processors] → bool] =
  (λ A : 2 * card(A) > card(fullset[processors]))
(* The following definitions and axioms are used to model a general class
of fault-tolerant computation schemes. The elaboration of these
uninterpreted functions, as well as those in rcp_defs, would be made
for a particular choice of application-dependent computation style and
voting pattern. Given suitable choices, the axioms can then be shown
to be theorems. *)
```

```
control_state: Type
cell: Type
cell_state: Type
c, d, e: Var cell
K: Var control_state
H: Var nat
```

```
succ: function[control_state → control_state]
```

f_k : function[Pstate \rightarrow control_state]
 f_t : function[Pstate, cell \rightarrow cell_state]

 f_c : function[inputs, Pstate \rightarrow Pstate]
 f_a : function[Pstate \rightarrow outputs] (* actuator output *)
 f_s : function[Pstate \rightarrow MB]
 f_v : function[Pstate, MBvec \rightarrow Pstate]

(* rec(c,K,H) = T iff cell c's state should have been recovered when in control state K with healthy count H; note that H-1 healthy frames will have occurred previously. *)

rec: function[cell, control_state, nat \rightarrow bool]

(* dep(c,d,K) = T iff cell c's value in the next state depends on cell d's value in the current state, when in control state K; if cell c is voted during K, or its computation takes only sensor inputs, there is no dependency; if c is not computed during K, c depends only on itself; otherwise, c depends on one or more cells for its new value. *)

dep: function[cell, cell, control_state \rightarrow bool]

dep_agree: function[cell, control_state, Pstate, Pstate \rightarrow bool] =
 $(\lambda c, K, X, Y : (\forall d : \text{dep}(c, d, K) \supset f_t(X, d) = f_t(Y, d)))$

(* Axioms to be satisfied by the generic application *)

succ_ax: Axiom $f_k(f_c(u, ps)) = \text{succ}(f_k(ps))$

full_recovery: Axiom $H \geq \text{recovery_period} \supset \text{rec}(c, K, H)$

initial_recovery: Axiom $\text{rec}(c, K, H) \supset H > 2$

dep_recovery: Axiom $\text{rec}(c, \text{succ}(K), H + 1) \wedge \text{dep}(c, d, K) \supset \text{rec}(d, K, H)$

components_equal: Axiom $f_k(X) = f_k(Y) \wedge (\forall c : f_t(X, c) = f_t(Y, c)) \supset X = Y$

control_recovered: Axiom
maj_condition(A) $\wedge (\forall p : p \in A \supset w(p) = f_s(ps)) \supset f_k(f_v(Y, w)) = f_k(ps)$

cell_recovered: Axiom
maj_condition(A)
 $\wedge (\forall p : p \in A \supset w(p) = f_s(f_c(u, ps)))$
 $\wedge f_k(X) = K \wedge f_k(ps) = K \wedge \text{dep_agree}(c, K, X, ps)$
 $\supset f_t(f_v(f_c(u, X), w), c) = f_t(f_c(u, ps), c)$

vote_maj: Axiom maj_condition(A) $\wedge (\forall p : p \in A \supset w(p) = f_s(ps))$
 $\supset f_v(ps, w) = ps$

(* Lemmas pertaining to sets and cardinalities *)

card_fullset: Lemma $\text{card}(\text{fullset}[\text{processors}]) > 0$

proc_extensionality: Lemma $(\forall p : p \in A = p \in B) \supset (A = B)$

Proof

discharge_finite: Prove

finite[processors] {f \leftarrow ($\lambda p \rightarrow \text{nat} : p$), N \leftarrow nrep}

nat_nit: Sublemma $k > 0 \Leftrightarrow k \neq 0$

p_nat_nit: Prove nat_nit

p_card_fullset: Prove card_fullset from
empty {a ← fullset[processors]},
card_empty {a ← fullset[processors]},
nat_nit {k ← card(fullset[processors])}

p_proc_extensionality: Prove proc_extensionality {p ← x@p1} from
extensionality {a ← A, b ← B}

End

RS: Module

Using generic_FT

Exporting all with generic_FT

Theory

rs_proc_state: Type = Record healthy : nat,
proc_state : Pstate
end record

RSstate: Type = array [processors] of rs_proc_state

rs0: RSstate

rsproc0: rs_proc_state

s, t: Var RSstate

u: Var inputs

w: Var MBvec

h: Var MBmatrix

p, q: Var processors

k: Var nat

A: Var set[processors]

working_proc: function[RSstate, processors → bool] =

(λ s, p : s(p).healthy ≥ recovery_period)

working_set: function[RSstate → set[processors]] =

(λ s : (λ p : working_proc(s, p)))

maj_working: function[RSstate → bool] =

(λ t : maj.condition(working_set(t)))

allowable_faults: function[RSstate, RSstate → bool] =

(λ s, t : maj_working(t)

∧ (∀ p : t(p).healthy > 0 ⊃ t(p).healthy = 1 + s(p).healthy))

good_values_sent: function[RSstate, inputs, MBvec → bool] =

(λ s, u, w : (∀ q :

s(q).healthy > 0 ⊃ w(q) = f_s(f_c(u, s(q).proc_state))))

voted_final_state: function[RSstate, RSstate, inputs, MBmatrix, processors
→ bool] =

(λ s, t, u, h, p : t(p).proc_state = f_v(f_c(u, s(p).proc_state), h(p)))

N_rs: Definition function[RSstate, RSstate, inputs → bool] =

```

(λ s, t, u : (∃ h :
  (∀ p :
    s(p).healthy > 0
    ⊃ good_values_sent(s, u, h(p))
    ∧ voted_final_state(s, t, u, h, p)))
  ∧ allowable_faults(s, t))
initial_rs: function[RSstate → bool] =
(λ s : (∀ p :
  s(p).healthy = recovery_period
  ∧ s(p).proc_state = initial_proc_state))

```

Proof

End

RS_to_US: Module

Using RS, US, RS_majority

Exporting all with RS, US, RS_majority

Theory

```

rs, s, t, x, y, z: Var RSstate
us, ps, X, Y: Var Pstate
p, i, j: Var processors
k, l, q: Var nat
u: Var inputs
w: Var MBvec
h: Var MBmatrix
MBmatrix0: MBmatrix
MBcons_fn: Type is function[processors → MBvec]
MBfn: Var MBcons_fn
RSstate_prop: Type is function[RSstate → bool]
rs_prop: Var RSstate_prop

```

```

RSmap: function[RSstate → Pstate] = (λ rs : maj(rs))
rs_measure: function[RSstate, nat → nat] == (λ rs, k : k)
reachable_in_n: function[RSstate, nat → bool] =
(λ t, k : if k = 0
  then initial_rs(t)
  else (∃ s, u : reachable_in_n(s, k - 1) ∧  $\mathcal{N}_{rs}(s, t, u)$ )
  end if) by rs_measure
reachable: function[RSstate → bool] = (λ t : (∃ k : reachable_in_n(t, k)))

frame_commutates: Theorem reachable(s) ∧  $\mathcal{N}_{rs}(s, t, u)$  ⊃  $\mathcal{N}_{us}(RSmap(s), RSmap(t), u)$ 

initial_maps: Theorem initial_rs(s) ⊃ initial_us(RSmap(s))

```

End

RS_majority: Module

Using US, RS, nat_inductions

Exporting all

Theory

```

k: Var nat

```

```

p: Var processors
us: Var Pstate
rs: Var RSstate
A: Var set[processors]

maj_exists: function[RSstate → bool] =
  (λ rs : (∃ A, us :
    maj_condition(A) ∧ (∀ p : p ∈ A ⊃ rs(p).proc_state = us)))
maj: function[RSstate → Pstate]

maj_ax: Axiom (∃ A :
  maj_condition(A) ∧ (∀ p : p ∈ A ⊃ rs(p).proc_state = us)
  ⊃ maj(rs) = us)

```

End

RS_lemmas: Module

Using RS_to_US

Exporting all with RS_to_US

Theory

```

rs, s, t, x, y, z: Var RSstate
us: Var Pstate
p, i, j: Var processors
k, l, q: Var nat
u: Var inputs
w: Var MBvec
h: Var MBmatrix
MBmatrix0: MBmatrix
MBcons.fn: Type is function[processors → MBvec]
MBfn: Var MBcons.fn
RSstate_prop: Type is function[RSstate → bool]
rs.prop: Var RSstate_prop
m, n, a, b: Var proc_plus
prop: Var function[proc_plus → bool]
c, d, e: Var cell
K: Var control_state
H: Var nat
A: Var set[processors]

initial_maj: Lemma
  initial_rs(s) ⊃ (∀ p : maj_exists(s) ∧ s(p).proc_state = maj(s))

initial_working: Lemma initial_rs(s) ⊃ working_set(s) = fullset[processors]

initial_maj_cond: Lemma initial_rs(s) ⊃ maj_condition(working_set(s))

control_recovery: function[RSstate → bool] =
  (λ s : (∀ p : s(p).healthy > 1 ⊃ fk(s(p).proc_state) = fk(maj(s))))
cell_recovery: function[RSstate → bool] =
  (λ s : (∀ p, c :
    rec(c, fk(s(p).proc_state), s(p).healthy)
    ⊃ fi(s(p).proc_state, c) = fi(maj(s), c)))
state_recovery: function[RSstate → bool] =
  (λ s : maj_exists(s) ∧ control_recovery(s) ∧ cell_recovery(s))
working_majority: function[RSstate → bool] =
  (λ s : (∀ p : p ∈ working_set(s) ⊃ s(p).proc_state = maj(s)))

```

consensus_prop: Lemma $\text{state_recovery}(s) \supset \text{working_majority}(s)$
working_set_healthy: Lemma $\text{working_set}(s)(p) \supset s(p).\text{healthy} > 0$
maj_sent: Lemma $\text{state_recovery}(s) \wedge \text{good_values_sent}(s, u, w)$
 $\supset (\forall p : p \in \text{working_set}(s) \supset w(p) = f_s(f_c(u, \text{maj}(s))))$
rec_maj_exists: Lemma
 $\text{maj_working}(s) \wedge \text{state_recovery}(s) \wedge \mathcal{N}_{r,s}(s, t, u) \supset \text{maj_exists}(t)$
rec_maj_f.c: Lemma
 $\text{maj_working}(s) \wedge \text{state_recovery}(s) \wedge \mathcal{N}_{r,s}(s, t, u) \supset \text{maj}(t) = f_c(u, \text{maj}(s))$

End

RS_invariants: **Module**

Using RS_lemmas, nat_inductions

Exporting all with RS_lemmas

Theory

rs, s, t, x, y, z : **Var** RSstate
 us : **Var** Pstate
 p, i, j : **Var** processors
 k, l, q : **Var** nat
 u : **Var** inputs
 w : **Var** MBvec
 h : **Var** MBmatrix
RSstate_prop: **Type is function**[RSstate \rightarrow bool]
rs_prop: **Var** RSstate_prop
 m, n, a, b : **Var** proc_plus
prop: **Var function**[proc_plus \rightarrow bool]
 c, d, e : **Var** cell
 K : **Var** control_state
 H : **Var** nat
 A : **Var** set[processors]
state_invariant: **function**[RSstate_prop \rightarrow bool] =
 $(\lambda rs_prop : (\forall t : \text{reachable}(t) \supset rs_prop(t)))$
state_induction: **Lemma**
 $(\forall x : \text{initial_rs}(x) \supset rs_prop(x))$
 $\wedge (\forall s, t, u : \text{reachable}(s) \wedge rs_prop(s) \wedge \mathcal{N}_{r,s}(s, t, u) \supset rs_prop(t))$
 $\supset \text{state_invariant}(rs_prop)$
maj_working_inv: **Lemma** $\text{state_invariant}(\text{maj_working})$
state_rec_inv: **Lemma** $\text{state_invariant}(\text{state_recovery})$

Proof

state_invariant_to_n: **function**[RSstate_prop, nat \rightarrow bool] =
 $(\lambda rs_prop, k : (\forall t : \text{reachable_in_n}(t, k) \supset rs_prop(t)))$
base_state_ind: **Lemma**
 $(\text{initial_rs}(x) \supset rs_prop(x)) \supset (\text{reachable_in_n}(x, 0) \supset rs_prop(x))$
ind_state_ind: **Lemma**
 $(\forall s, t, u : \text{reachable}(s) \wedge rs_prop(s) \wedge \mathcal{N}_{r,s}(s, t, u) \supset rs_prop(t))$
 $\supset (\forall k : \text{state_invariant_to_n}(rs_prop, k)$
 $\supset \text{state_invariant_to_n}(rs_prop, k + 1))$

p_base_state_ind: Prove base_state_ind from
 reachable.in.n $\{t \leftarrow x, k \leftarrow 0\}$

p_ind_state_ind: Prove
 ind_state_ind $\{s \leftarrow s@p3, t \leftarrow t@p2, u \leftarrow u@p3\}$ from
 state_invariant_to.n $\{k \leftarrow k, t \leftarrow s@p3\}$,
 state_invariant_to.n $\{k \leftarrow k + 1, t \leftarrow t\}$,
 reachable.in.n $\{t \leftarrow t, k \leftarrow k + 1\}$,
 reachable $\{t \leftarrow s@p3, k \leftarrow k\}$

p_state_induction: Prove
 state_induction
 $\{x \leftarrow t@p3,$
 $s \leftarrow s@p4,$
 $t \leftarrow t@p4,$
 $u \leftarrow u@p4\}$ from
 nat_induction
 $\{p \leftarrow (\lambda k : \text{state_invariant_to.n}(rs_prop, k)),$
 $n_2 \leftarrow k@p7\}$,
 base_state_ind $\{x \leftarrow t@p3\}$,
 state_invariant_to.n $\{t \leftarrow x, k \leftarrow 0\}$,
 ind_state_ind $\{k \leftarrow n_1@p1\}$,
 state_invariant_to.n $\{t \leftarrow t@p6, k \leftarrow k@p7\}$,
 state_invariant,
 reachable $\{t \leftarrow t@p6\}$

maj_working_inv_l1: Lemma initial_rs(s) \supset maj_working(s)

maj_working_inv_l2: Lemma $\mathcal{N}_{rs}(s, t, u) \supset$ maj_working(t)

p_maj_working_inv_l1: Prove maj_working_inv_l1 from
 maj_working $\{t \leftarrow s\}$, initial_maj_cond

p_maj_working_inv_l2: Prove maj_working_inv_l2 from \mathcal{N}_{rs} , allowable_faults

p_maj_working_inv: Prove maj_working_inv from
 state_induction $\{rs_prop \leftarrow \text{maj_working}\}$,
 maj_working_inv_l1 $\{s \leftarrow x@p1\}$,
 maj_working_inv_l2 $\{s \leftarrow s@p1, t \leftarrow t@p1, u \leftarrow u@p1\}$

state_rec_inv_l1: Lemma initial_rs(s) \supset state_recovery(s)

state_rec_inv_l2: Lemma
 maj_working(s) \wedge state_recovery(s) \wedge $\mathcal{N}_{rs}(s, t, u) \wedge$ maj(t) = $f_c(u, \text{maj}(s))$
 \supset control_recovery(t)

state_rec_inv_l3: Lemma
 maj_working(s) \wedge state_recovery(s)
 \wedge maj(t) = $f_c(u, \text{maj}(s))$
 \wedge $t(p).healthy = 1 + s(p).healthy$
 \wedge $f_k(s(p).proc_state) = f_k(\text{maj}(s))$
 \wedge $f_k(t(p).proc_state) = f_k(\text{maj}(t))$
 \wedge good_values_sent(s, u, h(p))
 \wedge rec(c, $f_k(t(p).proc_state)$, $t(p).healthy$)
 \supset $f_i(f_v(f_c(u, s(p).proc_state), h(p)), c) = f_i(f_c(u, \text{maj}(s)), c)$

state_rec_inv_l4: Lemma
 maj_working(s) \wedge state_recovery(s)
 \wedge $\mathcal{N}_{rs}(s, t, u) \wedge$ maj(t) = $f_c(u, \text{maj}(s)) \wedge$ control_recovery(t)
 \supset cell_recovery(t)

state_rec_inv_l5: Lemma

$\text{reachable}(s) \wedge \text{state_recovery}(s) \wedge \mathcal{N}_{r,s}(s, t, u) \supset \text{state_recovery}(t)$

p_state_rec_inv_l1: Prove state_rec_inv_l1 from

control_recovery,
cell_recovery,
state_recovery,
initial_maj $\{p \leftarrow p@p1\}$,
initial_maj $\{p \leftarrow p@p2\}$

p_state_rec_inv_l2: Prove state_rec_inv_l2 from

control_recovery $\{s \leftarrow t\}$,
 $\mathcal{N}_{r,s} \{p \leftarrow p@p1\}$,
control_recovered
 $\{ps \leftarrow f_c(u, \text{maj}(s)),$
 $A \leftarrow \text{working_set}(s),$
 $w \leftarrow ((h@p2)p@p1),$
 $Y \leftarrow f_c(u, (s(p@p1)).\text{proc_state})\}$,
maj_sent $\{p \leftarrow p@p3, w \leftarrow ((h@p2)p@p1)\}$,
maj_working $\{t \leftarrow s\}$,
state_recovery,
control_recovery $\{p \leftarrow p@p1\}$,
voted_final_state $\{h \leftarrow h@p2, p \leftarrow p@p1\}$,
allowable_faults $\{p \leftarrow p@p1\}$

p_state_rec_inv_l3: Prove state_rec_inv_l3 from

dep_agree $\{K \leftarrow f_k(\text{maj}(s)), X \leftarrow s(p).\text{proc_state}, Y \leftarrow \text{maj}(s)\}$,
cell_recovered
 $\{ps \leftarrow \text{maj}(s),$
 $w \leftarrow h(p),$
 $X \leftarrow s(p).\text{proc_state},$
 $A \leftarrow \text{working_set}(s),$
 $K \leftarrow f_k(\text{maj}(s))\}$,
maj_sent $\{p \leftarrow p@p2, w \leftarrow h(p)\}$,
maj_working $\{t \leftarrow s\}$,
state_recovery,
cell_recovery $\{p \leftarrow p, c \leftarrow d@p1\}$,
dep_recovery $\{d \leftarrow d@p1, K \leftarrow f_k(\text{maj}(s)), H \leftarrow s(p).\text{healthy}\}$,
succ_ax $\{ps \leftarrow \text{maj}(s)\}$

p_state_rec_inv_l4: Prove state_rec_inv_l4 from

cell_recovery $\{s \leftarrow t\}$,
 $\mathcal{N}_{r,s} \{p \leftarrow p@p1\}$,
state_rec_inv_l3 $\{p \leftarrow p@p1, h \leftarrow h@p2, c \leftarrow c@p1\}$,
state_recovery,
control_recovery $\{p \leftarrow p@p1\}$,
control_recovery $\{s \leftarrow t, p \leftarrow p@p1\}$,
voted_final_state $\{h \leftarrow h@p2, p \leftarrow p@p1\}$,
allowable_faults $\{p \leftarrow p@p1\}$,
initial_recovery
 $\{c \leftarrow c@p1,$
 $H \leftarrow (t(p@p1)).\text{healthy},$
 $K \leftarrow f_k((t(p@p1)).\text{proc_state})\}$,
succ_ax $\{ps \leftarrow \text{maj}(s)\}$

p_state_rec_inv_l5: Prove state_rec_inv_l5 from

```

state_rec_inv_l2,
rec_maj_exists,
rec_maj_f.c,
state_rec_inv_l4,
state_recovery {s ← t},
maj_working_inv,
state_invariant {rs_prop ← maj_working, t ← s}

```

```

p.state_rec_inv: Prove state_rec_inv from
state_induction {rs_prop ← state_recovery},
state_rec_inv_l1 {s ← x@p1},
state_rec_inv_l5 {s ← s@p1, t ← t@p1, u ← u@p1}

```

End

RS_top_proof: **Module**

Using RS_invariants

Exporting all

Theory

```

rs, s, t, x, y, z: Var RSstate
us: Var Pstate
p, i, j: Var processors
k, l, q: Var nat
u: Var inputs
w: Var MBvec
h: Var MBmatrix
c, d, e: Var cell
K: Var control.state
H: Var nat
A: Var set[processors]
MBmatrix0: MBmatrix
MBcons.fn: Type is function[processors → MBvec]
MBfn: Var MBcons.fn
RSstate_prop: Type is function[RSstate → bool]
rs_prop: Var RSstate_prop
m, n, a, b: Var proc.plus
prop: Var function[proc.plus → bool]

```

Proof

```

p.frame_commutates: Prove frame_commutates from
N_us {s ← maj(s), t ← maj(t)},
rec_maj_f.c,
consensus_prop,
maj_working_inv,
state_invariant {rs_prop ← maj_working, t ← s},
state_rec_inv,
state_invariant {rs_prop ← state_recovery, t ← s},
state_recovery,
RSmap {rs ← s},
RSmap {rs ← t}

```

```

p.initial_maps: Prove initial_maps from
maj_ax {A ← working_set(s), rs ← s, us ← initial_proc_state},
initial_us {s ← RSmap(s)},
initial_rs {p ← p@p1},
RSmap {rs ← s},

```

initial_maj_cond

p_initial_working: Prove initial_working from
extensionality $\{a \leftarrow \text{working_set}(s), b \leftarrow \text{fullset}[\text{processors}]\}$,
initial_rs $\{p \leftarrow x@p1\}$,
working_set $\{p \leftarrow x@p1\}$,
working_proc $\{p \leftarrow x@p1\}$

p_initial_maj_cond: Prove initial_maj_cond from
maj_condition $\{A \leftarrow \text{working_set}(s)\}$, initial_working, card_fullset

p_initial_maj: Prove initial_maj from
maj_ax
 $\{rs \leftarrow s,$
 $A \leftarrow \text{fullset}[\text{processors}],$
 $us \leftarrow \text{initial_proc_state}\}$,
maj_exists
 $\{rs \leftarrow s,$
 $A \leftarrow \text{fullset}[\text{processors}],$
 $us \leftarrow \text{initial_proc_state}\}$,
maj_condition $\{A \leftarrow \text{fullset}[\text{processors}]\}$,
initial_rs $\{p \leftarrow p@p1\}$,
initial_rs $\{p \leftarrow p@p2\}$,
initial_rs,
card_fullset

p_working_set_healthy: Prove working_set_healthy from
working_set, working_proc, recovery_period_ax

p_consensus_prop: Prove consensus_prop from
working_majority,
components_equal $\{X \leftarrow (s(p@p1)).\text{proc_state}, Y \leftarrow \text{maj}(s)\}$,
control_recovery $\{p \leftarrow p@p1\}$,
cell_recovery $\{p \leftarrow p@p1, c \leftarrow c@p2\}$,
full_recovery
 $\{c \leftarrow c@p2,$
 $K \leftarrow f_k((s(p@p1)).\text{proc_state}),$
 $H \leftarrow (s(p@p1)).\text{healthy}\}$,
state_recovery,
working_set $\{p \leftarrow p@p1\}$,
working_proc $\{p \leftarrow p@p1\}$,
recovery_period_ax

p_maj_sent: Prove maj_sent from
good_values_sent $\{q \leftarrow p\}$,
consensus_prop,
working_majority,
working_set_healthy

p_rec_maj_exists: Prove rec_maj_exists from

```

maj_exists {rs ← t, A ← working_set(s), us ← f_c(u, maj(s))},
N_rs {p ← p@p1},
vote_maj
  {ps ← f_c(u, maj(s)),
   w ← ((h@p2)p@p1),
   A ← working_set(s)},
maj_sent {p ← p@p3, w ← ((h@p2)p@p1)},
state_recovery,
consensus_prop,
working_majority {p ← p@p1},
voted_final_state {h ← h@p2, p ← p@p1},
working_set_healthy {p ← p@p1},
maj_working {t ← s}

```

p-rec-maj-f.c: Prove rec-maj-f.c from

```

maj_ax {rs ← t, A ← working_set(s), us ← f_c(u, maj(s))},
N_rs {p ← p@p1},
vote_maj
  {ps ← f_c(u, maj(s)),
   w ← ((h@p2)p@p1),
   A ← working_set(s)},
maj_sent {p ← p@p3, w ← ((h@p2)p@p1)},
state_recovery,
consensus_prop,
working_majority {p ← p@p1},
voted_final_state {h ← h@p2, p ← p@p1},
working_set_healthy {p ← p@p1},
maj_working {t ← s}

```

End

RS_tcc_proof: **Module**

Using rcp-defs.tcc

Exporting all

Theory

Proof

proc_plus_TCC1_PROOF: Prove proc_plus_TCC1 {p ← 0}

processors_TCC1_PROOF: Prove processors_TCC1 {p ← nrep} from
processors_exist_ax

End

RS_to_US_tcc: **Module**

Using RS_to_US

Exporting all with RS_to_US

Theory

s: **Var** RS.RSstate
t: **Var** RS.RSstate
k: **Var** naturalnumber

reachable_in_n_TCC1: **Formula** $(\neg(k = 0)) \supset (k - 1 \geq 0)$

reachable.in.n.TCC2: **Formula**
 $(\neg(k = 0)) \supset \text{rs_measure}(t, k) > \text{rs_measure}(s, k - 1)$

Proof

reachable.in.n.TCC1.PROOF: **Prove** reachable.in.n.TCC1

reachable.in.n.TCC2.PROOF: **Prove** reachable.in.n.TCC2

End RS.to.US.tcc

DS: Module

Using generic.FT

Exporting all with generic.FT

Theory

```

ds_proc_state: Type = Record healthy : nat,
                    proc_state : Pstate,
                    mailbox : MBvec
                    end record
ds_proc_array: Type = array [processors] of ds_proc_state
DSstate: Type = Record phase : phases,
                    proc : ds_proc_array
                    end record

ds0: DSstate
dsproc0: ds_proc_state
s, t, x, y, z: Var DSstate
u: Var inputs
w: Var MBvec
i, j, p, q, qq: Var processors
k: Var nat
ph: Var phases
A: Var set[processors]

working_proc: function[DSstate, processors → bool] =
  (λ s, p : s.proc(p).healthy ≥ recovery-period)
working_set: function[DSstate → set[processors]] =
  (λ s : (λ p : working_proc(s, p)))
maj_working: function[DSstate → bool] =
  (λ t : maj_condition(working_set(t)))

allowable_faults: function[DSstate, DSstate → bool] =
  (λ s, t : maj_working(t)
    ∧ (∀ i : t.proc(i).healthy > 0
      ⊃ t.proc(i).healthy = 1 + s.proc(i).healthy))
broadcast_received: function[DSstate, DSstate, processors → bool] =
  (λ s, t, p : (∀ qq :
    s.proc(qq).healthy > 0
    ⊃ t.proc(p).mailbox(qq) = s.proc(qq).mailbox(qq)))

 $\mathcal{N}_{ds}^c$ : function[DSstate, DSstate, inputs, processors → bool] =
  (λ s, t, u, i :
    s.proc(i).healthy > 0
    ⊃ t.proc(i).proc_state = fc(u, s.proc(i).proc_state)
    ∧ t.proc(i).mailbox(i) = fs(fc(u, s.proc(i).proc_state)))

```

```

 $\mathcal{N}_{ds}^b$ : function[DSstate, DSstate, processors  $\rightarrow$  bool] =
  (  $\lambda$  s, t, i : s.proc(i).healthy > 0
     $\supset$  t.proc(i).proc_state = s.proc(i).proc_state
     $\wedge$  broadcast_received(s, t, i))

 $\mathcal{N}_{ds}^v$ : function[DSstate, DSstate, processors  $\rightarrow$  bool] =
  (  $\lambda$  s, t, i : s.proc(i).healthy > 0
     $\supset$  (t.proc(i).mailbox = s.proc(i).mailbox
     $\wedge$  t.proc(i).proc_state
    =  $f_v$ (s.proc(i).proc_state, s.proc(i).mailbox)))

 $\mathcal{N}_{ds}^s$ : function[DSstate, DSstate, processors  $\rightarrow$  bool] =
  (  $\lambda$  s, t, i : s.proc(i).healthy > 0
     $\supset$  t.proc(i).proc_state = s.proc(i).proc_state)
   $\wedge$  (t.proc(i).healthy > 0
     $\supset$  t.proc(i).healthy = 1 + s.proc(i).healthy))

 $\mathcal{N}_{ds}$ : function[DSstate, DSstate, inputs  $\rightarrow$  bool] =
  (  $\lambda$  s, t, u : maj_working(t)
     $\wedge$  t.phase = next_phase(s.phase)
     $\wedge$  ( $\forall$  i :
      if s.phase = sync
      then  $\mathcal{N}_{ds}^s$ (s, t, i)
      else t.proc(i).healthy = s.proc(i).healthy
       $\wedge$  (s.phase = compute  $\supset$   $\mathcal{N}_{ds}^c$ (s, t, u, i))
       $\wedge$  (s.phase = broadcast  $\supset$   $\mathcal{N}_{ds}^b$ (s, t, i))
       $\wedge$  (s.phase = vote  $\supset$   $\mathcal{N}_{ds}^v$ (s, t, i))
      end if))

frame_N_ds: function[DSstate, DSstate, inputs  $\rightarrow$  bool] =
  (  $\lambda$  s, t, u : ( $\exists$  x, y, z :
     $\mathcal{N}_{ds}(s, x, u) \wedge \mathcal{N}_{ds}(x, y, u) \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u)$ ))

initial_ds: function[DSstate  $\rightarrow$  bool] =
  (  $\lambda$  s : s.phase = compute
     $\wedge$  ( $\forall$  i : s.proc(i).healthy = recovery_period
     $\wedge$  s.proc(i).proc_state = initial_proc_state))

```

End

DS.to.RS: Module

Using DS, RS

Exporting all with DS, RS

Theory

ds, s, t, x, y, z: Var DSstate

```

rs: Var RSstate
i, j: Var processors
p: Var nat
u: Var inputs
w: Var MBvec
h: Var MBmatrix
MBmatrix0: MBmatrix
MBcons.fn: Type is function[processors → MBvec]
MBfn: Var MBcons.fn
ssu_measure: function[DSstate, nat → nat] == (λ ds, p : p)
ss_update: Recursive function[DSstate, nat → RSstate] =
  (λ ds, p : if (p = 0) ∨ (p > nrep)
    then rs0
    else ss_update(ds, p - 1)
    with [(p) := rproc0
          with [(healthy) := ds.proc(p).healthy,
                (proc_state) := ds.proc(p).proc_state]]
    end if) by ssu_measure

DSmap: function[DSstate → RSstate] = (λ ds : ss_update(ds, nrep))
MBmc_measure: function[MBcons.fn, nat → nat] == (λ MBfn, p : p)
MBmatrix_cons: Recursive function[MBcons.fn, nat → MBmatrix] =
  (λ MBfn, p : if (p = 0) ∨ (p > nrep)
    then MBmatrix0
    else MBmatrix_cons(MBfn, p - 1)
    with [(p) := MBfn(p)]
    end if) by MBmc_measure

frame_commutates: Theorem
  s.phase = compute ∧ frame_N_ds(s, t, u) ⊃  $\mathcal{N}_{rs}$ (DSmap(s), DSmap(t), u)

initial_maps: Theorem initial_ds(s) ⊃ initial_rs(DSmap(s))

good_values_sent: function[DSstate, inputs, MBvec → bool] =
  (λ s, u, w : (∀ j :
    s.proc(j).healthy > 0 ⊃ w(j) = f_s(f_c(u, s.proc(j).proc_state))))

voted_final_state: function[DSstate, DSstate, inputs, MBmatrix, processors
  → bool] =
  (λ s, t, u, h, i :
    t.proc(i).proc_state = f_v(f_c(u, s.proc(i).proc_state), h(i)))

is_new_proc_state: function[DSstate, DSstate, inputs → bool] =
  (λ s, t, u : (∃ h :
    (∀ i : s.proc(i).healthy > 0
      ⊃ good_values_sent(s, u, h(i))
      ∧ voted_final_state(s, t, u, h, i))))

fr_com_1: Lemma s.phase = compute ∧ frame_N_ds(s, t, u)
  ⊃ is_new_proc_state(s, t, u) ∧ allowable_faults(s, t)

fr_com_2: Lemma is_new_proc_state(s, t, u) ∧ allowable_faults(s, t)
  ⊃  $\mathcal{N}_{rs}$ (DSmap(s), DSmap(t), u)

fc_A: Lemma s.phase = compute ∧ frame_N_ds(s, t, u)
  ⊃ is_new_proc_state(s, t, u)

fc_B: Lemma s.phase = compute ∧ frame_N_ds(s, t, u) ⊃ allowable_faults(s, t)

```

End

DS Lemmas: **Module**

Using DS.to_RS

Exporting all with DS.to_RS

Theory

```
ds: Var DSstate
rs: Var RSstate
p, q: Var nat
ph: Var phases
s, t, x, y, z: Var DSstate
i, j, jj: Var processors
u: Var inputs
w: Var MBvec
h: Var MBmatrix
MBfn: Var MBcons_fn
MB: Var MBvec
k, m, n, a, b: Var proc_plus
prop: Var function[proc_plus → bool]
```

```
half_frame_N_ds: function[DSstate, DSstate, inputs → bool] =
  (λ x, t, u : (∃ y, z :  $\mathcal{N}_{ds}(x, y, u) \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u)$ ))
```

```
quarter_frame_N_ds: function[DSstate, DSstate, inputs → bool] =
  (λ y, t, u : (∃ z :  $\mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u)$ ))
```

```
fc.A.1a: Lemma s.phase = compute ∧ frame_N_ds(s, t, u)
  ⊃ (∃ x, y, z :
    maj_working(x)
    ∧ (∀ i :
      x.phase = broadcast
      ∧ x.proc(i).healthy = s.proc(i).healthy ∧  $\mathcal{N}_{ds}^c(s, x, u, i)$ 
      ∧  $\mathcal{N}_{ds}(x, y, u) \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u)$ )
```

```
fc.A.1b: Lemma s.phase = compute ∧ frame_N_ds(s, t, u)
  ⊃ (∃ x, y, z :
    maj_working(x)
    ∧ maj_working(y)
    ∧ (∀ i :
      x.phase = broadcast
      ∧ x.proc(i).healthy = s.proc(i).healthy
      ∧  $\mathcal{N}_{ds}^c(s, x, u, i)$ 
      ∧ y.phase = next_phase(x.phase)
      ∧ y.proc(i).healthy = x.proc(i).healthy
      ∧  $\mathcal{N}_{ds}^b(x, y, i)$ 
      ∧  $\mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u)$ )
```

```
fc.A.1c: Lemma s.phase = compute ∧ frame_N_ds(s, t, u)
```


$$\begin{aligned}
& \supset (\exists x, y, z : \\
& \quad \text{maj_working}(x) \\
& \quad \wedge \text{maj_working}(y) \\
& \quad \wedge (\forall i : \\
& \quad \quad x.\text{phase} = \text{broadcast} \\
& \quad \quad \wedge x.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \\
& \quad \quad \wedge s.\text{proc}(i).\text{healthy} > 0 \\
& \quad \quad \quad \supset x.\text{proc}(i).\text{proc_state} \\
& \quad \quad \quad \quad = f_c(u, s.\text{proc}(i).\text{proc_state})) \\
& \quad \quad \wedge y.\text{phase} = \text{vote} \\
& \quad \quad \wedge y.\text{proc}(i).\text{healthy} = x.\text{proc}(i).\text{healthy} \\
& \quad \quad \wedge (x.\text{proc}(i).\text{healthy} > 0 \\
& \quad \quad \quad \supset y.\text{proc}(i).\text{proc_state} \\
& \quad \quad \quad \quad = x.\text{proc}(i).\text{proc_state} \\
& \quad \quad \quad \wedge (\forall j : \\
& \quad \quad \quad \quad x.\text{proc}(j).\text{healthy} > 0 \\
& \quad \quad \quad \quad \supset y.\text{proc}(i).\text{mailbox}(j) \\
& \quad \quad \quad \quad \quad = f_s(x.\text{proc}(j).\text{proc_state})))))) \\
& \quad \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))
\end{aligned}$$

fc.A.1d: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$

$$\begin{aligned}
& \supset (\exists x, y, z : \\
& \quad \text{maj_working}(x) \\
& \quad \wedge \text{maj_working}(y) \\
& \quad \wedge (\forall jj : \\
& \quad \quad x.\text{proc}(jj).\text{healthy} = s.\text{proc}(jj).\text{healthy} \\
& \quad \quad \wedge (s.\text{proc}(jj).\text{healthy} > 0 \\
& \quad \quad \quad \supset y.\text{proc}(jj).\text{proc_state} = x.\text{proc}(jj).\text{proc_state})) \\
& \quad \wedge (\forall i : \\
& \quad \quad y.\text{phase} = \text{vote} \\
& \quad \quad \wedge y.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \\
& \quad \quad \wedge s.\text{proc}(i).\text{healthy} > 0 \\
& \quad \quad \quad \supset y.\text{proc}(i).\text{proc_state} \\
& \quad \quad \quad \quad = f_c(u, s.\text{proc}(i).\text{proc_state}) \\
& \quad \quad \quad \wedge (\forall j : \\
& \quad \quad \quad \quad x.\text{proc}(j).\text{healthy} > 0 \\
& \quad \quad \quad \quad \supset y.\text{proc}(i).\text{mailbox}(j) \\
& \quad \quad \quad \quad \quad = f_s(x.\text{proc}(j).\text{proc_state})))))) \\
& \quad \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))
\end{aligned}$$

fc.A.1e: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$

$$\begin{aligned}
& \supset (\exists x, y, z : \\
& \quad \text{maj_working}(x) \\
& \quad \wedge \text{maj_working}(y) \\
& \quad \wedge (\forall i : \\
& \quad \quad y.\text{phase} = \text{vote} \\
& \quad \quad \wedge y.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \\
& \quad \quad \wedge s.\text{proc}(i).\text{healthy} > 0 \\
& \quad \quad \quad \supset y.\text{proc}(i).\text{proc_state} \\
& \quad \quad \quad \quad = f_c(u, s.\text{proc}(i).\text{proc_state}) \\
& \quad \quad \quad \wedge (\forall j : \\
& \quad \quad \quad \quad s.\text{proc}(j).\text{healthy} > 0 \\
& \quad \quad \quad \quad \supset y.\text{proc}(i).\text{mailbox}(j) \\
& \quad \quad \quad \quad \quad = f_s(y.\text{proc}(j).\text{proc_state})))))) \\
& \quad \wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u))
\end{aligned}$$

fc_A.1f: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists y, z :$
 maj_working(y)
 $\wedge (\forall i :$
 y.phase = vote
 $\wedge y.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset y.\text{proc}(i).\text{proc_state}$
 $= f_c(u, s.\text{proc}(i).\text{proc_state})$
 $\wedge (\forall j :$
 s.proc(j).healthy > 0
 $\supset y.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(y.\text{proc}(j).\text{proc_state}))))$
 $\wedge \mathcal{N}_{ds}(y, z, u) \wedge \mathcal{N}_{ds}(z, t, u)$

fc_A.2a: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists y, z :$
 maj_working(y)
 $\wedge \text{maj_working}(z)$
 $\wedge (\forall i :$
 y.phase = vote
 $\wedge y.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset y.\text{proc}(i).\text{proc_state}$
 $= f_c(u, s.\text{proc}(i).\text{proc_state})$
 $\wedge (\forall j :$
 s.proc(j).healthy > 0
 $\supset y.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(y.\text{proc}(j).\text{proc_state}))))$
 $\wedge z.\text{phase} = \text{next_phase}(y.\text{phase})$
 $\wedge z.\text{proc}(i).\text{healthy} = y.\text{proc}(i).\text{healthy}$
 $\wedge \mathcal{N}_{ds}^u(y, z, i)$
 $\wedge \mathcal{N}_{ds}(z, t, u)$

fc_A.2b: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists y, z :$
 maj_working(y)
 $\wedge \text{maj_working}(z)$
 $\wedge (\forall i :$
 z.phase = next_phase(vote)
 $\wedge z.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset y.\text{proc}(i).\text{proc_state}$
 $= f_c(u, s.\text{proc}(i).\text{proc_state})$
 $\wedge (\forall j :$
 s.proc(j).healthy > 0
 $\supset y.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(y.\text{proc}(j).\text{proc_state}))))$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset (z.\text{proc}(i).\text{mailbox} = y.\text{proc}(i).\text{mailbox}$
 $\wedge z.\text{proc}(i).\text{proc_state}$
 $= f_v(y.\text{proc}(i).\text{proc_state},$
 $y.\text{proc}(i).\text{mailbox}))))$
 $\wedge \mathcal{N}_{ds}(z, t, u)$

fc_A_2c: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists y, z :$
 $\text{maj_working}(y)$
 $\wedge \text{maj_working}(z)$
 $\wedge (\forall i :$
 $z.\text{phase} = \text{next_phase}(\text{vote})$
 $\wedge z.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset y.\text{proc}(i).\text{proc_state}$
 $= f_c(u, s.\text{proc}(i).\text{proc_state})$
 $\wedge z.\text{proc}(i).\text{proc_state}$
 $= f_v(y.\text{proc}(i).\text{proc_state},$
 $z.\text{proc}(i).\text{mailbox})$
 $\wedge (\forall j :$
 $s.\text{proc}(j).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(y.\text{proc}(j).\text{proc_state}))$
 $\wedge \mathcal{N}_{ds}(z, t, u))$

fc_A_2d: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists z : \text{maj_working}(z)$
 $\wedge (\forall i :$
 $z.\text{phase} = \text{sync}$
 $\wedge z.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{proc_state}$
 $= f_v(f_c(u, s.\text{proc}(i).\text{proc_state}),$
 $z.\text{proc}(i).\text{mailbox})$
 $\wedge (\forall j :$
 $s.\text{proc}(j).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(f_c(u, s.\text{proc}(j).\text{proc_state}))$
 $\wedge \mathcal{N}_{ds}(z, t, u))$

fc_A_3a: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists z : \text{maj_working}(t)$
 $\wedge \text{maj_working}(z)$
 $\wedge (\forall i :$
 $z.\text{phase} = \text{sync}$
 $\wedge z.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{proc_state}$
 $= f_v(f_c(u, s.\text{proc}(i).\text{proc_state}),$
 $z.\text{proc}(i).\text{mailbox})$
 $\wedge (\forall j :$
 $s.\text{proc}(j).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(f_c(u, s.\text{proc}(j).\text{proc_state}))$
 $\wedge t.\text{phase} = \text{next_phase}(z.\text{phase}) \wedge \mathcal{N}_{ds}^*(z, t, i))$

fc_A_3b: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists z : \text{maj_working}(t)$
 $\wedge (\forall i :$
 $t.\text{phase} = \text{next_phase}(\text{sync})$
 $\wedge z.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{proc_state}$
 $= f_v(f_c(u, s.\text{proc}(i).\text{proc_state}),$
 $z.\text{proc}(i).\text{mailbox})$
 $\wedge (\forall j :$
 $s.\text{proc}(j).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(f_c(u, s.\text{proc}(j).\text{proc_state}))))$
 $\wedge (z.\text{proc}(i).\text{healthy} > 0$
 $\supset t.\text{proc}(i).\text{proc_state} = z.\text{proc}(i).\text{proc_state})$
 $\wedge (t.\text{proc}(i).\text{healthy} > 0$
 $\supset t.\text{proc}(i).\text{healthy} = 1 + z.\text{proc}(i).\text{healthy}))$

fc_A_3c: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists z : \text{maj_working}(t)$
 $\wedge (\forall i :$
 $t.\text{phase} = \text{compute}$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset t.\text{proc}(i).\text{proc_state}$
 $= f_v(f_c(u, s.\text{proc}(i).\text{proc_state}),$
 $z.\text{proc}(i).\text{mailbox})$
 $\wedge (\forall j :$
 $s.\text{proc}(j).\text{healthy} > 0$
 $\supset z.\text{proc}(i).\text{mailbox}(j)$
 $= f_s(f_c(u, s.\text{proc}(j).\text{proc_state}))))$
 $\wedge (t.\text{proc}(i).\text{healthy} > 0$
 $\supset t.\text{proc}(i).\text{healthy} = 1 + s.\text{proc}(i).\text{healthy}))$

fc_A_3d: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset \text{maj_working}(t)$
 $\wedge (\exists h : (\forall i :$
 $t.\text{phase} = \text{compute}$
 $\wedge (t.\text{proc}(i).\text{healthy} > 0$
 $\supset t.\text{proc}(i).\text{healthy} = 1 + s.\text{proc}(i).\text{healthy})$
 $\wedge s.\text{proc}(i).\text{healthy} > 0$
 $\supset t.\text{proc}(i).\text{proc_state}$
 $= f_v(f_c(u, s.\text{proc}(i).\text{proc_state}), h(i))$
 $\wedge (\forall j :$
 $s.\text{proc}(j).\text{healthy} > 0$
 $\supset h(i)(j) = f_s(f_c(u, s.\text{proc}(j).\text{proc_state}))))))$

map_1: Lemma $(\text{DSmap}(s)(i)).\text{healthy} = s.\text{proc}(i).\text{healthy}$

map_2: Lemma $(\text{DSmap}(s)(i)).\text{proc_state} = s.\text{proc}(i).\text{proc_state}$

map_3: Lemma $\text{allowable_faults}(s, t) \supset \text{RS}(\text{allowable_faultsDSmap}(s), \text{DSmap}(t))$

map_4: Lemma $\text{RS}(\text{good_values_sentDSmap}(s), u, w) = \text{good_values_sent}(s, u, w)$

map_5: Lemma $\text{RS}(\text{voted_final_stateDSmap}(s), \text{DSmap}(t), u, h, i)$
 $= \text{voted_final_state}(s, t, u, h, i)$

map_7: Lemma $\text{RS}(\text{maj_workingDSmap}(s)) = \text{DS}(\text{maj_working}s)$

support_1: Lemma $(\forall i : s.\text{proc}(i).\text{healthy} = x.\text{proc}(i).\text{healthy})$
 $\wedge \text{allowable_faults}(x, y)$
 $\supset \text{allowable_faults}(s, y)$

support_4: Lemma $\mathcal{N}_{ds}(s, t, u) \supset t.\text{phase} = \text{next_phase}(s.\text{phase})$

support_5: Lemma $s.\text{phase} = ph \wedge ph \neq \text{sync} \wedge \mathcal{N}_{ds}(s, x, u)$
 $\supset (\forall i : s.\text{proc}(i).\text{healthy} = x.\text{proc}(i).\text{healthy})$

support_6: Lemma $s.\text{phase} = ph$
 $\wedge ph \neq \text{sync} \wedge \mathcal{N}_{ds}(s, x, u) \wedge \text{allowable_faults}(x, y)$
 $\supset \text{allowable_faults}(s, y)$

support_7: Lemma $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists x : \mathcal{N}_{ds}(s, x, u) \wedge x.\text{phase} = \text{broadcast} \wedge \text{half_frame_N_ds}(x, t, u))$

support_8: Lemma $x.\text{phase} = \text{broadcast} \wedge \text{half_frame_N_ds}(x, t, u)$
 $\supset (\exists y : \mathcal{N}_{ds}(x, y, u) \wedge y.\text{phase} = \text{vote} \wedge \text{quarter_frame_N_ds}(y, t, u))$

support_9: Lemma $y.\text{phase} = \text{vote} \wedge \text{quarter_frame_N_ds}(y, t, u)$
 $\supset (\exists z : \mathcal{N}_{ds}(y, z, u) \wedge z.\text{phase} = \text{sync} \wedge \mathcal{N}_{ds}(z, t, u))$

support_10: Lemma $s.\text{phase} = \text{sync} \wedge \mathcal{N}_{ds}(s, t, u) \supset \text{allowable_faults}(s, t)$

support_11: Lemma
 $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u) \supset \text{allowable_faults}(s, t)$

support_12: Lemma
 $s.\text{phase} = \text{compute} \wedge \text{frame_N_ds}(s, t, u)$
 $\supset (\exists z : z.\text{phase} = \text{sync} \wedge \mathcal{N}_{ds}(z, t, u))$

support_13: Lemma $\text{MBmatrix_cons}(\text{MBfn}, \text{nrep})(i) = \text{MBfn}(i)$

support_14: Lemma $\text{initial_ds}(s) \supset \text{working_set}(s) = \text{fullset}[\text{processors}]$

support_15: Lemma $\text{initial_ds}(s) \supset \text{maj_condition}(\text{working_set}(s))$

End

DS_top_proof: Module

Using DS_lemmas

Exporting all with DS_lemmas

Theory

ds : Var DSstate
 rs : Var RSstate
 p, q : Var nat
 ph : Var phases
 s, t, x, y, z : Var DSstate
 i, j, ii, jj : Var processors
 u : Var inputs
 w : Var MBvec
 h : Var MBmatrix
 k, m, n, a, b : Var proc_plus
 $prop$: Var function[proc_plus \rightarrow bool]

Proof

p_frame_commutates: Prove frame_commutates from fr_com_1, fr_com_2

p.initial_maps: Prove initial_maps from
 initial_ds {i ← p@p2},
 initial_rs {s ← DSmap(s)},
 map_1 {i ← p@p2},
 map_2 {i ← p@p2}

p.fr.com.1: Prove fr.com.1 from fc.A, fc.B

p.fr.com.2: Prove fr.com.2 from
 \mathcal{N}_{rs} {s ← DSmap(s), t ← DSmap(t), h ← h@p2},
 is_new_proc_state {s ← s, t ← t, i ← p@p1},
 map_3 {s ← s, t ← t},
 map_4 {s ← s, w ← h@p2(p@p1)},
 map_5 {s ← s, t ← t, h ← h@p2, i ← p@p1},
 map_1 {s ← s, i ← p@p1}

p.fc.A: Prove fc.A from
 fc.A_3d {i ← i@p2, j ← j@p3},
 is_new_proc_state {h ← h@p1},
 DS.to.RS.good_values_sent {w ← h@p1(i@p2)},
 DS.to.RS.voted_final_state {i ← i@p2, h ← h@p1}

p.fc.B: Prove fc.B from support_11

p.fc.A.1a: Prove fc.A.1a {x ← x@p1, y ← y@p1, z ← z@p1} from
 frame.N_ds,
 \mathcal{N}_{ds} {s ← s@p1, t ← x@p1},
 next_phase {ph ← s.phase},
 distinct_phases

p.fc.A.1b: Prove fc.A.1b {x ← x@p1, y ← y@p1, z ← z@p1} from
 fc.A.1a, \mathcal{N}_{ds} {s ← x@p1, t ← y@p1}, distinct_phases

p.fc.A.1c: Prove fc.A.1c {x ← x@p1, y ← y@p1, z ← z@p1} from
 fc.A.1b,
 fc.A.1b {i ← j},
 distinct_phases,
 next_phase {ph ← x.phase},
 \mathcal{N}_{ds}^c {t ← x},
 \mathcal{N}_{ds}^c {t ← x, i ← j},
 \mathcal{N}_{ds}^b {s ← x, t ← y},
 broadcast_received {s ← x, t ← y, p ← i, qq ← j}

p.fc.A.1d: Prove fc.A.1d {x ← x@p1, y ← y@p1, z ← z@p1} from
 fc.A.1c, fc.A.1c {i ← jj}

p.fc.A.1e: Prove fc.A.1e {x ← x@p1, y ← y@p1, z ← z@p1} from
 fc.A.1d {jj ← j}, fc.A.1d

p.fc.A.1f: Prove fc.A.1f {y ← y@p1, z ← z@p1} from fc.A.1e

p.fc.A.2a: Prove fc.A.2a {y ← y@p1, z ← z@p1} from
 fc.A.1f, \mathcal{N}_{ds} {s ← y@p1, t ← z@p1}, distinct_phases

p.fc.A.2b: Prove fc.A.2b {y ← y@p1, z ← z@p1} from
 fc.A.2a, \mathcal{N}_{ds}^v {s ← y, t ← z, i ← i@C}

p.fc.A.2c: Prove fc.A.2c {y ← y@p1, z ← z@p1} from fc.A.2b

p.fc.A.2d: Prove fc.A.2d {z ← z@p1} from
 fc.A.2c, next_phase {ph ← vote}, distinct_phases, fc.A.2c {i ← j}

p.fc.A.3a: Prove fc.A.3a $\{z \leftarrow z@p1\}$ from
 fc.A.2d, $\mathcal{N}_{ds} \{s \leftarrow z@p1, t \leftarrow t@p1\}$, distinct_phases

p.fc.A.3b: Prove fc.A.3b $\{z \leftarrow z@p1\}$ from
 fc.A.3a, $\mathcal{N}'_{ds} \{s \leftarrow z, i \leftarrow i@c\}$

p.fc.A.3c: Prove fc.A.3c $\{z \leftarrow z@p1\}$ from
 fc.A.3b, next_phase $\{ph \leftarrow \text{sync}\}$, distinct_phases

p.fc.A.3d: Prove fc.A.3d
 $\{h \leftarrow \text{MBmatrix_cons}((\lambda i : z@p1.\text{proc}(i).\text{mailbox}), \text{nrep})\}$ from
 fc.A.3c
 $\{j \leftarrow j@c,$
 $i \leftarrow i@c,$
 $u \leftarrow u@c,$
 $t \leftarrow t@c,$
 $s \leftarrow s@c\},$
 support_13 $\{MBfn \leftarrow (\lambda i : z@p1.\text{proc}(i).\text{mailbox}), i \leftarrow i\}$

End

DS_map_proof: **Module**

Using DS_lemmas, nat_inductions

Exporting all with DS_lemmas

Theory

ds: Var DSstate
rs: Var RSstate
p, qq: Var nat
ph: Var phases
s, t, x, y, z: Var DSstate
i, j: Var processors
u: Var inputs
w: Var MBvec
h: Var MBmatrix
k, m, n, a, b: Var proc_plus
prop: Var function[proc_plus \rightarrow bool]

Proof

ml1_prop: function[DSstate, processors \rightarrow function[proc_plus \rightarrow bool]] =
 $(\lambda ds, i : (\lambda k :$
 $\text{ss_update}(ds, k)(i).\text{healthy}$
 $= \text{if } i \leq k \text{ then } ds.\text{proc}(i).\text{healthy} \text{ else } rs_0(i).\text{healthy} \text{ end if}))$

ml1_base: Lemma ml1_prop(*s, i*)(0)

ml1_ind: Lemma $k < \text{nrep} \wedge \text{ml1_prop}(s, i)(k) \supset \text{ml1_prop}(s, i)(k + 1)$

p_ml1_base: Prove ml1_base from
 ml1_prop $\{ds \leftarrow s, i \leftarrow i, k \leftarrow 0\},$
 ss_update $\{ds \leftarrow s, p \leftarrow 0\}$

p_ml1_ind: Prove ml1_ind from
 ml1_prop $\{ds \leftarrow s, i \leftarrow i, k \leftarrow k\},$
 ml1_prop
 $\{ds \leftarrow s,$
 $i \leftarrow i,$
 $k \leftarrow \text{if } k = \text{nrep} \text{ then } \text{nrep} \text{ else } k + 1 \text{ end if}\},$
 ss_update $\{ds \leftarrow s, p \leftarrow k + 1\}$

p_map_1: Prove map.1 from

```
DSmap {ds ← s},
processors_induction {prop ← ml1_prop(s, i), n ← nrep},
ml1_prop {ds ← s, i ← i, k ← nrep},
ml1_base {s ← s, i ← i},
ml1_ind {s ← s, i ← i, k ← m@P2}
```

ml2_prop: function[DSstate, processors → function[proc_plus → bool]] =

```
(λ ds, i : (λ k :
  ss_update(ds, k)(i).proc_state
  = if i ≤ k
    then ds.proc(i).proc_state
    else rs0(i).proc_state
  end if))
```

ml2_base: Lemma ml2_prop(s, i)(0)

ml2_ind: Lemma $k < nrep \wedge ml2_prop(s, i)(k) \supset ml2_prop(s, i)(k + 1)$

p_ml2_base: Prove ml2_base from

```
ml2_prop {ds ← s, i ← i, k ← 0},
ss_update {ds ← s, p ← 0}
```

p_ml2_ind: Prove ml2_ind from

```
ml2_prop {ds ← s, i ← i, k ← k},
ml2_prop
  {ds ← s,
   i ← i,
   k ← if k = nrep then nrep else k + 1 end if},
ss_update {ds ← s, p ← k + 1}
```

p_map_2: Prove map.2 from

```
DSmap {ds ← s},
processors_induction {prop ← ml2_prop(s, i), n ← nrep},
ml2_prop {ds ← s, i ← i, k ← nrep},
ml2_base {s ← s, i ← i},
ml2_ind {s ← s, i ← i, k ← m@P2}
```

p_map_3: Prove map.3 from

```
RS.allowable_faults {s ← DSmap(s), t ← DSmap(t)},
DS.allowable_faults {s ← s, t ← t, i ← p@p1},
map.7 {s ← t},
map.1 {s ← s, i ← p@p1},
map.1 {s ← t, i ← p@p1}
```

p_map_4: Prove map.4 from

```
RS.good_values_sent {s ← DSmap(s), q ← j@P2},
DS.to_RS.good_values_sent {j ← q@P1S},
map.1 {i ← j@p2},
map.2 {i ← j@p2},
map.1 {i ← q@P1},
map.2 {i ← q@P1}
```

p_map_5: Prove map.5 from

```
RS.voted_final_state {s ← DSmap(s), t ← DSmap(t), p ← i},
DS.to_RS.voted_final_state,
map.1 {i ← i},
map.1 {s ← t, i ← i},
map.2 {i ← i},
map.2 {s ← t, i ← i}
```


p_map_7: Prove map_7 from
proc_extensionality
{A ← RS(.working_setDSmap(s)),
B ← DS(.working_sets)},
RS.maj_working {t ← DSmap(s)},
RS.working_set {s ← DSmap(s), p ← p@p1},
RS.working_proc {s ← DSmap(s), p ← p@p1},
DS.maj_working {t ← s},
DS.working_set {s ← s, p ← p@p1},
DS.working_proc {s ← s, p ← p@p1},
map_1 {i ← p@p1}

End

DS_support_proof: Module

Using DS_lemmas, nat_inductions

Exporting all with DS_lemmas

Theory

ds: Var DSstate
rs: Var RSstate
p, q: Var nat
ph: Var phases
s, t, x, y, z: Var DSstate
i, j: Var processors
u: Var inputs
w: Var MBvec
h: Var MBmatrix
MBfn: Var MBcons_fn
k, m, n, a, b: Var proc.plus
prop: Var function[proc.plus → bool]

Proof

p_support_1: Prove support_1 {i ← i@p2} from
DS.allowable_faults {s ← x, t ← y, i ← i@p2},
DS.allowable_faults {s ← s, t ← y}

p_support_4: Prove support_4 from \mathcal{N}_{ds}

p_support_5: Prove support_5 from
member_phases {phases_var ← ph},
 \mathcal{N}_{ds} {s ← s, t ← x, u ← u, i ← i}

p_support_6: Prove support_6 from support_1, support_5 {i ← i@p1}

p_support_7: Prove support_7 {x ← x@p1} from
frame_N_ds,
half_frame_N_ds {x ← x@p1, y ← y@p1, z ← z@p1},
support_4 {s ← s, t ← x@p1, u ← u},
next_phase {ph ← compute}

p_support_8: Prove support_8 {y ← y@p1} from
half_frame_N_ds,
quarter_frame_N_ds {y ← y@p1, z ← z@p1},
support_4 {s ← x, t ← y@p1, u ← u},
next_phase {ph ← broadcast},
distinct_phases

p_support_9: Prove support_9 $\{z \leftarrow z@p1\}$ from
quarter_frame_N_ds,
support_4 $\{s \leftarrow y, t \leftarrow z@p1, u \leftarrow u\}$,
next_phase $\{ph \leftarrow vote\}$,
distinct_phases

p_support_10: Prove support_10 from
DS.allowable_faults, \mathcal{N}_{ds} $\{i \leftarrow i@p1\}$, \mathcal{N}'_{ds} $\{i \leftarrow i@p1\}$

p_support_11: Prove support_11 from
support_6 $\{s \leftarrow s, x \leftarrow x@p4, y \leftarrow t, ph \leftarrow compute\}$,
support_6 $\{s \leftarrow x@p4, x \leftarrow y@p5, y \leftarrow t, ph \leftarrow broadcast\}$,
support_6 $\{s \leftarrow y@p5, x \leftarrow z@p6, y \leftarrow t, ph \leftarrow vote\}$,
support_7,
support_8 $\{x \leftarrow x@p4\}$,
support_9 $\{y \leftarrow y@p5\}$,
support_10 $\{s \leftarrow z@p6\}$,
distinct_phases

p_support_12: Prove support_12 $\{z \leftarrow z@p3\}$ from
support_7, support_8 $\{x \leftarrow x@p1\}$, support_9 $\{y \leftarrow y@p2\}$

sl13_prop: function[MBcons_fn, processors \rightarrow function[proc_plus \rightarrow bool]] =
 $(\lambda MBfn, i : (\lambda k :$
MBmatrix_cons(MBfn, k)(i)
= if $i \leq k$ then MBfn(i) else MBmatrix0(i) end if)

sl13_base: Lemma sl13_prop(MBfn, i)(0)

sl13_ind: Lemma $m < nrep \wedge sl13_prop(MBfn, i)(m) \supset$
sl13_prop(MBfn, i)(m + 1)

p_sl13_base: Prove sl13_base from
sl13_prop $\{k \leftarrow 0, i \leftarrow i\}$, MBmatrix_cons $\{p \leftarrow 0\}$

p_sl13_ind: Prove sl13_ind from
sl13_prop $\{k \leftarrow m, i \leftarrow i\}$,
sl13_prop $\{i \leftarrow i, k \leftarrow \text{if } m = nrep \text{ then } nrep \text{ else } m + 1 \text{ end if}\}$,
MBmatrix_cons $\{p \leftarrow m + 1\}$

p_support_13: Prove support_13 from
processors_induction $\{prop \leftarrow sl13_prop(MBfn, i), n \leftarrow nrep\}$,
sl13_prop $\{k \leftarrow nrep, i \leftarrow i\}$,
sl13_base $\{i \leftarrow i\}$,
sl13_ind $\{i \leftarrow i, m \leftarrow m@p1\}$

p_support_14: Prove support_14 from
proc_extensionality $\{A \leftarrow working_set(s), B \leftarrow fullset[processors]\}$,
initial_ds $\{i \leftarrow p@p1\}$,
DS.working_set $\{p \leftarrow p@p1\}$,
DS.working_proc $\{p \leftarrow p@p1\}$

p_support_15: Prove support_15 from
maj_condition $\{A \leftarrow working_set(s)\}$, support_14, card_fullset

End

DS_to_RS.tcc: **Module**

Using DS_to_RS

Exporting all with DS_to_RS

Theory

ds: **Var** DS.DSstate
p: **Var** naturalnumber
MBfn: **Var** function[rcp_defs.processors \rightarrow rcp_defs.MBvec]
ss_update_TCC1: **Formula** ($\neg((p = 0) \vee (p > nrep))$) $\supset (p - 1 \geq 0)$
ss_update_TCC2: **Formula** ($\neg((p = 0) \vee (p > nrep))$) $\supset ((p > 0) \wedge (p \leq nrep))$
ss_update_TCC3: **Formula**
($\neg((p = 0) \vee (p > nrep))$) \supset **ssu.measure**(*ds*, *p*) > **ssu.measure**(*ds*, *p* - 1)
MBmatrix_cons_TCC1: **Formula**
($\neg((p = 0) \vee (p > nrep))$)
 \supset **MBmc.measure**(*MBfn*, *p*) > **MBmc.measure**(*MBfn*, *p* - 1)

Proof

ss_update_TCC1_PROOF: **Prove** **ss_update_TCC1**
ss_update_TCC2_PROOF: **Prove** **ss_update_TCC2**
ss_update_TCC3_PROOF: **Prove** **ss_update_TCC3**
MBmatrix_cons_TCC1_PROOF: **Prove** **MBmatrix_cons_TCC1**

End DS.to_RS.tcc

DS.support_proof.tcc: **Module**

Using DS.support_proof

Exporting all with DS.support_proof

Theory

p: **Var** rcp_defs.processors
m: **Var** rcp_defs.proc_plus
z: **Var** DS.DSstate
y: **Var** DS.DSstate
x: **Var** DS.DSstate
i: **Var** rcp_defs.processors
p_sl13_base_TCC1: **Formula** ((0 \geq 0) \wedge (0 \leq nrep))
p_sl13_ind_TCC1: **Formula**
((**if** *m* = nrep **then** nrep **else** *m* + 1 **end if** \geq 0)
 \wedge (**if** *m* = nrep **then** nrep **else** *m* + 1 **end if** \leq nrep))
p_support_13_TCC1: **Formula** ((nrep \geq 0) \wedge (nrep \leq nrep))

Proof

p_sl13_base_TCC1_PROOF: **Prove** **p_sl13_base_TCC1**
p_sl13_ind_TCC1_PROOF: **Prove** **p_sl13_ind_TCC1**
p_support_13_TCC1_PROOF: **Prove** **p_support_13_TCC1**

End DS.support_proof.tcc

DS.map_proof.tcc: **Module**

Using DS.map_proof

Exporting all with DS.map_proof

Theory

```
k: Var rcp_defs.proc_plus
q: Var rcp_defs.processors
j: Var rcp_defs.processors
p: Var rcp_defs.processors
m: Var rcp_defs.proc_plus
```

```
p.ml1_base.TCC1: Formula ((0 ≥ 0) ∧ (0 ≤ nrep))
```

```
p.ml1_ind.TCC1: Formula
  (( if k = nrep then nrep else k + 1 end if ≥ 0)
    ∧ ( if k = nrep then nrep else k + 1 end if ≤ nrep))
```

```
p.map.1.TCC1: Formula ((nrep ≥ 0) ∧ (nrep ≤ nrep))
```

Proof

```
p.ml1_base.TCC1.PROOF: Prove p.ml1_base.TCC1
```

```
p.ml1_ind.TCC1.PROOF: Prove p.ml1_ind.TCC1
```

```
p.map.1.TCC1.PROOF: Prove p.map.1.TCC1
```

End DS_map_proof.tcc

DA: Module

Using clkmod, generic_FT

Exporting all with clkmod, generic_FT

Theory

```
max_comm_delay: realtime (* max broadcast delivery time *)
da_proc_state: Type = Record healthy : nat,
  proc_state : Pstate,
  mailbox : MBvec,
  lclock : logical_clocktime,
  cum_delta : number (* = Corr; added to logical
    end record to obtain physical *)
da_proc_array: Type = array [processors] of da_proc_state
```

```
DAstate: Type = Record phase : phases,
  sync_period : nat, (* = idealized frame count *)
  proc : da_proc_array
  end record
```

```
s, t, x, y, z, da: Var DAstate
u: Var inputs
w: Var MBvec
i, j, p, q, qq: Var processors
k: Var nat
ph: Var phases
ps: Var da_proc_state
T: Var logical_clocktime
A: Var set[processors]
```

```
Corr_implementation: Lemma s.proc(p).cum_delta = Corrp(s.sync_period)
```

```

working_proc: function[DAstate, processors → bool] =
  (λ s, p : s.proc(p).healthy ≥ recovery_period)
working_set: function[DAstate → set[processors]] =
  (λ s : (λ p : working_proc(s, p)))
maj_working: function[DAstate → bool] =
  (λ t : maj_condition(working_set(t)))

enough_hardware: function[DAstate → bool] =
  (λ t : maj_working(t) ∧ enough_clocks(t.sync_period))

da_rt: function[DAstate, processors, logical_clocktime → realtime] =
  (λ da, p, T : c_p(T + da.proc(p).cum_delta))
unknown: fraction
ν: fraction = unknown (* variability of processor run rates *)
X, Y: Var logical_clocktime
D: Var number
clock_advanced: function[logical_clocktime, logical_clocktime, number
  → bool] =
  (λ X, Y, D : X + D * (1 - ν) ≤ Y ∧ Y ≤ X + D * (1 + ν))
duration: function[phases → logical_clocktime]

broadcast_duration: Axiom
  (1 - Rho) * |duration(broadcast) - 2 * ν * duration(compute) - ν * duration(broadcast)| - δ
  ≥ max_comm_delay

broadcast_duration2: Axiom
  duration(broadcast) - 2 * ν * duration(compute) - ν * duration(broadcast) ≥ 0
all_durations: Axiom
  (1 + ν) * duration(compute) + (1 + ν) * duration(broadcast) ≤ frame_time

pos_durations: Axiom
  0 ≤ (1 - ν) * duration(compute)
  ∧ 0 ≤ (1 - ν) * duration(broadcast)
  ∧ 0 ≤ (1 - ν) * duration(vote) ∧ 0 ≤ (1 - ν) * duration(sync)

broadcast_received: function[DAstate, DAstate, processors → bool] =
  (λ s, t, p : (∀ qq :
    s.proc(qq).healthy > 0
    ∧ da_rt(s, qq, s.proc(qq).lclock) + max_comm_delay
    ≤ da_rt(t, p, t.proc(p).lclock)
    ⊃ t.proc(p).mailbox(qq) = s.proc(qq).mailbox(qq)))

N_da^c: function[DAstate, DAstate, inputs, processors → bool] =
  (λ s, t, u, i :
    s.proc(i).healthy > 0
    ⊃ t.proc(i).proc_state = f_c(u, s.proc(i).proc_state)
    ∧ t.proc(i).mailbox(i) = f_o(f_c(u, s.proc(i).proc_state)))

N_da^b: function[DAstate, DAstate, processors → bool] =
  (λ s, t, i : s.proc(i).healthy > 0
    ⊃ t.proc(i).proc_state = s.proc(i).proc_state
    ∧ broadcast_received(s, t, i))

```

```

 $\mathcal{N}_{da}^v$ : function[Dstate, Dstate, processors  $\rightarrow$  bool] =
  ( $\lambda$  s, t, i : s.proc(i).healthy > 0
     $\supset$  t.proc(i).mailbox = s.proc(i).mailbox
       $\wedge$  t.proc(i).proc_state
        =  $f_v$ (s.proc(i).proc_state, s.proc(i).mailbox))
  )

```

```

 $\mathcal{N}_{da}^s$ : function[Dstate, Dstate, processors  $\rightarrow$  bool] =
  ( $\lambda$  s, t, i : s.proc(i).healthy > 0
     $\supset$  t.proc(i).proc_state = s.proc(i).proc_state)
   $\wedge$  (t.proc(i).healthy > 0
     $\supset$  t.proc(i).healthy = 1 + s.proc(i).healthy
       $\wedge$  nonfaulty_clock(i, t.sync_period))
   $\wedge$  t.sync_period = 1 + s.sync_period
   $\wedge$  (nonfaulty_clock(i, s.sync_period)
     $\supset$  t.proc(i).lclock = (1 + s.sync_period) * frame_time
       $\wedge$  t.proc(i).cum_delta
        = s.proc(i).cum_delta +  $\Delta_i^{(s.sync\_period)}$ )
  )

```

```

 $\mathcal{N}_{da}$ : function[Dstate, Dstate, inputs  $\rightarrow$  bool] =
  ( $\lambda$  s, t, u : enough_hardware(t)
     $\wedge$  t.phase = next_phase(s.phase)
     $\wedge$  ( $\forall$  i :
      if s.phase = sync
        then  $\mathcal{N}_{da}^s$ (s, t, i)
      else t.proc(i).healthy = s.proc(i).healthy
         $\wedge$  t.proc(i).cum_delta = s.proc(i).cum_delta
           $\wedge$  t.sync_period = s.sync_period
             $\wedge$  (nonfaulty_clock(i, s.sync_period)
               $\supset$  clock_advanced(s.proc(i).lclock,
                t.proc(i).lclock,
                duration(s.phase)))
           $\wedge$  (s.phase = compute  $\supset$   $\mathcal{N}_{da}^c$ (s, t, u, i))
             $\wedge$  (s.phase = broadcast  $\supset$   $\mathcal{N}_{da}^b$ (s, t, i))
               $\wedge$  (s.phase = vote  $\supset$   $\mathcal{N}_{da}^v$ (s, t, i))
        end if)
  )

```

```

initial_da: function[Dstate  $\rightarrow$  bool] =
  ( $\lambda$  s : s.phase = compute
     $\wedge$  s.sync_period = 0
     $\wedge$  ( $\forall$  i :
      s.proc(i).healthy = recovery_period
         $\wedge$  s.proc(i).proc_state = initial_proc_state
           $\wedge$  s.proc(i).cum_delta = 0
             $\wedge$  s.proc(i).lclock = 0  $\wedge$  nonfaulty_clock(i, 0))
  )

```

End

DA.to_DS: Module

Using DA, DS

Exporting all with DA, DS

Theory

```

da, s, t, x, y, z: Var DAstate
ds: Var DSstate
p, i, j: Var processors
k, l: Var nat
u: Var inputs
w: Var MBvec
h: Var MBmatrix
ph: Var phases
MBmatrix0: MBmatrix
MBcons_fn: Type is function[processors → MBvec]
MBfn: Var MBcons_fn
T, T1, T2, BB: Var logical_clocktime
DAstate_prop: Type is function[DAstate → bool]
da_prop: Var DAstate_prop
da_measure: function[DAstate, nat → nat] == (λ da, k : k)
ss_update: Recursive function[DAstate, nat → DSstate] =
  (λ da, k : if (k = 0) ∨ (k > nrep)
    then ds0
    else ss_update(da, k - 1)
    with [(proc)(k) := dsproc0
          with [(healthy) := da.proc(k).healthy,
                (proc_state) := da.proc(k).proc_state,
                (mailbox) := da.proc(k).mailbox]]
    end if) by da_measure

DAmap: function[DAstate → DSstate] =
  (λ da : ss_update(da, nrep) with [(phase) := da.phase])
MBmc_measure: function[MBcons_fn, nat → nat] == (λ MBfn, k : k)
MBmatrix_cons: Recursive function[MBcons_fn, nat → MBmatrix] =
  (λ MBfn, k : if (k = 0) ∨ (k > nrep)
    then MBmatrix0
    else MBmatrix_cons(MBfn, k - 1)
    with [(k) := MBfn(k)]
    end if) by MBmc_measure

reachable_in_n: function[DAstate, nat → bool] =
  (λ t, k : if k = 0
    then initial_da(t)
    else (∃ s, u : reachable_in_n(s, k - 1) ∧  $\mathcal{N}_{da}(s, t, u)$ )
    end if) by da_measure
reachable: function[DAstate → bool] = (λ t : (∃ k : reachable_in_n(t, k)))

phase_commutates: Theorem reachable(s) ∧  $\mathcal{N}_{da}(s, t, u) \supset \mathcal{N}_{ds}(DAmap(s), DAmap(t), u)$ 
initial_maps: Theorem initial_da(s) ⊃ initial_ds(DAmap(s))

```

End

DA_invariants: Module

Using DA_to_DS, nat_inductions, DA_lemmas

Exporting all with DA_to_DS

Theory

da, s, t, x, y, z : Var DАstate
 ds : Var DSstate
 p, i, j : Var processors
 k, l : Var nat
 u : Var inputs
 w : Var MBvcc
 h : Var MBmatrix
 ph : Var phases
 cdv : Var number
 ii : Var period
 T, T_1, T_2, BB : Var logical_clocktime
DАstate_prop: Type is function[DАstate \rightarrow bool]
 da_prop : Var DАstate_prop
state_invariant: function[DАstate_prop \rightarrow bool] =
 $(\lambda da_prop : (\forall t : \text{reachable}(t) \supset da_prop(t)))$
state_induction: Lemma
 $(\forall x : \text{initial_da}(x) \supset da_prop(x))$
 $\wedge (\forall s, t, u : \text{reachable}(s) \wedge da_prop(s) \wedge \mathcal{N}_{da}(s, t, u) \supset da_prop(t))$
 $\supset \text{state_invariant}(da_prop)$
enough_inv: Lemma state_invariant($(\lambda s : \text{enough_hardware}(s))$)
nf_clks: function[DАstate \rightarrow bool] =
 $(\lambda s : (\forall i :$
 $s.\text{proc}(i).\text{healthy} > 0 \supset \text{nonfaulty_clock}(i, s.\text{sync_period}))$)
nfclk_inv: Lemma state_invariant($(\lambda s : \text{nf_clks}(s))$)
lclock_eq: function[DАstate \rightarrow bool] =
 $(\lambda s : (\forall i, j :$
 $\text{nonfaulty_clock}(i, s.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(j, s.\text{sync_period}) \wedge s.\text{phase} = \text{compute}$
 $\supset s.\text{proc}(i).\text{lclock} = s.\text{proc}(j).\text{lclock})$)
lclock_inv: Lemma state_invariant($(\lambda s : \text{lclock_eq}(s))$)
lclock_val: function[DАstate \rightarrow bool] =
 $(\lambda s : (\forall i :$
 $\text{nonfaulty_clock}(i, s.\text{sync_period}) \wedge s.\text{phase} = \text{compute}$
 $\supset s.\text{proc}(i).\text{lclock} = s.\text{sync_period} * \text{frame.time})$)
clkval_inv: Lemma state_invariant($(\lambda s : \text{lclock_val}(s))$)
rtl1: Lemma $\text{reachable}(da) \wedge \text{nonfaulty_clock}(p, da.\text{sync_period})$
 $\supset da.\text{proc}(p).\text{cum_delta} = \text{Corr}_p^{(da.\text{sync_period})}$
da_rt_lem: Lemma $\text{reachable}(da) \wedge \text{nonfaulty_clock}(p, da.\text{sync_period})$
 $\supset da.\text{rt}(da, p, T) = \text{rt}_p^{(da.\text{sync_period})}(T)$
cum_delta_val: function[DАstate \rightarrow bool] =
 $(\lambda s : (\forall p :$
 $\text{nonfaulty_clock}(p, s.\text{sync_period})$
 $\supset s.\text{proc}(p).\text{cum_delta} = \text{Corr}_p^{(s.\text{sync_period})})$)
Corr_lem: Lemma $ii > 0 \supset \text{Corr}_p^{(ii)} = \text{Corr}_p^{(ii-1)} + \Delta_p^{(\text{pred}(ii))}$
cdl1: Lemma $\mathcal{N}_{da}(s, t, u) \wedge s.\text{proc}(p).\text{cum_delta} = cdv$
 $\supset t.\text{proc}(p).\text{cum_delta} = cdv + \Delta_p^{(\text{pred}(ii))}$
cum_delta_inv: Lemma state_invariant($(\lambda s : \text{cum_delta_val}(s))$)

Proof

state_invariant_to_n: function[DAstate_prop, nat → bool] =
 (λ da_prop, k : (∀ t : reachable_in_n(t, k) ⊃ da_prop(t)))

base_state_ind: Lemma
 (initial_da(x) ⊃ da_prop(x)) ⊃ (reachable_in_n(x, 0) ⊃ da_prop(x))

ind_state_ind: Lemma
 (∀ s, t, u : reachable(s) ∧ da_prop(s) ∧ $\mathcal{N}_{da}(s, t, u)$ ⊃ da_prop(t))
 ⊃ (∀ k : state_invariant_to_n(da_prop, k)
 ⊃ state_invariant_to_n(da_prop, k + 1))

p_base_state_ind: Prove base_state_ind from
 reachable_in_n {t ← x, k ← 0}

p_ind_state_ind: Prove
 ind_state_ind {s ← s@p3, t ← t@p2, u ← u@p3} from
 state_invariant_to_n {k ← k, t ← s@p3},
 state_invariant_to_n {k ← k + 1, t ← t},
 reachable_in_n {t ← t, k ← k + 1},
 reachable {t ← s@p3, k ← k}

p_state_induction: Prove
 state_induction
 {x ← t@p3,
 s ← s@p4,
 t ← t@p4,
 u ← u@p4} from
 nat_induction
 {p ← (λ k : state_invariant_to_n(da_prop, k)),
 n2 ← k@p7},
 base_state_ind {x ← t@p3},
 state_invariant_to_n {t ← x, k ← 0},
 ind_state_ind {k ← n1@p1},
 state_invariant_to_n {t ← t@p6, k ← k@p7},
 state_invariant,
 reachable {t ← t@p6}

enough_inv_l1: Lemma initial_da(s) ⊃ enough_hardware(s)

enough_inv_l2: Lemma $\mathcal{N}_{da}(s, t, u) \wedge$ enough_hardware(s) ⊃ enough_hardware(t)

p_enough_inv_l1: Prove enough_inv_l1 from
 enough_hardware {t ← s},
 enough_clocks {i ← s.sync_period},
 DA.maj_working {t ← s},
 support_14,
 support_15,
 processors_exist_ax

p_enough_inv_l2: Prove enough_inv_l2 from \mathcal{N}_{da}

p_enough_inv: Prove enough_inv from
 state_induction {da_prop ← (λ s : enough_hardware(s))},
 enough_inv_l1 {s ← x@p1},
 enough_inv_l2 {s ← s@p1, t ← t@p1, u ← u@p1}

nfclk_inv_l1: Lemma initial_da(s) ⊃ nf.clks(s)

nfclk_inv_l2: Lemma $\mathcal{N}_{da}(s, t, u) \wedge$ nf.clks(s) ⊃ nf.clks(t)

p_nfclk_inv_l1: Prove nfclk_inv_l1 from nf_clks, initial_da {i ← i@p1}

p_nfclk_inv_l2: Prove nfclk_inv_l2 from

\mathcal{N}_{da} {i ← i@p3}, nf_clks {i ← i@p3}, nf_clks {s ← t}, \mathcal{N}_{da}^s {i ← i@p3}

p_nfclk_inv: Prove nfclk_inv from

state_induction {da_prop ← (λ s : nf_clks(s))},

nfclk_inv_l1 {s ← x@p1},

nfclk_inv_l2 {s ← s@p1, t ← t@p1, u ← u@p1}

lclock_inv_l1: Lemma initial_da(s) ⊃ lclock_eq(s)

lclock_inv_l2: Lemma $\mathcal{N}_{da}(s, t, u) \wedge s.\text{phase} = \text{sync} \supset \text{lclock_eq}(t)$

lclock_inv_l2b: Lemma

$\mathcal{N}_{da}(s, t, u) \supset (s.\text{phase} = \text{sync} \supset t.\text{phase} = \text{compute})$

$\wedge (s.\text{phase} = \text{compute} \supset t.\text{phase} = \text{broadcast})$

$\wedge (s.\text{phase} = \text{broadcast} \supset t.\text{phase} = \text{vote})$

$\wedge (s.\text{phase} = \text{vote} \supset t.\text{phase} = \text{sync})$

p_lclock_inv_l2b: Prove lclock_inv_l2b from

\mathcal{N}_{da} ,

distinct_phases,

next_phase {ph ← compute},

next_phase {ph ← vote},

next_phase {ph ← broadcast},

next_phase {ph ← sync}

lclock_inv_l2c: Lemma $\mathcal{N}_{da}(s, t, u) \wedge s.\text{phase} \neq \text{sync} \supset t.\text{phase} \neq \text{compute}$

p_lclock_inv_l2c: Prove lclock_inv_l2c from

lclock_inv_l2b,

distinct_phases,

member_phases {phases_var ← t.phase},

member_phases {phases_var ← s.phase}

lclock_inv_l3: Lemma

$\mathcal{N}_{da}(s, t, u) \wedge s.\text{phase} \neq \text{sync} \supset t.\text{sync_period} = s.\text{sync_period}$

lclock_inv_l4: Lemma

$\mathcal{N}_{da}(s, t, u) \wedge s.\text{phase} \neq \text{sync} \wedge \text{lclock_eq}(s) \supset \text{lclock_eq}(t)$

p_lclock_inv_l1: Prove lclock_inv_l1 from

lclock_eq, initial_da {i ← i@p1}, initial_da {i ← j@p1}

p_lclock_inv_l2: Prove lclock_inv_l2 from

lclock_eq {s ← t},

\mathcal{N}_{da} {i ← i@p1},

\mathcal{N}_{da} {i ← j@p1},

\mathcal{N}_{da}^s {i ← i@p1},

\mathcal{N}_{da}^s {i ← j@p1},

nf_lem {p ← i@p1, i ← s.sync_period},

nf_lem {p ← j@p1, i ← s.sync_period}

p_lclock_inv_l3: Prove lclock_inv_l3 from \mathcal{N}_{da} {i ← i}, lclock_inv_l2b

p_lclock_inv_l4: Prove lclock_inv_l4 from

lclock_eq {s ← t},

lclock_eq {i ← i@p1, j ← j@p1},

lclock_inv_l2c,

distinct_phases,

lclock_inv_l3

p.lclock_inv: Prove lclock_inv from

state_induction {da_prop ← (λ s : lclock_eq(s))},
lclock_inv_l1 {s ← x@p1},
lclock_inv_l2 {s ← s@p1, t ← t@p1, u ← u@p1},
lclock_inv_l4 {s ← s@p1, t ← t@p1, u ← u@p1}

clkval_inv_l1: Lemma initial_da(s) ⊃ lclock_val(s)

clkval_inv_l2: Lemma reachable(s) ∧ $\mathcal{N}_{da}(s, t, u)$ ⊃ lclock_val(t)

p.clkval_inv_l1: Prove clkval_inv_l1

from lclock_val, initial_da {i ← i@p1}

p.clkval_inv_l2: Prove clkval_inv_l2 from

lclock_val {s ← t},
 \mathcal{N}_{da} {i ← i@p1},
 \mathcal{N}_{da}^s {i ← i@p1},
support_16 {ph ← s.phase},
prev_phase {ph ← t.phase},
nfc_lem {p ← i@p1, i ← s.sync_period}

p.clkval_inv: Prove clkval_inv from

state_induction {da_prop ← (λ s : lclock_val(s))},
clkval_inv_l1 {s ← x@p1},
clkval_inv_l2 {s ← s@p1, t ← t@p1, u ← u@p1}

p.rtl1: Prove rtl1 from

cum_delta_inv,
state_invariant {da_prop ← (λ s : cum_delta_val(s)), t ← da},
cum_delta_val {s ← da}

p.da_rt_lem: Prove da_rt_lem from

da_rt {p ← p}, $rtl_{x1}^{(x2)}(*3)$ {i ← da.sync_period, p ← p}, rtl1

cum_delta_inv_l1: Lemma initial_da(s) ⊃ cum_delta_val(s)

p.cum_delta_inv_l1: Prove cum_delta_inv_l1 from

initial_da {i ← p@p2}, cum_delta_val, $Corr_{x1}^{(x2)}$ {p ← p@p2, i ← 0}

cum_delta_inv_l2: Lemma

$\mathcal{N}_{da}(s, t, u) \wedge s.phase = sync \wedge cum_delta_val(s) \supset cum_delta_val(t)$

pt, ps: Var period

cdi_l2a: Lemma $pt = ps + 1 \supset Corr_p^{(pt)} = Corr_p^{(ps)} + \Delta_p^{(ps)}$

p.cdi_l2a: Prove cdi_l2a from $Corr_{x1}^{(x2)}$ {i ← pt, p ← p}

p.cum_delta_inv_l2: Prove cum_delta_inv_l2 from

cum_delta_val {s ← s, p ← p@p2},
cum_delta_val {s ← t},
 \mathcal{N}_{da} {i ← p@p2},
 \mathcal{N}_{da}^s {i ← p@p2},
cdi_l2a {p ← p@p2, pt ← t.sync_period, ps ← s.sync_period},
nfc_lem {p ← p@p2, i ← s.sync_period}

cum_delta_inv_l4: Lemma

$\mathcal{N}_{da}(s, t, u) \wedge s.phase \neq sync \wedge cum_delta_val(s) \supset cum_delta_val(t)$

p_cum_delta_inv_l4: Prove cum_delta_inv_l4 from
 cum_delta_val {s ← t},
 cum_delta_val {p ← p@p1},
 N_{da} {i ← p@p1},
 distinct_phases

p_cum_delta_inv: Prove cum_delta_inv from
 state_induction {da_prop ← (λ s : cum_delta_val(s))},
 cum_delta_inv_l1 {s ← x@p1},
 cum_delta_inv_l2 {s ← s@p1, t ← t@p1, u ← u@p1},
 cum_delta_inv_l4 {s ← s@p1, t ← t@p1, u ← u@p1}

End

DA_lemmas: Module

Using DA_to_DS, clkprop

Exporting all with DA_to_DS, clkprop

Theory

ds: Var DSstate
 da: Var DAsate
 k: Var nat
 ph: Var phases
 s, t, x, y, z: Var DAsate
 ss, tt: Var DSstate
 p, q, i, j: Var processors
 u: Var inputs
 w: Var MBvec
 h: Var MBmatrix
 MBfn: Var MBcons.fn
 m, n, a, b: Var proc.plus
 prop: Var function[proc.plus → bool]
 T, T₁, T₂, BB, bb: Var logical_clocktime
 DAsate_prop: Type is function[DAsate → bool]
 da_prop: Var DAsate.prop

phase_com_compute: Lemma
 s.phase = compute ∧ N_{da}(s, t, u) ⊃ N_{da}(DAmapping(s), DAmapping(t), u)

hide1: function[DAsate, DAsate, inputs → bool] =
 (λ s, t, u : (enough_hardware(t)
 ∧ t.phase = next_phase(s.phase)
 ∧ t.sync_period = s.sync_period
 ∧ (∀ i :
 t.proc(i).healthy = s.proc(i).healthy
 ∧ t.proc(i).cum_delta = s.proc(i).cum_delta
 ∧ t.sync_period = s.sync_period
 ∧ (nonfaulty_clock(i, s.sync_period)
 ⊃ clock_advanced(s.proc(i).lclock,
 t.proc(i).lclock,
 duration(s.phase))))
 ∧ N_{da}^c(s, t, u, i))))

phase_com_lx1: Lemma s.phase = compute ∧ N_{da}(s, t, u) ⊃ hide1(s, t, u)

phase_com_lx2: Lemma

$$\begin{aligned} & s.\text{phase} = \text{compute} \\ & \quad \wedge (\text{maj_working}(\text{DAm}ap(t)) \\ & \quad \quad \wedge (\forall i : \\ & \quad \quad \quad \text{DAm}ap(t).\text{phase} = \text{next_phase}(\text{DAm}ap(s).\text{phase}) \\ & \quad \quad \quad \wedge \text{DAm}ap(t).\text{proc}(i).\text{healthy} = \text{DAm}ap(s).\text{proc}(i).\text{healthy} \\ & \quad \quad \quad \wedge \mathcal{N}_{da}^c(\text{DAm}ap(s), \text{DAm}ap(t), u, i))) \\ & \supset \mathcal{N}_{da}(\text{DAm}ap(s), \text{DAm}ap(t), u) \end{aligned}$$

phase_com_lx4: Lemma

$$\begin{aligned} & s.\text{phase} = \text{compute} \\ & \quad \wedge (\text{maj_working}(\text{DAm}ap(t)) \\ & \quad \quad \wedge (\forall i : \\ & \quad \quad \quad t.\text{phase} = \text{next_phase}(s.\text{phase}) \\ & \quad \quad \quad \wedge t.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \wedge \mathcal{N}_{da}^c(s, t, u, i))) \\ & \supset \mathcal{N}_{da}(\text{DAm}ap(s), \text{DAm}ap(t), u) \end{aligned}$$

phase_com_lx7: Lemma

$$\begin{aligned} & s.\text{phase} = \text{compute} \wedge \mathcal{N}_{da}(s, t, u) \\ & \quad \supset (\text{maj_working}(\text{DAm}ap(t)) \\ & \quad \quad \wedge (\forall i : \\ & \quad \quad \quad t.\text{phase} = \text{next_phase}(s.\text{phase}) \\ & \quad \quad \quad \wedge t.\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy} \wedge \mathcal{N}_{da}^c(s, t, u, i))) \end{aligned}$$

phase_com_broadcast: Lemma

$$\text{reachable}(s) \wedge s.\text{phase} = \text{broadcast} \wedge \mathcal{N}_{da}(s, t, u) \supset \mathcal{N}_{da}(\text{DAm}ap(s), \text{DAm}ap(t), u)$$

com_broadcast_1: Lemma

$$s.\text{phase} = \text{broadcast} \wedge \mathcal{N}_{da}(s, t, u) \supset (\forall i : \mathcal{N}_{da}^b(s, t, i))$$

com_broadcast_2: Lemma

$$\begin{aligned} & s.\text{phase} = \text{broadcast} \\ & \quad \wedge \text{reachable}(s) \\ & \quad \quad \wedge s.\text{proc}(i).\text{healthy} = t.\text{proc}(i).\text{healthy} \\ & \quad \quad \quad \wedge \mathcal{N}_{da}(s, t, u) \wedge \mathcal{N}_{da}^b(s, t, i) \\ & \supset \mathcal{N}_{da}^b(\text{DAm}ap(s), \text{DAm}ap(t), i) \end{aligned}$$

com_broadcast_3: Lemma

$$\begin{aligned} & ss.\text{phase} = \text{broadcast} \\ & \quad \wedge tt.\text{phase} = \text{next_phase}(ss.\text{phase}) \\ & \quad \quad \wedge (\forall i : \mathcal{N}_{da}^b(ss, tt, i) \wedge tt.\text{proc}(i).\text{healthy} = ss.\text{proc}(i).\text{healthy}) \\ & \quad \quad \quad \wedge \text{DS.maj_working}(tt) \\ & \supset \mathcal{N}_{da}(ss, tt, u) \end{aligned}$$

com_broadcast_4: Lemma

$$\begin{aligned} & s.\text{phase} = \text{broadcast} \wedge \mathcal{N}_{da}(s, t, u) \\ & \quad \supset \text{DAm}ap(t).\text{phase} = \text{next_phase}(\text{DAm}ap(s).\text{phase}) \\ & \quad \quad \wedge \text{DAm}ap(t).\text{proc}(i).\text{healthy} = \text{DAm}ap(s).\text{proc}(i).\text{healthy} \\ & \quad \quad \quad \wedge \text{DS.maj_working}(\text{DAm}ap(t)) \end{aligned}$$

com_broadcast_5: Lemma

$$\begin{aligned} & \text{reachable}(s) \wedge \mathcal{N}_{da}(s, t, u) \\ & \quad \wedge s.\text{phase} = \text{broadcast} \\ & \quad \quad \wedge s.\text{proc}(i).\text{healthy} > 0 \wedge \text{broadcast_received}(s, t, i) \\ & \supset \text{broadcast_received}(\text{DAm}ap(s), \text{DAm}ap(t), i) \end{aligned}$$

phase_com_vote: Lemma

$$s.\text{phase} = \text{vote} \wedge \mathcal{N}_{da}(s, t, u) \supset \mathcal{N}_{da}(\text{DAm}ap(s), \text{DAm}ap(t), u)$$

com_vote_1: Lemma $s.\text{phase} = \text{vote} \wedge \mathcal{N}_{da}(s, t, u) \supset (\forall i : \mathcal{N}_{da}^v(s, t, i))$

com_vote_2: Lemma $\mathcal{N}_{da}^v(s, t, i) \supset \mathcal{N}_{da}^v(\text{DAm}ap(s), \text{DAm}ap(t), i)$

com_vote_3: Lemma

$$\begin{aligned} & ss.\text{phase} = \text{vote} \wedge tt.\text{phase} = \text{next_phase}(ss.\text{phase}) \\ & \wedge (\forall i : \mathcal{N}_{da}^v(ss, tt, i) \wedge tt.\text{proc}(i).\text{healthy} = ss.\text{proc}(i).\text{healthy}) \\ & \wedge \text{DS.maj_working}(tt) \\ & \supset \mathcal{N}_{da}(ss, tt, u) \end{aligned}$$

com_vote_4: Lemma

$$\begin{aligned} & s.\text{phase} = \text{vote} \wedge \mathcal{N}_{da}(s, t, u) \\ & \supset \text{DAm}ap(t).\text{phase} = \text{next_phase}(\text{DAm}ap(s).\text{phase}) \\ & \wedge \text{DAm}ap(t).\text{proc}(i).\text{healthy} = \text{DAm}ap(s).\text{proc}(i).\text{healthy} \\ & \wedge \text{DS.maj_working}(\text{DAm}ap(t)) \end{aligned}$$

phase_com_sync: Lemma

$$s.\text{phase} = \text{sync} \wedge \mathcal{N}_{da}(s, t, u) \supset \mathcal{N}_{da}(\text{DAm}ap(s), \text{DAm}ap(t), u)$$

com_sync_1: Lemma $s.\text{phase} = \text{sync} \wedge \mathcal{N}_{da}(s, t, u) \supset (\forall i : \mathcal{N}_{da}^s(s, t, i))$

com_sync_2: Lemma $\mathcal{N}_{da}^s(s, t, i) \supset \mathcal{N}_{da}^s(\text{DAm}ap(s), \text{DAm}ap(t), i)$

com_sync_3: Lemma

$$\begin{aligned} & ss.\text{phase} = \text{sync} \wedge tt.\text{phase} = \text{next_phase}(ss.\text{phase}) \\ & \wedge (\forall i : \mathcal{N}_{da}^s(ss, tt, i) \wedge \text{DS.maj_working}(tt)) \\ & \supset \mathcal{N}_{da}(ss, tt, u) \end{aligned}$$

com_sync_4: Lemma

$$\begin{aligned} & s.\text{phase} = \text{sync} \wedge \mathcal{N}_{da}(s, t, u) \\ & \supset \text{DAm}ap(t).\text{phase} = \text{next_phase}(\text{DAm}ap(s).\text{phase}) \wedge \text{DS.maj_working}(\text{DAm}ap(t)) \end{aligned}$$

earliest_later_time: Lemma

$$\begin{aligned} & T_2 = T_1 + BB \wedge (T_1 \geq T^0) \\ & \wedge (BB \geq T^0) \\ & \wedge \text{nonfaulty_clock}(i, da.\text{sync_period}) \\ & \wedge \text{nonfaulty_clock}(j, da.\text{sync_period}) \\ & \wedge \text{enough_clocks}(da.\text{sync_period}) \\ & \wedge T_2 \in R^{(da.\text{sync_period})} \wedge T_1 \in R^{(da.\text{sync_period})} \\ & \supset \text{rt}_i^{(da.\text{sync_period})}(T_2) \\ & \geq \text{rt}_j^{(da.\text{sync_period})}(T_1) + (1 - \text{Rho}) * |BB| - \delta \end{aligned}$$

ELT: Lemma $T_2 \geq T_1 + bb$

$$\begin{aligned} & \wedge (T_1 \geq T^0) \\ & \wedge (bb \geq T^0) \\ & \wedge \text{nonfaulty_clock}(p, da.\text{sync_period}) \\ & \wedge \text{nonfaulty_clock}(q, da.\text{sync_period}) \\ & \wedge \text{enough_clocks}(da.\text{sync_period}) \\ & \wedge T_2 \in R^{(da.\text{sync_period})} \wedge T_1 \in R^{(da.\text{sync_period})} \\ & \supset \text{rt}_p^{(da.\text{sync_period})}(T_2) \\ & \geq \text{rt}_q^{(da.\text{sync_period})}(T_1) + (1 - \text{Rho}) * |bb| - \delta \end{aligned}$$

elt_a: Lemma $(bb \geq T^0) \wedge BB \geq bb \supset (1 - \text{Rho}) * |BB| \geq (1 - \text{Rho}) * |bb|$

map_1: Lemma $\text{DAm}ap(s).\text{proc}(i).\text{healthy} = s.\text{proc}(i).\text{healthy}$

map_2: Lemma $\text{DAm}ap(s).\text{proc}(i).\text{proc_state} = s.\text{proc}(i).\text{proc_state}$

map_3: Lemma $\text{DAm}ap(s).\text{phase} = s.\text{phase}$

map_4: Lemma $\text{DAm}ap(s).\text{proc}(i).\text{mailbox} = s.\text{proc}(i).\text{mailbox}$

map_7: Lemma $\text{DS}(\text{maj_working}(\text{DAm}ap(s))) = \text{DA}(\text{maj_workings}(s))$

support_1: Lemma $\text{initial_da}(s) \supset \text{working_set}(s) = \text{fullset}[\text{processors}]$
support_4: Lemma $s.\text{phase} = ph \wedge \mathcal{N}_{da}(s, x, u) \supset x.\text{phase} = \text{next_phase}(ph)$
support_5: Lemma $s.\text{phase} = ph \wedge ph \neq \text{sync} \wedge \mathcal{N}_{da}(s, x, u)$
 $\supset (\forall i : s.\text{proc}(i).\text{healthy} = x.\text{proc}(i).\text{healthy})$
support_13: Lemma $\text{MBmatrix_cons}(MBfn, \text{nrep})(i) = \text{MBfn}(i)$
support_14: Lemma $\text{initial_da}(s) \supset \text{maj_condition}(\text{working_set}(s))$
support_15: Lemma $\text{initial_da}(s) \supset \text{num_good_clocks}(s.\text{sync_period}, \text{nrep}) = \text{nrep}$
support_16: Lemma $\text{prev_phase}(\text{next_phase}(ph)) = ph$

End

DA_top-proof: **Module**

Using DA_lemmas, DA_invariants

Exporting all with DA_lemmas

Theory

ds: Var DSstate
da: Var DAstate
k: Var nat
ph: Var phases
s, t, x, y, z: Var DAstate
ss, tt: Var DSstate
p, q, i, j: Var processors
u: Var inputs
w: Var MBvec
h: Var MBmatrix
MBfn: Var MBcons_fn
m, n, a, b: Var proc_plus
prop: Var function[proc_plus \rightarrow bool]
T, X, Y: Var logical_clocktime

Proof

p_phase_commutates: Prove phase_commutates from
phase_com_compute,
phase_com_broadcast,
phase_com_vote,
phase_com_sync,
member_phases {phases_var \leftarrow s.phase}

p_initial_maps: Prove initial_maps from
initial_da {i \leftarrow i@p2},
initial_ds {s \leftarrow DAmapping(s)},
map_1 {i \leftarrow i@p2},
map_2 {i \leftarrow i@p2},
map_3

p_phase_com_compute: Prove phase_com_compute from
phase_com_lx4, phase_com_lx7 {i \leftarrow i@p1}

p_phase_com_lx1: Prove phase_com_lx1 from
 \mathcal{N}_{da} {i \leftarrow i@p3}, distinct_phases, hide1

p_phase_com_lx2: Prove phase_com_lx2 $\{i \leftarrow i@p1\}$ from
 $\mathcal{N}_{da} \{s \leftarrow \text{DAm}ap(s), t \leftarrow \text{DAm}ap(t)\}$, distinct_phases, map_3

p_phase_com_lx4: Prove phase_com_lx4 $\{i \leftarrow i@p1\}$ from
phase_com_lx2,
 $\mathcal{N}_{da}^c \{s \leftarrow \text{DAm}ap(s), t \leftarrow \text{DAm}ap(t)\}$,
 $\mathcal{N}_{da}^c \{s \leftarrow s@c, t \leftarrow t@c\}$,
map_1,
map_2,
map_3,
map_4,
map_1 $\{s \leftarrow t\}$,
map_2 $\{s \leftarrow t\}$,
map_3 $\{s \leftarrow t\}$,
map_4 $\{s \leftarrow t\}$

p_phase_com_lx7: Prove phase_com_lx7 from
phase_com_lx1, map_7 $\{s \leftarrow t\}$, hide1, enough_hardware

p_phase_com_broadcast: Prove phase_com_broadcast from
com_broadcast_1 $\{i \leftarrow i@p3\}$,
com_broadcast_2 $\{i \leftarrow i@p3\}$,
com_broadcast_3 $\{ss \leftarrow \text{DAm}ap(s), tt \leftarrow \text{DAm}ap(t)\}$,
com_broadcast_4 $\{i \leftarrow i@p3\}$,
map_1 $\{s \leftarrow s, i \leftarrow i@p2\}$,
map_1 $\{s \leftarrow t, i \leftarrow i@p2\}$,
map_3 $\{\}$

p_com_broadcast_1: Prove com_broadcast_1 from
 \mathcal{N}_{da} , next_phase $\{ph \leftarrow \text{broadcast}\}$, distinct_phases

p_com_broadcast_2: Prove com_broadcast_2 from
com_broadcast_5,
 \mathcal{N}_{da}^b ,
 $\mathcal{N}_{da}^b \{s \leftarrow \text{DAm}ap(s), t \leftarrow \text{DAm}ap(t)\}$,
map_1 $\{s \leftarrow s\}$,
map_1 $\{s \leftarrow t\}$,
map_2 $\{s \leftarrow s\}$,
map_2 $\{s \leftarrow t\}$

p_com_broadcast_3: Prove com_broadcast_3 $\{i \leftarrow i@p1\}$ from
 $\mathcal{N}_{da} \{s \leftarrow ss, t \leftarrow tt\}$, distinct_phases

p_com_broadcast_4: Prove com_broadcast_4 from
 \mathcal{N}_{da} ,
map_1 $\{s \leftarrow s\}$,
map_1 $\{s \leftarrow t\}$,
map_3 $\{s \leftarrow s\}$,
map_3 $\{s \leftarrow t\}$,
map_7 $\{s \leftarrow t\}$,
distinct_phases,
enough_hardware

p_earliest_later_time: Prove earliest_later_time from
GOAL $\{p \leftarrow i, q \leftarrow j, i \leftarrow da.\text{sync_period}\}$

p_elt.a: Prove elt.a from

```
| * 1 | {x ← bb},
| * 1 | {x ← BB},
*1 × *2 {y ← (1 - Rho), x ← |bb|},
*1 × *2 {y ← (1 - Rho), x ← |BB|},
mult_leq {z ← (1 - Rho), x ← |BB|, y ← |bb|}
```

p_ELt: Prove ELT from

```
earliest_later_time {BB ← T2 - T1, i ← p, j ← q@C},
elt_a {BB ← T2 - T1},
*1 × *2 {x ← (1 - Rho), y ← |bb|}
```

p_phase_com_vote: Prove phase_com_vote from

```
com_vote.1 {i ← i@p3},
com_vote.2 {i ← i@p3},
com_vote.3 {ss ← DAmapping(s), tt ← DAmapping(t)},
com_vote.4 {i ← i@p3},
map.3 {}
```

p_com_vote.1: Prove com_vote.1 from \mathcal{N}_{da} , distinct_phases

p_com_vote.2: Prove com_vote.2 from

```
 $\mathcal{N}_{da}^v$  {s ← DAmapping(s), t ← DAmapping(t)},
 $\mathcal{N}_{da}^v$ ,
map.1 {s ← s},
map.1 {s ← t},
map.2 {s ← s},
map.2 {s ← t},
map.4 {s ← s},
map.4 {s ← t}
```

p_com_vote.3: Prove com_vote.3 {i ← i@p1} from

```
 $\mathcal{N}_{da}$  {s ← ss, t ← tt}, distinct_phases
```

p_com_vote.4: Prove com_vote.4 from

```
 $\mathcal{N}_{da}$ ,
enough_hardware,
map.1 {s ← s},
map.1 {s ← t},
map.3 {s ← s},
map.3 {s ← t},
map.7 {s ← t},
distinct_phases
```

p_phase_com_sync: Prove phase_com_sync from

```
com_sync.1 {i ← i@p3},
com_sync.2 {i ← i@p3},
com_sync.3 {ss ← DAmapping(s), tt ← DAmapping(t)},
com_sync.4 {},
map.3 {}
```

p_com_sync.1: Prove com_sync.1 from \mathcal{N}_{da}

p_com_sync.2: Prove com_sync.2 from

```
 $\mathcal{N}_{da}^s$  {s ← DAmapping(s), t ← DAmapping(t)},
 $\mathcal{N}_{da}^s$ ,
map.1 {s ← s},
map.1 {s ← t},
map.2 {s ← s},
map.2 {s ← t}
```

p_com_sync_3: Prove com_sync_3 $\{i \leftarrow i@p1\}$ from $\mathcal{N}_{da} \{s \leftarrow ss, t \leftarrow tt\}$

p_com_sync_4: Prove com_sync_4 from

\mathcal{N}_{da} , enough_hardware, map_3 $\{s \leftarrow s\}$, map_3 $\{s \leftarrow t\}$, map_7 $\{s \leftarrow t\}$

End

DA_map_proof: **Module**

Using DA_lemmas, nat_inductions

Exporting all with DA_lemmas

Theory

ds: Var DSstate
da: Var DAsate
k, q: Var nat
ph: Var phases
s, t, x, y, z: Var DAsate
p, i, j: Var processors
u: Var inputs
w: Var MBvec
h: Var MBmatrix
MBfn: Var MBcons_fn
m, n, a, b: Var proc_plus
prop: Var function[proc_plus \rightarrow bool]

Proof

m1_prop: function[DAsate, processors \rightarrow function[proc_plus \rightarrow bool]] =

($\lambda da, i : (\lambda a :$
 ss_update(*da*, *a*).proc(*i*).healthy
 = if $i \leq a$
 then *da*.proc(*i*).healthy
 else *ds*₀.proc(*i*).healthy
 end if))

m1_base: Lemma m1_prop(*s*, *i*)(0)

m1_ind: Lemma $a < nrep \wedge m1_prop(s, i)(a) \supset m1_prop(s, i)(a + 1)$

p_m1_base: Prove m1_base from

m1_prop $\{da \leftarrow s, i \leftarrow i, a \leftarrow 0\}$,
ss_update $\{da \leftarrow s, k \leftarrow 0\}$

p_m1_ind: Prove m1_ind from

m1_prop $\{da \leftarrow s, i \leftarrow i, a \leftarrow a\}$,
m1_prop
 $\{da \leftarrow s,$
 $i \leftarrow i,$
 $a \leftarrow \text{if } a = nrep \text{ then } nrep \text{ else } a + 1 \text{ end if}\}$,
ss_update $\{da \leftarrow s, k \leftarrow a + 1\}$

p_map_1: Prove map_1 from

DAmapping $\{da \leftarrow s\}$,
processors.induction $\{prop \leftarrow m1_prop(s, i), n \leftarrow nrep\}$,
m1_prop $\{da \leftarrow s, i \leftarrow i, a \leftarrow nrep\}$,
m1_base $\{s \leftarrow s, i \leftarrow i\}$,
m1_ind $\{s \leftarrow s, i \leftarrow i, a \leftarrow m@P2\}$

```

ml2_prop: function[DASate, processors → function[proc_plus → bool]] =
  (λ da, i : (λ a :
    ss_update(da, a).proc(i).proc_state
    = if i ≤ a
      then da.proc(i).proc_state
      else ds0.proc(i).proc_state
    end if))

```

ml2_base: Lemma ml2_prop(s, i)(0)

ml2_ind: Lemma $a < nrep \wedge ml2_prop(s, i)(a) \supset ml2_prop(s, i)(a + 1)$

p_ml2_base: Prove ml2_base from
 ml2_prop {da ← s, i ← i, a ← 0},
 ss_update {da ← s, k ← 0}

p_ml2_ind: Prove ml2_ind from
 ml2_prop {da ← s, i ← i, a ← a},
 ml2_prop
 {da ← s,
 i ← i,
 a ← if a = nrep then nrep else a + 1 end if},
 ss_update {da ← s, k ← a + 1}

p_map_2: Prove map_2 from
 DAmapping {da ← s},
 processors_induction {prop ← ml2_prop(s, i), n ← nrep},
 ml2_prop {da ← s, i ← i, a ← nrep},
 ml2_base {s ← s, i ← i},
 ml2_ind {s ← s, i ← i, a ← m@P2}

p_map_3: Prove map_3 from DAmapping {da ← s}

```

ml4_prop: function[DASate, processors → function[proc_plus → bool]] =
  (λ da, i : (λ a :
    ss_update(da, a).proc(i).mailbox
    = if i ≤ a
      then da.proc(i).mailbox
      else ds0.proc(i).mailbox
    end if))

```

ml4_base: Lemma ml4_prop(s, i)(0)

ml4_ind: Lemma $a < nrep \wedge ml4_prop(s, i)(a) \supset ml4_prop(s, i)(a + 1)$

p_ml4_base: Prove ml4_base from
 ml4_prop {da ← s, i ← i, a ← 0},
 ss_update {da ← s, k ← 0}

p_ml4_ind: Prove ml4_ind from
 ml4_prop {da ← s, i ← i, a ← a},
 ml4_prop
 {da ← s,
 i ← i,
 a ← if a = nrep then nrep else a + 1 end if},
 ss_update {da ← s, k ← a + 1}

p_map-4: Prove map_4 from
 DAmap {da ← s},
 processors_induction {prop ← ml4_prop(s, i), n ← nrep},
 ml4_prop {da ← s, i ← i, a ← nrep},
 ml4_base {s ← s, i ← i},
 ml4_lind {s ← s, i ← i, a ← m@P2}

p_map-7: Prove map_7 from
 proc_extensionality
 {A ← DS(.working_setDAmap(s)),
 B ← DA(.working_sets)},
 DS.maj_working {t ← DAmap(s)},
 DS.working_set {s ← DAmap(s), p ← p@p1},
 DS.working_proc {s ← DAmap(s), p ← p@p1},
 DA.maj_working {t ← s},
 DA.working_set {s ← s, p ← p@p1},
 DA.working_proc {s ← s, p ← p@p1},
 map_1 {i ← p@p1}

End

DA_support_proof: Module

Using DA_lemmas, nat_inductions, DA_invariants

Exporting all with DA_lemmas

Theory

ds: Var DSstate
 da: Var DAsate
 k, q: Var nat
 ph: Var phases
 s, t, x, y, z: Var DAsate
 p, i, j: Var processors
 u: Var inputs
 w: Var MBvec
 h: Var MBmatrix
 MBfn: Var MBcons_fn
 m, n, a, b: Var proc_plus
 prop: Var function[proc_plus → bool]

Proof

p_support_1: Prove support_1 from
 proc_extensionality {A ← working_set(s), B ← fullset[processors]},
 initial_da {i ← p@p1},
 DA.working_set {p ← p@p1},
 DA.working_proc {p ← p@p1}

p_support_4: Prove support_4 from \mathcal{N}_{da} {s ← s, t ← x}

p_support_5: Prove support_5 from
 member_phases {phases_var ← ph},
 \mathcal{N}_{da} {s ← s, t ← x, u ← u, i ← i}

sl13_prop: function[MBcons_fn, processors → function[proc_plus → bool]] =
 (λ MBfn, i : (λ a :
 MBmatrix_cons(MBfn, a)(i)
 = if i ≤ a then MBfn(i) else MBmatrix0(i) end if))

sl13_base: Lemma sl13_prop(MBfn, i)(0)

sl13.ind: Lemma $a < nrep \wedge sl13_prop(MBfn, i)(a) \supset sl13_prop(MBfn, i)(a + 1)$

p.sl13.base: Prove sl13.base from

sl13_prop {a ← 0, i ← i}, MBmatrix_cons {k ← 0}

p.sl13.ind: Prove sl13.ind from

sl13_prop {a ← a, i ← i},
sl13_prop {i ← i, a ← if a = nrep then nrep else a + 1 end if},
MBmatrix_cons {k ← a + 1}

p.support.13: Prove support.13 from

processors_induction {prop ← sl13_prop(MBfn, i), n ← nrep},
sl13_prop {a ← nrep, i ← i},
sl13_base {i ← i},
sl13_ind {i ← i, a ← m@p1}

p.support.14: Prove support.14 from

maj_condition {A ← working_set(s)}, support.1, card_fullset

sl15_prop: function[DAstate → function[nat → bool]] =

(λ s : (λ q :
initial_da(s)
⊃ num_good_clocks(s.sync_period, q)
= if q ≤ nrep then q else 0 end if))

sl15_base: Lemma sl15_prop(s)(0)

sl15.ind: Lemma sl15_prop(s)(q) ⊃ sl15_prop(s)(q + 1)

p.sl15.base: Prove sl15_base from

sl15_prop {s ← s, q ← 0},
num_good_clocks {i ← s.sync_period, k ← 0}

p.sl15.ind: Prove sl15.ind from

sl15_prop {s ← s, q ← q},
sl15_prop {s ← s, q ← q + 1},
num_good_clocks {i ← s.sync_period, k ← q + 1},
initial_da {s ← s, i ← if q < nrep then q + 1 else nrep end if}

p.support.15: Prove support.15 from

nat_induction {p ← sl15_prop(s), n2 ← nrep},
sl15_prop {s ← s, q ← nrep},
sl15_base {s ← s},
sl15_ind {s ← s, q ← n1@p1}

p.support.16: Prove support.16 from

next_phase,
prev_phase {ph ← next_phase(ph)},
distinct_phases,
member_phases {phases_var ← ph}

End

DA.broadcast_prf: Module

Using DA.lemmas, DA.invariants

Exporting all with DA.lemmas

Theory

ds: Var DSstate
da: Var DAsate
k: Var nat
ph: Var phases
r, s, t, x: Var DAsate
ss, tt: Var DSstate
p, q, pp, qq: Var processors
u, u₁, u₂: Var inputs
w: Var MBvec
h: Var MBmatrix
MBfn: Var MBcons_fn
m, n, a, b: Var proc_plus
prop: Var function[proc_plus → bool]
T, X, Y, T₁, T₂, BB: Var logical_clocktime
bb, xx, yy, zz: Var clocktime
Tp, Sq, Rq, Rp, Epsi: Var clocktime

int5: Lemma $r.\text{phase} = \text{compute}$

$\wedge \text{reachable}(r)$
 $\wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2) \wedge \text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset r.\text{proc}(q).\text{lclock} \in R^{(r.\text{sync_period})}$
 $\wedge s.\text{proc}(q).\text{lclock} \in R^{(s.\text{sync_period})}$
 $\wedge t.\text{proc}(q).\text{lclock} \in R^{(t.\text{sync_period})}$

pdurc: Var clocktime

qdurc: Var clocktime

pdurb: Var clocktime

qdurb: Var clocktime

dur: Var clocktime

near: function[clocktime, phases → bool] ==

$(\lambda \text{dur}, \text{ph} : (1 - \nu) * \text{duration}(\text{ph}) \leq \text{dur} \wedge \text{dur} \leq (1 + \nu) * \text{duration}(\text{ph}))$

br1: Lemma $r.\text{phase} = \text{compute} \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$

$\supset (s.\text{phase} = \text{broadcast}$
 $\wedge t.\text{phase} = \text{vote}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\forall pp :$
 $\text{nonfaulty_clock}(pp, r.\text{sync_period})$
 $\supset \text{clock_advanced}(r.\text{proc}(pp).\text{lclock},$
 $s.\text{proc}(pp).\text{lclock},$
 $\text{duration}(\text{compute}))$
 $\wedge \text{clock_advanced}(s.\text{proc}(pp).\text{lclock},$
 $t.\text{proc}(pp).\text{lclock},$
 $\text{duration}(\text{broadcast}))))$

br1a: Lemma $r.\text{phase} = \text{compute} \wedge \mathcal{N}_{da}(r, s, u_1)$

$\supset (s.\text{phase} = \text{broadcast}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge (\forall pp :$
 $\text{nonfaulty_clock}(pp, r.\text{sync_period})$
 $\supset \text{clock_advanced}(r.\text{proc}(pp).\text{lclock},$
 $s.\text{proc}(pp).\text{lclock},$
 $\text{duration}(\text{compute}))))$

br2: Lemma $r.\text{phase} = \text{compute} \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$
 $\supset s.\text{phase} = \text{broadcast}$
 $\wedge t.\text{phase} = \text{vote}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(p, r.\text{sync_period})$
 $\supset (\exists \text{pdurc} :$
 $\text{near}(\text{pdurc}, \text{compute})$
 $\wedge s.\text{proc}(p).\text{lclock} = r.\text{proc}(p).\text{lclock} + \text{pdurc})$
 $\wedge (\exists \text{pdurb} :$
 $\text{near}(\text{pdurb}, \text{broadcast})$
 $\wedge t.\text{proc}(p).\text{lclock} = s.\text{proc}(p).\text{lclock} + \text{pdurb}))$

br3: Lemma $r.\text{phase} = \text{compute} \wedge \text{reachable}(r) \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$
 $\supset s.\text{phase} = \text{broadcast}$
 $\wedge t.\text{phase} = \text{vote}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(p, r.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset r.\text{proc}(p).\text{lclock} = r.\text{proc}(q).\text{lclock}$
 $\wedge (\exists \text{pdurc} :$
 $\text{near}(\text{pdurc}, \text{compute})$
 $\wedge s.\text{proc}(p).\text{lclock} = r.\text{proc}(p).\text{lclock} + \text{pdurc})$
 $\wedge (\exists \text{pdurb} :$
 $\text{near}(\text{pdurb}, \text{broadcast})$
 $\wedge t.\text{proc}(p).\text{lclock} = s.\text{proc}(p).\text{lclock} + \text{pdurb})$
 $\wedge (\exists \text{qdurc} :$
 $\text{near}(\text{qdurc}, \text{compute})$
 $\wedge s.\text{proc}(q).\text{lclock}$
 $= r.\text{proc}(q).\text{lclock} + \text{qdurc})$
 $\wedge (\exists \text{qdurb} :$
 $\text{near}(\text{qdurb}, \text{broadcast})$
 $\wedge t.\text{proc}(q).\text{lclock}$
 $= s.\text{proc}(q).\text{lclock} + \text{qdurb}))$

br3_aa: Lemma $r.\text{phase} = \text{compute}$
 $\wedge \text{reachable}(r)$
 $\wedge \text{nonfaulty_clock}(p, r.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset r.\text{proc}(p).\text{lclock} = r.\text{proc}(q).\text{lclock}$

Proof

p_br1: Prove br1 from

br1a,
 $\mathcal{N}_{da} \{s \leftarrow s, t \leftarrow t, u \leftarrow u_2, i \leftarrow pp\},$
 $\text{next_phase} \{ph \leftarrow \text{broadcast}\},$
 distinct_phases

p_br1a: Prove br1a from

$\mathcal{N}_{da} \{s \leftarrow r, t \leftarrow s, u \leftarrow u_1, i \leftarrow pp\},$
 $\text{next_phase} \{ph \leftarrow \text{compute}\},$
 distinct_phases

p_br2: Prove br2

{pdurc ← s.proc(p).lclock - r.proc(p).lclock,
pdurb ← t.proc(p).lclock - s.proc(p).lclock} from

br1 {pp ← p},

clock_advanced

{X ← r.proc(p).lclock,
Y ← s.proc(p).lclock,
D ← duration(compute)},

clock_advanced

{X ← s.proc(p).lclock,
Y ← t.proc(p).lclock,
D ← duration(broadcast)}

p_br3_aa: Prove br3_aa from

state_invariant {t ← r, da_prop ← (λ s : lclock_eq(s))},

lclock_inv,

lclock_eq {s ← r, i ← p, j ← q}

p_br3: Prove br3

{pdurc ← pdurc@p1,
pdurb ← pdurb@p1,
qdurc ← pdurc@p2,
qdurb ← pdurb@p2} from br2, br2 {p ← q}, br3_aa

br4: Lemma r.phase = compute ∧ reachable(r) ∧ $\mathcal{N}_{da}(r, s, u_1)$ ∧ $\mathcal{N}_{da}(s, t, u_2)$

⊃ s.phase = broadcast

∧ t.phase = vote

∧ s.sync_period = r.sync_period

∧ t.sync_period = s.sync_period

∧ (nonfaulty_clock(p, r.sync_period)

∧ nonfaulty_clock(q, r.sync_period)

∧ Rq = r.proc(q).lclock

∧ Rp = r.proc(p).lclock

∧ Sq = s.proc(q).lclock ∧ Tp = t.proc(p).lclock

⊃ (∃ pdurc, pdurb, qdurc, qdurb :

near(pdurc, compute)

∧ near(pdurb, broadcast)

∧ near(qdurc, compute)

∧ near(qdurb, broadcast)

∧ Rp = Rq

∧ Sq = Rq + qdurc

∧ Tp = Sq - qdurc + pdurc + pdurb))

p_br4: Prove br4

{pdurc ← pdurc@p1,

pdurb ← pdurb@p1,

qdurc ← qdurc@p1,

qdurb ← qdurb@p1} from br3

br5: Lemma $r.\text{phase} = \text{compute} \wedge \text{reachable}(r) \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$
 $\supset s.\text{phase} = \text{broadcast}$
 $\wedge t.\text{phase} = \text{vote}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(p, r.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset s.\text{proc}(q).\text{lclock} \in R^{(s.\text{sync_period})}$
 $\wedge t.\text{proc}(p).\text{lclock} \in R^{(t.\text{sync_period})}$
 $\wedge t.\text{proc}(p).\text{lclock}$
 $\geq s.\text{proc}(q).\text{lclock} + \text{duration}(\text{broadcast})$
 $\quad - 2 * \nu * \text{duration}(\text{compute})$
 $\quad - \nu * \text{duration}(\text{broadcast}))$

p-br5: Prove br5 from

br4

$\{Rq \leftarrow r.\text{proc}(q).\text{lclock},$
 $Rp \leftarrow r.\text{proc}(p).\text{lclock},$
 $Sq \leftarrow s.\text{proc}(q).\text{lclock},$
 $Tp \leftarrow t.\text{proc}(p).\text{lclock}\},$

int5,

int5 $\{q \leftarrow p\}$

br6: Lemma $(\exists r :$

$r.\text{phase} = \text{compute}$

$\wedge \text{reachable}(r) \wedge \mathcal{N}_{da}(r, s, u_1) \wedge s.\text{sync_period} = r.\text{sync_period})$

$\wedge \mathcal{N}_{da}(s, t, u_2)$

$\supset s.\text{phase} = \text{broadcast}$

$\wedge t.\text{phase} = \text{vote}$

$\wedge t.\text{sync_period} = s.\text{sync_period}$

$\wedge (\text{nonfaulty_clock}(p, s.\text{sync_period})$

$\wedge \text{nonfaulty_clock}(q, s.\text{sync_period})$

$\supset s.\text{proc}(q).\text{lclock} \in R^{(s.\text{sync_period})}$

$\wedge t.\text{proc}(p).\text{lclock} \in R^{(t.\text{sync_period})}$

$\wedge t.\text{proc}(p).\text{lclock}$

$\geq s.\text{proc}(q).\text{lclock} + \text{duration}(\text{broadcast})$

$\quad - 2 * \nu * \text{duration}(\text{compute})$

$\quad - \nu * \text{duration}(\text{broadcast}))$

p-br6: Prove br6 from br5

br7: Lemma $\mathcal{N}_{da}(x, s, u)$

$\supset x.\text{phase} = \text{prev_phase}(s.\text{phase})$

$\wedge (x.\text{phase} \neq \text{sync} \supset x.\text{sync_period} = s.\text{sync_period})$

p-br7: Prove br7 from

support_16 $\{ph \leftarrow x.\text{phase}\},$

$\mathcal{N}_{da} \{s \leftarrow x, t \leftarrow s, u \leftarrow u\},$

distinct_phases

br8: Lemma $\text{reachable}(s) \wedge s.\text{phase} = \text{broadcast}$

$\supset (\exists x, u :$

$\mathcal{N}_{da}(x, s, u)$

$\wedge \text{reachable}(x) \wedge x.\text{phase} = \text{compute} \wedge x.\text{sync_period} = s.\text{sync_period})$

p.br8: Prove br8 $\{x \leftarrow s@p2, u \leftarrow u@p2\}$ from
 reachable $\{t \leftarrow s\}$,
 reachable_in_n $\{t \leftarrow s, k \leftarrow k@p1\}$,
 reachable $\{t \leftarrow s@p2, k \leftarrow \text{if } k@p1 = 0 \text{ then } 0 \text{ else } k@p1 - 1 \text{ end if}\}$,
 initial_da $\{s \leftarrow s\}$,
 br7 $\{x \leftarrow s@p2, s \leftarrow s, u \leftarrow u@p2\}$,
 prev_phase $\{ph \leftarrow s.\text{phase}\}$,
 distinct_phases

br9: Lemma $\text{reachable}(s) \wedge \mathcal{N}_{da}(s, t, u_2) \wedge s.\text{phase} = \text{broadcast}$
 $\supset t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(p, s.\text{sync_period}) \wedge \text{nonfaulty_clock}(q, s.\text{sync_period})$
 $\supset s.\text{proc}(q).\text{lclock} \in R^{(s.\text{sync_period})}$
 $\wedge t.\text{proc}(p).\text{lclock} \in R^{(t.\text{sync_period})}$
 $\wedge t.\text{proc}(p).\text{lclock}$
 $\geq s.\text{proc}(q).\text{lclock} + \text{duration}(\text{broadcast})$
 $- 2 * \nu * \text{duration}(\text{compute})$
 $- \nu * \text{duration}(\text{broadcast}))$

p.br9: Prove br9 from br6 $\{r \leftarrow x@p2, u_1 \leftarrow u@p2\}$, br8

rtp0: Lemma $Sq \in R^{(s.\text{sync_period})} \supset Sq \geq 0$

rtp0a: Lemma $T \geq 0 \supset \text{frame_time} * k + T \geq 0$

p.rtp0a: Prove rtp0a from
 mult_non_neg $\{x \leftarrow \text{frame_time}, y \leftarrow k\}$,
 $*1 * *2 \{x \leftarrow \text{frame_time}, y \leftarrow k\}$

p.rtp0: Prove rtp0 from
 $*1 \in R^{(*2)} \{T \leftarrow Sq, i \leftarrow s.\text{sync_period}, \Pi \leftarrow 0\}$,
 $T^{(*1)} \{i \leftarrow s.\text{sync_period}\}$,
 rtp0a $\{T \leftarrow \Pi@p1, k \leftarrow s.\text{sync_period}\}$

rtpl: Lemma $\text{reachable}(s) \wedge \mathcal{N}_{da}(s, t, u_2) \wedge s.\text{phase} = \text{broadcast}$
 $\supset t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(p, t.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(q, s.\text{sync_period})$
 $\wedge \text{enough_clocks}(s.\text{sync_period})$
 $\wedge Tp = t.\text{proc}(p).\text{lclock}$
 $\wedge Sq = s.\text{proc}(q).\text{lclock}$
 $\wedge Epsi$
 $= 2 * \nu * \text{duration}(\text{compute}) + \nu * \text{duration}(\text{broadcast})$
 $\wedge \text{duration}(\text{broadcast}) - Epsi \geq 0$
 $\supset rtp_p^{(s.\text{sync_period})}(Tp)$
 $\geq rtp_q^{(s.\text{sync_period})}(Sq)$
 $+ (1 - \text{Rho}) * |\text{duration}(\text{broadcast}) - Epsi|$
 $- \delta)$

p.rtpl: Prove rtpl from
 rtp0,
 br9,
 ELT
 $\{da \leftarrow s,$
 $T_2 \leftarrow Tp,$
 $T_1 \leftarrow Sq,$
 $q \leftarrow q,$
 $bb \leftarrow \text{duration}(\text{broadcast}) - Epsi\}$

rtp2: Lemma reachable(s)
 $\wedge \mathcal{N}_{da}(s, t, u_2)$
 $\wedge s.\text{phase} = \text{broadcast}$
 $\wedge \text{nonfaulty_clock}(p, s.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(q, s.\text{sync_period})$
 $\wedge \text{enough_clocks}(s.\text{sync_period})$
 $\wedge Tp = t.\text{proc}(p).\text{lclock}$
 $\wedge Sq = s.\text{proc}(q).\text{lclock}$
 $\wedge Epsi$
 $= 2 * \nu * \text{duration}(\text{compute}) + \nu * \text{duration}(\text{broadcast})$
 $\wedge \text{duration}(\text{broadcast}) - Epsi \geq 0$
 $\supset rt_p^{(s.\text{sync_period})}(Tp)$
 $\geq rt_q^{(s.\text{sync_period})}(Sq) + (1 - Rho) * |\text{duration}(\text{broadcast}) - Epsi|$
 $- \delta$

p_rtp2: Prove rtp2 from rtp1

rtp3: Lemma reachable(s)
 $\wedge \mathcal{N}_{da}(s, t, u_2)$
 $\wedge s.\text{phase} = \text{broadcast}$
 $\wedge \text{nonfaulty_clock}(p, s.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(q, s.\text{sync_period})$
 $\wedge \text{enough_clocks}(s.\text{sync_period}) \wedge t.\text{sync_period} = s.\text{sync_period}$
 $\supset rt_p^{(t.\text{sync_period})}(t.\text{proc}(p).\text{lclock})$
 $\geq rt_q^{(s.\text{sync_period})}(s.\text{proc}(q).\text{lclock}) + \text{max_comm_delay}$

p_rtp3: Prove rtp3 from

rtp2
 $\{Epsi \leftarrow 2 * \nu * \text{duration}(\text{compute}) + \nu * \text{duration}(\text{broadcast}),$
 $Sq \leftarrow s.\text{proc}(q).\text{lclock},$
 $Tp \leftarrow t.\text{proc}(p).\text{lclock}\},$
broadcast_duration,
broadcast_duration2

rtp4: Lemma reachable(s)
 $\wedge \mathcal{N}_{da}(s, t, u_2)$
 $\wedge s.\text{phase} = \text{broadcast}$
 $\wedge \text{nonfaulty_clock}(p, s.\text{sync_period})$
 $\wedge \text{nonfaulty_clock}(q, s.\text{sync_period}) \wedge \text{enough_clocks}(s.\text{sync_period})$
 $\supset da_rt(t, p, t.\text{proc}(p).\text{lclock})$
 $\geq da_rt(s, q, s.\text{proc}(q).\text{lclock}) + \text{max_comm_delay}$

rtp4a: Lemma reachable(s) $\wedge \mathcal{N}_{da}(s, t, u_2) \wedge s.\text{phase} = \text{broadcast}$
 $\supset t.\text{sync_period} = s.\text{sync_period}$

p_rtp4a: Prove rtp4a from

$\mathcal{N}_{da} \{s \leftarrow s, t \leftarrow t, u \leftarrow u_2\}, \text{distinct_phases}$

rtp4b: Lemma reachable(s) $\wedge \mathcal{N}_{da}(s, t, u) \supset \text{reachable}(t)$

p_rtp4b: Prove rtp4b from

reachable $\{k \leftarrow k@p3\},$
reachable $\{t \leftarrow s\},$
reachable.in.n $\{k \leftarrow k@p2 + 1, s \leftarrow s, u \leftarrow u\}$

p.rtp4: Prove rtp4 from

rtp3,
rtp4b { $u \leftarrow u_2$ },
rtp4a,
da_rt_lem { $da \leftarrow t, p \leftarrow p, T \leftarrow t.proc(p).lclock$ },
da_rt_lem { $da \leftarrow s, p \leftarrow q, T \leftarrow s.proc(q).lclock$ }

rtp5: Lemma reachable(s)

$\wedge \mathcal{N}_{da}(s, t, u)$
 $\wedge s.phase = broadcast$
 $\wedge s.proc(p).healthy > 0$
 $\wedge broadcast_received(s, t, p)$
 $\wedge (\forall q :$
 $s.proc(q).healthy > 0$
 $\supset da_rt(s, q, s.proc(q).lclock) + max_comm_delay$
 $\leq da_rt(t, p, t.proc(p).lclock)$)
 $\supset broadcast_received(DAmap(s), DAmap(t), p)$

p.rtp5: Prove rtp5 { $q \leftarrow qq@p2$ } from

distinct_phases,
DS.broadcast_received { $s \leftarrow DAmap(s), t \leftarrow DAmap(t), qq \leftarrow q$ },
DA.broadcast_received { $qq \leftarrow qq@p2$ },
map_1 { $s \leftarrow s, i \leftarrow qq@p2$ },
map_4 { $s \leftarrow s, i \leftarrow qq@p2$ },
map_4 { $s \leftarrow t, i \leftarrow p$ },
 \mathcal{N}_{da}

rtp6: Lemma reachable(s) $\wedge s.proc(p).healthy > 0$

$\supset nonfaulty_clock(p, s.sync_period)$

p.rtp6: Prove rtp6 from

nfclk_inv,
state_invariant { $t \leftarrow s, da_prop \leftarrow (\lambda s : nf_clks(s))$ },
nf_clks { $i \leftarrow p$ }

rtp7: Lemma reachable(s) $\wedge \mathcal{N}_{da}(s, t, u) \wedge s.phase = broadcast$

$\supset enough_clocks(s.sync_period) \wedge t.phase = vote$

p.rtp7: Prove rtp7 from

\mathcal{N}_{da} ,
state_invariant { $da_prop \leftarrow (\lambda s : enough_hardware(s)), t \leftarrow s$ },
enough_inv,
enough_hardware { $t \leftarrow s$ },
next_phase { $ph \leftarrow s.phase$ },
distinct_phases

p.com_broadcast_5: Prove com_broadcast_5 from

rtp4 { $u_2 \leftarrow u, q \leftarrow q@p2, p \leftarrow i$ },
rtp5 { $p \leftarrow i$ },
rtp6 { $p \leftarrow i$ },
rtp6 { $p \leftarrow q@p2$ },
rtp7

End

DA_intervals: Module

Using DA_broadcast_prf

Exporting all with DA_lemmas

Theory

ds: Var DSstate
da: Var DAsstate
k: Var nat
ph: Var phases
r, s, t, z: Var DAsstate
ss, tt: Var DSstate
p, q, pp, qq: Var processors
u, u₁, u₂: Var inputs
w: Var MBvec
h: Var MBmatrix
MBfn: Var MBcons.fn
m, n, a, b: Var proc.plus
prop: Var function[proc.plus → bool]
T, X, Y, T₁, T₂, BB: Var logical_clocktime
bb, xx, yy, zz, x₂, y₂: Var clocktime
T_p, S_q, Epsi: Var clocktime
pdurc: Var clocktime
qdurc: Var clocktime
pdurb: Var clocktime
qdurb: Var clocktime
dur: Var clocktime

Proof

br_int: Lemma $r.\text{phase} = \text{compute} \wedge \text{reachable}(r) \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$
 $\supset s.\text{phase} = \text{broadcast}$
 $\wedge t.\text{phase} = \text{vote}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset (\exists \text{qdurc}, \text{qdurb} :$
 $\text{near}(\text{qdurc}, \text{compute})$
 $\wedge \text{near}(\text{qdurb}, \text{broadcast})$
 $\wedge s.\text{proc}(q).\text{lclock} = r.\text{proc}(q).\text{lclock} + \text{qdurc}$
 $\wedge t.\text{proc}(q).\text{lclock} = s.\text{proc}(q).\text{lclock} + \text{qdurb}))$

p_br_int: Prove **br_int** { $\text{qdurc} \leftarrow \text{qdurc}@p1, \text{qdurb} \leftarrow \text{qdurb}@p1$ } from
br3 { $p \leftarrow q$ }

int0: Lemma $r.\text{phase} = \text{compute}$
 $\wedge \text{reachable}(r)$
 $\wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2) \wedge \text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset r.\text{proc}(q).\text{lclock} = r.\text{sync_period} * \text{frame_time}$
 $\wedge r.\text{proc}(q).\text{lclock} \in R^{(r.\text{sync_period})}$

p_int0: Prove **int0** from
 clkval_inv ,
 $\text{state_invariant} \{ \text{da_prop} \leftarrow (\lambda r : \text{lclock_val}(r)), t \leftarrow r \}$,
 $\text{lclock_val} \{ i \leftarrow q, s \leftarrow r \}$,
 $*1 \in R^{(*2)} \{ T \leftarrow r.\text{proc}(q).\text{lclock}, i \leftarrow r.\text{sync_period}, \Pi \leftarrow 0 \}$,
 $T^{(*1)} \{ i \leftarrow r.\text{sync_period} \}$

int1: Lemma $r.\text{phase} = \text{compute} \wedge \text{reachable}(r) \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$
 $\supset s.\text{phase} = \text{broadcast}$
 $\wedge t.\text{phase} = \text{vote}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset r.\text{proc}(q).\text{lclock} = r.\text{sync_period} * \text{frame_time}$
 $\wedge r.\text{proc}(q).\text{lclock} \in R^{(r.\text{sync_period})}$
 $\wedge s.\text{proc}(q).\text{lclock} \in R^{(s.\text{sync_period})})$

int1a: Lemma $xx \leq yy \wedge yy \leq zz \supset xx \leq zz$

p.int1a: Prove int1a

p.int1: Prove int1 from

int0,
br_int,
 $*1 \in R^{(*2)}$
 $\{T \leftarrow s.\text{proc}(q).\text{lclock},$
 $i \leftarrow s.\text{sync_period},$
 $\Pi \leftarrow \text{qdurc}@p2\},$
 $T^{(*1)} \{i \leftarrow s.\text{sync_period}\},$
pos_durations,
all_durations,
int1a
 $\{zz \leftarrow \text{qdurc}@p2,$
 $yy \leftarrow (1 - \nu) * \text{duration}(\text{compute}),$
 $xx \leftarrow 0\},$
int1a
 $\{xx \leftarrow \text{qdurc}@p2,$
 $yy \leftarrow (1 + \nu) * \text{duration}(\text{compute}),$
 $zz \leftarrow \text{frame_time}\}$

int2: Lemma $r.\text{phase} = \text{compute} \wedge \text{reachable}(r) \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$
 $\supset s.\text{phase} = \text{broadcast}$
 $\wedge t.\text{phase} = \text{vote}$
 $\wedge s.\text{sync_period} = r.\text{sync_period}$
 $\wedge t.\text{sync_period} = s.\text{sync_period}$
 $\wedge (\text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\wedge r.\text{proc}(q).\text{lclock} = r.\text{sync_period} * \text{frame_time}$
 $\supset t.\text{proc}(q).\text{lclock} \in R^{(t.\text{sync_period})})$

int2a: Lemma $\text{near}(\text{qdurc}, \text{compute}) \wedge \text{near}(\text{qdurc}, \text{broadcast})$
 $\supset 0 \leq \text{qdurc} + \text{qdurc} \wedge \text{qdurc} + \text{qdurc} \leq \text{frame_time}$

p.int2a: Prove int2a from

pos_durations,
all_durations,
 $*1 \times *2 \{x \leftarrow (1 - \nu), y \leftarrow \text{duration}(\text{compute})\},$
 $*1 \times *2 \{x \leftarrow (1 - \nu), y \leftarrow \text{duration}(\text{broadcast})\}$

p.int2: Prove int2 from

$T^{(*1)} \{i \leftarrow t.\text{sync_period}\},$
br_int,
 $*1 \in R^{(*2)}$
 $\{T \leftarrow t.\text{proc}(q).\text{lclock},$
 $i \leftarrow t.\text{sync_period},$
 $\Pi \leftarrow \text{qdurc}@p2 + \text{qdurc}@p2\},$
int2a $\{\text{qdurc} \leftarrow \text{qdurc}@p2, \text{qdurc} \leftarrow \text{qdurc}@p2\}$

int3: **Lemma** $r.\text{phase} = \text{compute} \wedge \text{reachable}(r) \wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2)$
 $\supset (\text{nonfaulty_clock}(q, r.\text{sync_period})$
 $\supset r.\text{proc}(q).\text{lclock} = r.\text{sync_period} * \text{frame_time}$
 $\wedge r.\text{proc}(q).\text{lclock} \in R^{(r.\text{sync_period})}$
 $\wedge s.\text{proc}(q).\text{lclock} \in R^{(s.\text{sync_period})}$
 $\wedge t.\text{proc}(q).\text{lclock} \in R^{(t.\text{sync_period})})$

p.int3: **Prove** int3 from int1, int2

int4: **Lemma** $(r.\text{phase} = \text{compute}$
 $\wedge \text{reachable}(r)$
 $\wedge \mathcal{N}_{da}(r, s, u_1) \wedge \mathcal{N}_{da}(s, t, u_2) \wedge \text{nonfaulty_clock}(q, r.\text{sync_period}))$
 $\supset (r.\text{proc}(q).\text{lclock} = r.\text{sync_period} * \text{frame_time}$
 $\wedge r.\text{proc}(q).\text{lclock} \in R^{(r.\text{sync_period})}$
 $\wedge s.\text{proc}(q).\text{lclock} \in R^{(s.\text{sync_period})}$
 $\wedge t.\text{proc}(q).\text{lclock} \in R^{(t.\text{sync_period})})$

p.int4: **Prove** int4 from int3

p.int5: **Prove** int5 from int4

End

clk_types: **Module**

Exporting all

Theory

realtime: **Type is number**
logical_clocktime: **Type is number**
physical_clocktime: **Type is number**
clocktime: **Type is number**
x: **Var number**
posnum: **Type from number with** $(\lambda x : x > 0)$
pos_logical_clocktime: **Type is posnum**
posrealtime: **Type is posnum**
fraction: **Type from number with** $(\lambda x : 1 \geq x \wedge x \geq 0 \wedge x \neq 1)$
period: **Type is nat**

End

clkmod: **Module**

Using rcp_defs, absmod, clk_types

Exporting all with rcp_defs, clk_types, absmod **Theory**

$\epsilon, \delta_0, \delta$: posrealtime
 Σ, Δ : pos_logical_clocktime
frame_time, sync_time: pos_logical_clocktime (* Changed from R, S *)
i: **Var period**
k: **Var nat**
 T^0 : logical_clocktime == 0
 $T^{(*)}$: function[period \rightarrow logical_clocktime] = $(\lambda i : T^0 + i * \text{frame_time})$
T_next: **Lemma** $T^{(i+1)} = T^{(i)} + \text{frame_time}$

T, Π : **Var** logical_clocktime
 T_1, T_2, T_0, T_N : **Var** physical_clocktime
 $*1 \in R^{(*2)}$: function[logical_clocktime, period \rightarrow boolean] =
 $(\lambda T, i : (\exists \Pi : 0 \leq \Pi \wedge \Pi \leq \text{frame_time} \wedge T = T^{(i)} + \Pi))$
 $*1 \in S^{(*2)}$: function[logical_clocktime, period \rightarrow boolean] =
 $(\lambda T, i : (\exists \Pi :$
 $0 \leq \Pi \wedge \Pi \leq \text{sync_time} \wedge T = T^{(i)} + \text{frame_time} - \text{sync_time} + \Pi))$
 p, q, r : **Var** processors
 $c_{*1}(*2)$: function[processors, physical_clocktime \rightarrow realtime]
log_to_phys: function[logical_clocktime \rightarrow physical_clocktime] ==
 $(\lambda T \rightarrow \text{physical_clocktime} : T)$
 x : **Var** number
 $\frac{x}{2}$: function[number \rightarrow number] == $(\lambda x : x/2)$
 ρ : fraction
Rho: fraction = $\frac{\rho}{2}$

goodclock: function[processors, physical_clocktime, physical_clocktime
 \rightarrow bool] =

$(\lambda p, T_0, T_N :$
 $(\forall T_1, T_2 :$
 $T_0 \leq T_1 \wedge T_0 \leq T_2 \wedge T_1 \leq T_N \wedge T_2 \leq T_N$
 $\supset |c_p(T_1) - c_p(T_2) - (T_1 - T_2)| \leq \text{Rho} * |T_1 - T_2|))$

monotonicity: **Theorem**

$(\exists T_0, T_N : \text{goodclock}(p, T_0, T_N) \wedge T_0 \leq T_1 \wedge T_0 \leq T_2 \wedge T_1 \leq T_N \wedge T_2 \leq T_N$
 $\supset (T_1 > T_2 \supset c_p(T_1) \geq c_p(T_2))$)

$\Delta_{*1}^{(*2)}$: function[processors, period \rightarrow clocktime]
(* mean of the skews within tolerance *)

Delta2: function[processors, processors, period \rightarrow clocktime]
(* measured skew *)

initial_Corr: function[processors \rightarrow clocktime] == $(\lambda p \rightarrow \text{number} : 0)$

second_arg: function[processors, period \rightarrow nat] == $(\lambda p, i : i)$

$\text{Corr}_{*1}^{(*2)}$: Recursive function[processors, period \rightarrow clocktime] =
 $(\lambda p, i : \text{if } i > 0$

then $\text{Corr}_p^{(\text{pred}(i))} + \Delta_p^{(\text{pred}(i))}$

else initial_Corr(p)

end if) by second_arg

$A_{*1}^{(*2)}(*3)$: function[processors, period, logical_clocktime
 \rightarrow physical_clocktime] == $(\lambda p, i, T : T + \text{Corr}_p^{(i)})$

$rt_{*1}^{(*2)}(*3)$: function[processors, period, logical_clocktime \rightarrow realtime] =
 $(\lambda p, i, T : c_p(A_p^{(i)}(T)))$

skew: function[processors, processors, clocktime, period \rightarrow clocktime] ==
 $(\lambda p, q, T, i \rightarrow \text{clocktime} : |rt_p^{(i)}(T) - rt_q^{(i)}(T)|)$

nonfaulty_clock: function[processors, period \rightarrow boolean] =
 $(\lambda p, i : \text{goodclock}(p, A_p^{(0)}(T^{(0)}), A_p^{(i)}(T^{(i+1)})))$

num.measure: function[period, nat \rightarrow nat] == (λ i, k : k
num.good_clocks: Recursive function[period, nat \rightarrow nat] =
 (λ i, k : if k = 0 \vee k > nrep
 then 0
 elseif nonfaulty_clock(k, i)
 then 1 + num.good_clocks(i, k - 1)
 else num.good_clocks(i, k - 1)
 end if) by num.measure

enough_clocks: function[period \rightarrow bool] =
 (λ i : 3 * num.good_clocks(i, nrep) > 2 * nrep)

S1A: function[period \rightarrow bool] == (λ i : enough_clocks(i))
 (* in current clock sync theory =
 (LAMBDA i :
 (FORALL r : (m + 1 <= r AND r <= n) IMPLIES nonfaulty_clock(r, i)))
 *)

S1C: function[processors, processors, period \rightarrow bool] =
 (λ p, q, i : ($\forall T$:
 nonfaulty_clock(p, i) \wedge nonfaulty_clock(q, i) $\wedge T \in R^{(i)}$
 \supset skew(p, q, T, i) $\leq \delta$)

S1C_lemma: Lemma S1C(p, q, i) \supset S1C(q, p, i)

S1: function[period \rightarrow bool] = (λ i : S1A(i) \supset ($\forall p, q$: S1C(p, q, i)))

S2: function[processors, period \rightarrow bool] =
 (λ p, i : (|Corr_p⁽ⁱ⁺¹⁾ - Corr_p⁽ⁱ⁾| < Σ))

(* The following three theorems were proved in the clock sync theory.
 They are taken as axioms here. *)
adj_always_pos: Axiom $A_p^{(k)}(T^{(k)}) \geq T^0$

Theorem.1: Axiom $S_1(i)$

(* THEOREM *)

Theorem.2: Axiom $S_2(p, i)$

(* THEOREM *)

A0: Axiom skew(p, q, T⁽⁰⁾, 0) < δ_0

A1: Lemma nonfaulty_clock(p, i) = goodclock(p, A_p⁽⁰⁾(T⁽⁰⁾), A_p⁽ⁱ⁾(T⁽ⁱ⁺¹⁾))

A2: Axiom nonfaulty_clock(p, i)
 \wedge nonfaulty_clock(q, i) \wedge S1C(p, q, i) \wedge S2(p, i)
 \supset $|\Delta_{qp}^{(i)}| \leq \text{sync_time}$
 \wedge ($\exists T_0 : T_0 \in S^{(i)} \wedge |\tau t_p^{(i)}(T_0 + \Delta_{qp}^{(i)}) - \tau t_q^{(i)}(T_0)| < \epsilon$)

A2_aux: Axiom $\Delta_{pp}^{(i)} = 0$

m: processors (* maximum number of faulty clocks *)

C0: Axiom $m < \text{nrep} \wedge m \leq \text{nrep} - \text{num_good_clocks}(i, \text{nrep})$

C1: Axiom frame_time $\geq 3 * \text{sync_time}$

C2: **Axiom** $\text{sync_time} \geq \Sigma$

C3: **Axiom** $\Sigma \geq \Delta$

C4: **Axiom** $\Delta \geq \delta + \epsilon + \frac{\rho}{2} * \text{sync_time}$

C5: **Axiom** $\delta \geq \delta_0 + \rho * \text{frame_time}$

C6: **Axiom** $\delta \geq 2 * (\epsilon + \rho * \text{sync_time}) + 2 * m * \Delta / (\text{nrep} - m)$
 $+ \text{nrep} * \rho * \text{frame_time} / (\text{nrep} - m)$
 $+ \rho * \Delta$
 $+ \text{nrep} * \rho * \Sigma / (\text{nrep} - m)$

sync_thm: Theorem

enough_clocks(i)

$\supset (\forall p, q :$
 $(\forall T : \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T \in R^{(i)}$
 $\supset |rt_p^{(i)}(T) - rt_q^{(i)}(T)| \leq \delta))$

Proof

p_sync_thm: Prove sync_thm from

Theorem.1 $\{i \leftarrow i\}$, **S1** $\{i \leftarrow i\}$, **S1C** $\{i \leftarrow i\}$

End

clkprop: Module

Using clkmod, DA

Exporting all

Theory

$T, T_1, T_2, T_3, T_4, BB, T_0, T_N, TX, TY$: **Var** logical_clocktime

p, q : **Var** processors

da : **Var** DAstate

i : **Var** period

ft2: Lemma $\text{goodclock}(q, T^0, T_1 + BB) \wedge (T_1 \geq T^0) \wedge (BB \geq T^0)$
 $\supset |c_q(T_1 + BB) - c_q(T_1) - BB| \leq \text{Rho} * |BB|$

ft3: Lemma $\text{goodclock}(q, T^0, T_1 + BB) \wedge (T_1 \geq T^0) \wedge (BB \geq T^0)$
 $\supset (1 - \text{Rho}) * |BB| \leq c_q(T_1 + BB) - c_q(T_1)$
 $\wedge c_q(T_1 + BB) - c_q(T_1) \leq (1 + \text{Rho}) * |BB|$

ft4: Lemma $\text{enough_clocks}(i)$
 $\wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T \in R^{(i)}$
 $\supset -\delta \leq rt_p^{(i)}(T) - rt_q^{(i)}(T) \wedge rt_p^{(i)}(T) - rt_q^{(i)}(T) \leq \delta$

ft5: Lemma $\text{goodclock}(q, T^0, T_1 + \text{Corr}_q^{(i)} + BB)$
 $\wedge (T_1 \geq T^0) \wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0) \wedge (BB \geq T^0)$
 $\supset (1 - \text{Rho}) * |BB| \leq rt_q^{(i)}(T_1 + BB) - rt_q^{(i)}(T_1)$
 $\wedge rt_q^{(i)}(T_1 + BB) - rt_q^{(i)}(T_1) \leq (1 + \text{Rho}) * |BB|$

ft6: Lemma $T_2 = T_1 + BB$
 $\wedge \text{goodclock}(q, T^0, T_1 + \text{Corr}_q^{(i)} + BB)$
 $\wedge (T_1 \geq T^0)$
 $\wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0)$
 $\wedge (BB \geq T^0)$
 $\wedge \text{enough_clocks}(i)$
 $\wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T_2 \in R^{(i)}$
 $\supset rt_p^{(i)}(T_2) \geq rt_q^{(i)}(T_1) + (1 - \text{Rho}) * |BB| - \delta$

ft7: Lemma $T_3 \leq T_4 \wedge \text{goodclock}(q, T^0, T_4) \supset \text{goodclock}(q, T^0, T_3)$

ft8: Lemma $T_1 + BB \leq T^{(i+1)} \wedge \text{nonfaulty_clock}(q, i)$
 $\supset \text{goodclock}(q, T^0, T_1 + \text{Corr}_q^{(i)} + BB)$

ft9: Lemma $T_2 = T_1 + BB$
 $\wedge T_2 \leq T^{(i+1)}$
 $\wedge (T_1 \geq T^0)$
 $\wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0)$
 $\wedge (BB \geq T^0)$
 $\wedge \text{enough_clocks}(i)$
 $\wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T_2 \in R^{(i)}$
 $\supset \text{rt}_p^{(i)}(T_2) \geq \text{rt}_q^{(i)}(T_1) + (1 - \text{Rho}) * |BB| - \delta$

ft10: Lemma $T_2 \in R^{(i)} \supset T_2 \leq T^{(i+1)}$

ft11: Lemma $T_2 = T_1 + BB$
 $\wedge (T_1 \geq T^0)$
 $\wedge (T_1 + \text{Corr}_q^{(i)} \geq T^0)$
 $\wedge (BB \geq T^0)$
 $\wedge \text{enough_clocks}(i)$
 $\wedge \text{nonfaulty_clock}(p, i) \wedge \text{nonfaulty_clock}(q, i) \wedge T_2 \in R^{(i)}$
 $\supset \text{rt}_p^{(i)}(T_2) \geq \text{rt}_q^{(i)}(T_1) + (1 - \text{Rho}) * |BB| - \delta$

ft12: Lemma $T_1 \in R^{(i)} \supset (T_1 + \text{Corr}_q^{(i)} \geq T^0)$

GOAL: Lemma $T_2 = T_1 + BB$
 $\wedge (T_1 \geq T^0)$
 $\wedge (BB \geq T^0)$
 $\wedge \text{nonfaulty_clock}(p, i)$
 $\wedge \text{nonfaulty_clock}(q, i) \wedge \text{enough_clocks}(i) \wedge T_2 \in R^{(i)} \wedge T_1 \in R^{(i)}$
 $\supset \text{rt}_p^{(i)}(T_2) \geq \text{rt}_q^{(i)}(T_1) + (1 - \text{Rho}) * |BB| - \delta$

nfc_lem: Lemma $\text{nonfaulty_clock}(p, i + 1) \supset \text{nonfaulty_clock}(p, i)$

Proof

nfc.a: Lemma $T^{(i+1)} + \text{Corr}_p^{(i)} \leq T^{(i+2)} + \text{Corr}_p^{(i+1)}$

p_nfc.a: Prove nfc.a from

$T^{(i+1)} \{i \leftarrow i + 1\},$

$T^{(i+1)} \{i \leftarrow i + 2\},$

Theorem.2 $\{i \leftarrow i\},$

S2 $\{i \leftarrow i\},$

abs_main $\{x \leftarrow \text{Corr}_p^{(i+1)} - \text{Corr}_p^{(i)}, z \leftarrow \Sigma\},$

C1,

C2

p_nfc_lem: Prove nfc_lem from

$\text{nonfaulty_clock},$

$\text{nonfaulty_clock} \{i \leftarrow i + 1\},$

goodclock

$\{T_N \leftarrow T^{(i+2)} + \text{Corr}_p^{(i+1)},$

$T_0 \leftarrow T^{(0)} + \text{Corr}_p^{(0)},$

$T_1 \leftarrow T_1 @ p4,$

$T_2 \leftarrow T_2 @ p4\},$

$\text{goodclock} \{T_N \leftarrow T^{(i+1)} + \text{Corr}_p^{(i)}, T_0 \leftarrow T^{(0)} + \text{Corr}_p^{(0)}\},$

nfc.a

p.ft2: Prove ft2 from
goodclock

$\{p \leftarrow q,$
 $T_0 \leftarrow T^0,$
 $T_N \leftarrow T_1 + BB,$
 $T_2 \leftarrow T_1,$
 $T_1 \leftarrow (T_1 + BB)\}$

p.ft3: Prove ft3 from
ft2,

$\text{abs_leq } \{x \leftarrow c_q(T_1 + BB) - c_q(T_1) - BB, z \leftarrow \text{Rho} * |BB|\},$
 $\text{abs_ge0 } \{x \leftarrow BB\}$

p.ft4: Prove ft4 from

$\text{sync_thm } \{i \leftarrow i\}, \text{abs_leq } \{x \leftarrow \text{rt}_p^{(i)}(T) - \text{rt}_q^{(i)}(T), z \leftarrow \delta\}$

p.ft5: Prove ft5 from

$\text{ft3 } \{T_1 \leftarrow T_1 + \text{Corr}_q^{(i)}\},$
 $\text{rt}_{*1}^{(*2)}(*3) \{p \leftarrow q, T \leftarrow T_1, i \leftarrow i\},$
 $\text{rt}_{*1}^{(*2)}(*3) \{p \leftarrow q, T \leftarrow T_1 + BB, i \leftarrow i\}$

p.ft6: Prove ft6 from ft4 $\{T \leftarrow T_2\}$, ft5

p.ft7: Prove ft7 from

goodclock
 $\{p \leftarrow q,$
 $T_0 \leftarrow T^0,$
 $T_1 \leftarrow T_1 \oplus p2,$
 $T_2 \leftarrow T_2 \oplus p2,$
 $T_N \leftarrow T_4\},$
goodclock $\{p \leftarrow q, T_0 \leftarrow T^0, T_N \leftarrow T_3\}$

ft8a: Lemma $\text{Corr}_q^{(0)} = 0$

p.ft8: Prove ft8 from

$\text{nonfaulty_clock } \{p \leftarrow q, i \leftarrow i\},$
 $\text{ft7 } \{T_3 \leftarrow T_1 + \text{Corr}_q^{(i)} + BB, T_4 \leftarrow T^{(i+1)} + \text{Corr}_q^{(i)}\},$
 $T^{(*1)} \{i \leftarrow 0\},$
ft8a

p.ft8a: Prove ft8a from $\text{Corr}_{*1}^{(*2)} \{i \leftarrow 0, p \leftarrow q\}$

p.ft9: Prove ft9 from ft6, ft8

p.ft10: Prove ft10 from

$*1 \in R^{(*2)} \{T \leftarrow T_2, \Pi \leftarrow \text{frame_time}\}, T^{(*1)} \{i \leftarrow i\}, T^{(*1)} \{i \leftarrow i + 1\}$

p.ft11: Prove ft11 from ft10 $\{i \leftarrow i\}$, ft9

p.ft12: Prove ft12 from

$\text{adj_always_pos } \{k \leftarrow i, p \leftarrow q\}, *1 \in R^{(*2)} \{T \leftarrow T_1, i \leftarrow i\}$

p.GOAL: Prove GOAL from ft11, ft12

End

DA_invariants.tcc: **Module**

Using DA_invariants

Exporting all with DA_invariants

Theory

```

ii: Var naturalnumber
p: Var rcp_defs.processors
j: Var rcp_defs.processors
i: Var rcp_defs.processors
z: Var DA.DAstate
n1: Var naturalnumber
k: Var naturalnumber
u: Var rcp_defs.inputs
t: Var DA.DAstate
s: Var DA.DAstate
Corr_lem_TCC1: Formula (ii > 0)  $\supset$  (ii - 1  $\geq$  0)

```

Proof

```
Corr_lem_TCC1.PROOF: Prove Corr_lem_TCC1
```

End DA_invariants_tcc

DA_map_proof.tcc: **Module**

Using DA_map_proof

Exporting all with DA_map_proof

Theory

```

a: Var rcp_defs.proc_plus
p: Var rcp_defs.processors
m: Var rcp_defs.proc_plus  p_mll_base_TCC1: Formula ((0  $\geq$  0)  $\wedge$  (0  $\leq$  nrep))

```

```

p_mll_ind_TCC1: Formula
  (( if a = nrep then nrep else a + 1 end if  $\geq$  0)
     $\wedge$  ( if a = nrep then nrep else a + 1 end if  $\leq$  nrep))

```

```
p_map_1_TCC1: Formula ((nrep  $\geq$  0)  $\wedge$  (nrep  $\leq$  nrep))
```

Proof

```
p_mll_base_TCC1.PROOF: Prove p_mll_base_TCC1
```

```
p_mll_ind_TCC1.PROOF: Prove p_mll_ind_TCC1
```

```
p_map_1_TCC1.PROOF: Prove p_map_1_TCC1
```

End DA_map_proof.tcc

DA_support_proof.tcc: **Module**

Using DA_support_proof

Exporting all with DA_support_proof

Theory

```

q: Var naturalnumber
a: Var rcp_defs.proc_plus
n1: Var naturalnumber
m: Var rcp_defs.proc_plus
p: Var rcp_defs.processors
p_sl13_base_TCC1: Formula ((0  $\geq$  0)  $\wedge$  (0  $\leq$  nrep))

```

```

p_sl13_ind_TCC1: Formula
  (( if a = nrep then nrep else a + 1 end if  $\geq$  0)
     $\wedge$  ( if a = nrep then nrep else a + 1 end if  $\leq$  nrep))

```

p_sl13_base_TCC1: **Formula** $((nrep \geq 0) \wedge (nrep \leq nrep))$
p_sl15_ind_TCC1: **Formula**
 $((\text{if } q < nrep \text{ then } q + 1 \text{ else } nrep \text{ end if } > 0)$
 $\wedge (\text{if } q < nrep \text{ then } q + 1 \text{ else } nrep \text{ end if } \leq nrep))$

Proof

p_sl13_base_TCC1_PROOF: **Prove** p_sl13_base_TCC1
p_sl13_ind_TCC1_PROOF: **Prove** p_sl13_ind_TCC1
p_support_13_TCC1_PROOF: **Prove** p_support_13_TCC1
p_sl15_ind_TCC1_PROOF: **Prove** p_sl15_ind_TCC1

End DA_support_proof.tcc

DA_to_DS.tcc: **Module**

Using DA_to_DS

Exporting all with DA_to_DS

Theory

da: **Var** DA.DAstate
s: **Var** DA.DAstate
t: **Var** DA.DAstate
k: **Var** naturalnumber
MBfn: **Var** function[rcp.defs.processors \rightarrow rcp.defs.MBvec]
ss_update_TCC1: **Formula** $(\neg((k = 0) \vee (k > nrep))) \supset (k - 1 \geq 0)$
ss_update_TCC2: **Formula** $(\neg((k = 0) \vee (k > nrep))) \supset ((k > 0) \wedge (k \leq nrep))$
ss_update_TCC3: **Formula**
 $(\neg((k = 0) \vee (k > nrep))) \supset da_measure(da, k) > da_measure(da, k - 1)$
MBmatrix_cons_TCC1: **Formula**
 $(\neg((k = 0) \vee (k > nrep)))$
 $\supset MBmc_measure(MBfn, k) > MBmc_measure(MBfn, k - 1)$
reachable_in_n_TCC1: **Formula** $(\neg(k = 0)) \supset (k - 1 \geq 0)$
reachable_in_n_TCC2: **Formula**
 $(\neg(k = 0)) \supset da_measure(t, k) > da_measure(s, k - 1)$

Proof

ss_update_TCC1_PROOF: **Prove** ss_update_TCC1
ss_update_TCC2_PROOF: **Prove** ss_update_TCC2
ss_update_TCC3_PROOF: **Prove** ss_update_TCC3
MBmatrix_cons_TCC1_PROOF: **Prove** MBmatrix_cons_TCC1
reachable_in_n_TCC1_PROOF: **Prove** reachable_in_n_TCC1
reachable_in_n_TCC2_PROOF: **Prove** reachable_in_n_TCC2

End DA_to_DS.tcc

DA_tcc_proof: **Module**

Using clk_types.tcc, clkmod.tcc, DA_map_proof.tcc, DA_support_proof.tcc

Exporting all

Theory

Proof

posnum_TCC1.PROOF: Prove posnum_TCC1 { $x \leftarrow 1$ }
fraction_TCC1.PROOF: Prove fraction_TCC1 { $x \leftarrow 0$ }
C6_TCC1.PROOF: Prove C6_TCC1 from C0
p_sl15_ind_TCC1.PROOF: Prove p_sl15_ind_TCC1 from processors_exist_ax
Rho_TCC1.PROOF: Prove Rho_TCC1 (* needs printerpdivide = yes *)

End

DA_broadcast_prf_tcc: **Module**

Using DA_broadcast_prf

Exporting all with DA_broadcast_prf

Theory

i: **Var** rcp_defs.processors
q: **Var** rcp_defs.processors
qq: **Var** rcp_defs.processors
ll: **Var** number
x: **Var** DA.DAstate
k: **Var** naturalnumber
u: **Var** rcp_defs.inputs
s: **Var** DA.DAstate
qdur: **Var** number
pdur: **Var** number
pdurc: **Var** number
p_br8_TCC1: **Formula** (if $k = 0$ then 0 else $k - 1$ end if ≥ 0)

Proof

p_br8_TCC1.PROOF: Prove p_br8_TCC1

End DA_broadcast_prf_tcc

clk_types_tcc: **Module**

Using clk_types

Exporting all with clk_types

Theory

x: **Var** number
posnum_TCC1: **Formula** ($\exists x : x > 0$)
fraction_TCC1: **Formula** ($\exists x : 1 \geq x \wedge x \geq 0 \wedge x \neq 1$)

Proof

posnum_TCC1.PROOF: Prove posnum_TCC1

fraction_TCC1.PROOF: Prove fraction_TCC1

End clk_types.tcc

clkmod.tcc: **Module**

Using clkmod

Exporting all with clkmod

Theory

i: **Var** naturalnumber

k: **Var** naturalnumber

p: **Var** rcp_defs.processors

half.TCC1: **Formula** ($2 \neq 0$)

Rho.TCC1: **Formula** ($1 \geq \frac{\rho}{2} \wedge \frac{\rho}{2} \geq 0 \wedge \frac{\rho}{2} \neq 1$)

Corr.TCC1: **Formula** ($i > 0$) \supset $\text{second_arg}(p, i) > \text{second_arg}(p, \text{pred}(i))$

num_good_clocks.TCC1: **Formula**

$(\neg(k = 0 \vee k > \text{nrep})) \supset ((k > 0) \wedge (k \leq \text{nrep}))$

num_good_clocks.TCC2: **Formula**

$(\text{nonfaulty_clock}(k, i) \wedge (\neg(k = 0 \vee k > \text{nrep}))) \supset (k - 1 \geq 0)$

num_good_clocks.TCC3: **Formula**

$(\neg(\text{nonfaulty_clock}(k, i))) \wedge (\neg(k = 0 \vee k > \text{nrep})) \supset (k - 1 \geq 0)$

num_good_clocks.TCC4: **Formula**

$(\text{nonfaulty_clock}(k, i) \wedge (\neg(k = 0 \vee k > \text{nrep})))$
 $\supset \text{num_measure}(i, k) > \text{num_measure}(i, k - 1)$

num_good_clocks.TCC5: **Formula**

$(\neg(\text{nonfaulty_clock}(k, i))) \wedge (\neg(k = 0 \vee k > \text{nrep}))$
 $\supset \text{num_measure}(i, k) > \text{num_measure}(i, k - 1)$

C6.TCC1: **Formula** $((\text{nrep} - m) \neq 0)$

Proof

half.TCC1.PROOF: **Prove** half.TCC1

Rho.TCC1.PROOF: **Prove** Rho.TCC1

Corr.TCC1.PROOF: **Prove** Corr.TCC1

num_good_clocks.TCC1.PROOF: **Prove** num_good_clocks.TCC1

num_good_clocks.TCC2.PROOF: **Prove** num_good_clocks.TCC2

num_good_clocks.TCC3.PROOF: **Prove** num_good_clocks.TCC3

num_good_clocks.TCC4.PROOF: **Prove** num_good_clocks.TCC4

num_good_clocks.TCC5.PROOF: **Prove** num_good_clocks.TCC5

C6.TCC1.PROOF: **Prove** C6.TCC1

End clkmod.tcc

top: **Module**

Using rcp_defs, generic_FT, sets[processors], cardinality[processors],
nat_inductions, noetherian[proc_plus, lessp], US, RS, RS_majority, RS_to_US,
RS_lemmas, RS_invariants, RS_top_proof, RS_tcc_proof, rcp_defs_tcc,
RS_to_US_tcc, DS, DS_to_RS, DS_lemmas, DS_top_proof, DS_map_proof,
DS_support_proof, multiplication, absmod, clk_types, clkmod, DA, DA_to_DS,
DA_invariants, clkprop, DA_lemmas, DA_top_proof, DA_map_proof,
DA_support_proof, DA_broadcast_prf, DA_intervals, rcp_defs_tcc,
DS_to_RS_tcc, DS_support_proof_tcc, DS_map_proof_tcc, DA_invariants_tcc,
DA_map_proof_tcc, DA_support_proof_tcc, DA_to_DS_tcc, clk_types_tcc,
clkmod_tcc, DA_tcc_proof, DA_broadcast_prf_tcc

Theory

u: Var inputs
us1, us2: Var Pstate
rs1, rs2: Var RSstate
ds1, ds2: Var DSstate
da1, da2: Var DAsate

RS.frame_commutates: Theorem

$\text{reachable}(rs1) \wedge \mathcal{N}_{rs}(rs1, rs2, u) \supset \mathcal{N}_{us}(RSmap(rs1), RSmap(rs2), u)$

RS.initial_maps: Theorem $\text{initial_rs}(rs1) \supset \text{initial_us}(RSmap(rs1))$

DS.frame_commutates: Theorem

$\text{ds1.phase} = \text{compute} \wedge \text{frame_N_ds}(ds1, ds2, u)$
 $\supset \mathcal{N}_{rs}(DSmap(ds1), DSmap(ds2), u)$

DS.initial_maps: Theorem $\text{initial_ds}(ds1) \supset \text{initial_rs}(DSmap(ds1))$

DA.phase_commutates: Theorem

$\text{reachable}(da1) \wedge \mathcal{N}_{da}(da1, da2, u) \supset \mathcal{N}_{ds}(DAmmap(da1), DAmmap(da2), u)$

DA.initial_maps: Theorem $\text{initial_da}(da1) \supset \text{initial_ds}(DAmmap(da1))$

Proof

p-RS.frame_commutates: Prove RS.frame_commutates from
RS_to_US.frame_commutates {s ← rs1, t ← rs2}

p-RS.initial_maps: Prove RS.initial_maps from
RS_to_US.initial_maps {s ← rs1}

p-DS.frame_commutates: Prove DS.frame_commutates from
DS_to_RS.frame_commutates {s ← ds1, t ← ds2}

p-DS.initial_maps: Prove DS.initial_maps from
DS_to_RS.initial_maps {s ← ds1}

p-DA.phase_commutates: Prove DA.phase_commutates from
DA_to_DS.phase_commutates {s ← da1, t ← da2}

p-DA.initial_maps: Prove DA.initial_maps from
DA_to_DS.initial_maps {s ← da1}

End

rcp_defs: Module

Exporting all

Theory

```

p: Var nat
Pstate: Type (* computation state of a single processor *)
inputs: Type (* type of external sensor input *)
outputs: Type (* actuator output type *)
MB: Type (* mailbox exchange type *)
nrep: nat (* number of replicated processors *)
initial_proc_state: Pstate (* assumes each processor begins identically *)
recovery_period: nat (* number of healthy frames required to recover
from transient fault plus one *)
recovery_period_ax: Axiom recovery_period > 2

processors_exist_ax: Axiom nrep > 0

processors: Type from nat with ( $\lambda p : (p > 0) \wedge (p \leq nrep)$ )
MBvec: Type = array [processors] of MB
MBmatrix: Type = array [processors] of MBvec
phases: Type = (compute, broadcast, vote, sync)
ph: Var phases
next_phase: function[phases  $\rightarrow$  phases] =
( $\lambda ph : \text{if } ph = \text{compute}$ 
then broadcast
elseif  $ph = \text{broadcast}$  then vote elseif  $ph = \text{vote}$  then sync else compute
end if)
prev_phase: function[phases  $\rightarrow$  phases] =
( $\lambda ph : \text{if } ph = \text{compute}$ 
then sync
elseif  $ph = \text{broadcast}$ 
then compute
elseif  $ph = \text{vote}$  then broadcast else vote
end if)
proc_plus: Type from nat with ( $\lambda p : (p \geq 0) \wedge (p \leq nrep)$ )
k, m, a, n, b: Var proc_plus
prop: Var function[proc_plus  $\rightarrow$  bool]
lessp: function[proc_plus, proc_plus  $\rightarrow$  bool] == ( $\lambda m, n : m < n$ )

processors_induction: Lemma
( $\forall prop : \text{prop}(0) \wedge (\forall m : m < nrep \wedge \text{prop}(m) \supset \text{prop}(m + 1))$ 
 $\supset (\forall n : \text{prop}(n))$ )

```

Proof

Using noetherian[proc_plus, lessp]

reachability: Lemma $a \neq 0 \Leftrightarrow (\exists b : a = b + 1)$

p-processors_induction: Prove processors_induction $\{m \leftarrow b @ P2\}$ from
general_induction $\{p \leftarrow prop, d \leftarrow n, d_2 \leftarrow m\}$,
reachability $\{a \leftarrow d_1 @ P1\}$

p-well_founded: Prove well_founded $\{\text{measure} \leftarrow (\lambda k \rightarrow \text{nat} : k)\}$

p_reachability: Prove reachability $\{b \leftarrow \text{if } a = 0 \text{ then } 0 \text{ else } a - 1 \text{ end if}\}$

End

sets: Module [T: Type]

Exporting all

Theory

```

set: Type is function[T → bool]
x, y, z: Var T
a, b: Var set
*1 ∪ *2: function[set, set → set] == (λ a, b : (λ x : a(x) ∨ b(x)))
*1 ∩ *2: function[set, set → set] == (λ a, b : (λ x : a(x) ∧ b(x)))
*1 \ *2: function[set, set → set] == (λ a, b : (λ x : a(x) ∧ ¬b(x)))
add: function[T, set → set] == (λ x, a : (λ y : x = y ∨ a(y)))
singleton: function[T → set] == (λ x : (λ y : y = x))
*1 ⊂ *2: function[set, set → bool] = (λ a, b : (∀ z : a(z) ⊃ b(z)))
*1 ∈ *2: function[T, set → bool] == (λ x, b : b(x))
empty: function[set → bool] = (λ a : (∀ x : ¬a(x)))
φ: set == (λ x : false)
fullset: set == (λ x : true)

```

extensionality: Axiom $(\forall x : x \in a = x \in b) \supset (a = b)$

End sets

cardinality: Module [T: Type]

Using sets[T]

Exporting all

Assuming

```

x, y, z: Var T
N: Var nat
f: Var function[T → nat]

```

finite: Formula $(\exists N, f : (\forall x, y : f(x) \leq N \wedge (f(x) = f(y) \supset x = y)))$

Theory

```

a, b, c: Var set
card: function[set → nat]

```

card_ax: Axiom $\text{card}(a \cup b) + \text{card}(a \cap b) = \text{card}(a) + \text{card}(b)$

card_subset: Axiom $a \subset b \supset \text{card}(a) \leq \text{card}(b)$

card_empty: Axiom $\text{card}(a) = 0 \Leftrightarrow \text{empty}(a)$

empty_prop: Lemma $\text{card}(a) > 0 \supset (\exists x : x \in a)$

card_prop: Lemma $a \subset c \wedge b \subset c \wedge 2 * \text{card}(a) > \text{card}(c) \wedge 2 * \text{card}(b) > \text{card}(c) \supset \text{card}(a \cap b) > 0$

Proof

empty_prop_proof: Prove empty_prop {x ← x@p2} from card_empty, empty

subset_union: Sublemma $a \subset c \wedge b \subset c \supset a \cup b \subset c$

subset_union_proof: Prove subset_union from

```

*1 ⊂ *2 {z ← z@p3, b ← c},
*1 ⊂ *2 {z ← z@p3, a ← b, b ← c},
*1 ⊂ *2 {a ← a ∪ b, b ← c}

```

m, n, p: Var nat

twice_prop: Sublemma $2 * m > p \wedge 2 * n > p \supset m + n > p$

twice_proof: Prove twice_prop

card_proof: Prove card_prop from
twice_prop {m ← card(a), n ← card(b), p ← card(c)},
card_ax,
subset_union,
card_subset {a ← a ∪ b, b ← c}

End cardinality

nat_inductions: Module

Theory

i, j: Var nat
n₁, n₂, n₃: Var nat
p: Var function[nat → bool]

nat_complete: Axiom
 $(\forall n_1 : (\forall n_3 : (n_3 \neq n_1) \supset p(n_3)) \supset p(n_1)) \supset (\forall n_2 : p(n_2))$

nat_induction: Axiom $(p(0) \wedge (\forall n_1 : p(n_1) \supset p(n_1 + 1))) \supset (\forall n_2 : p(n_2))$

nat_induct_by_2: Axiom
 $(p(0) \wedge p(1) \wedge (\forall n_1 : p(n_1) \supset p(n_1 + 2))) \supset (\forall n_2 : p(n_2))$

End nat_inductions

noetherian: Module [dom: Type, <: function[dom, dom → bool]]

Assuming

measure: Var function[dom → nat]
a, b: Var dom

well_founded: Formula $(\exists \text{measure} : a < b \supset \text{measure}(a) < \text{measure}(b))$

Theory

p, A, B: Var function[dom → bool]
d, d₁, d₂: Var dom

general_induction: Axiom
 $(\forall d_1 : (\forall d_2 : d_2 < d_1 \supset p(d_2)) \supset p(d_1)) \supset (\forall d : p(d))$

End noetherian

multiplication: Module

Exporting all

Theory

x, y, z, x₁, y₁, z₁, x₂, y₂, z₂: Var number
*1 × *2: function[number, number → number] = (λ x, y : (x + y))

mult_ldistrib: Lemma $x \times (y + z) = x \times y + x \times z$

mult_ldistrib_minus: Lemma $x \times (y - z) = x \times y - x \times z$

mult_rident: Lemma $x \times 1 = x$

mult_lident: Lemma $1 \times x = x$

distrib: Lemma $(x + y) \times z = x \times z + y \times z$

distrib_minus: Lemma $(x - y) \times z = x \times z - y \times z$

mult_non_neg: Axiom $((x \geq 0 \wedge y \geq 0) \vee (x \leq 0 \wedge y \leq 0)) \Leftrightarrow x \times y \geq 0$
mult_pos: Axiom $((x > 0 \wedge y > 0) \vee (x < 0 \wedge y < 0)) \Leftrightarrow x \times y > 0$
mult_com: Lemma $x \times y = y \times x$
pos_product: Lemma $x \geq 0 \wedge y \geq 0 \supset x \times y \geq 0$
mult_leq: Lemma $z \geq 0 \wedge x \geq y \supset x \times z \geq y \times z$
mult_leq_2: Lemma $z \geq 0 \wedge x \geq y \supset z \times x \geq z \times y$
mult_l0: Axiom $0 \times x = 0$
mult_gt: Lemma $z > 0 \wedge x > y \supset x \times z > y \times z$

Proof

mult_gt_pr: Prove mult_gt from
 mult_pos $\{x \leftarrow x - y, y \leftarrow z\}$, distrib_minus
distrib_minus_pr: Prove distrib_minus from
 mult_ldistrib_minus $\{x \leftarrow z, y \leftarrow x, z \leftarrow y\}$,
 mult_com $\{x \leftarrow x - y, y \leftarrow z\}$,
 mult_com $\{y \leftarrow z\}$,
 mult_com $\{x \leftarrow y, y \leftarrow z\}$
mult_leq_2_pr: Prove mult_leq_2 from
 mult_ldistrib_minus $\{x \leftarrow z, y \leftarrow x, z \leftarrow y\}$,
 mult_non_neg $\{x \leftarrow z, y \leftarrow x - y\}$
mult_leq_pr: Prove mult_leq from
 distrib_minus, mult_non_neg $\{x \leftarrow x - y, y \leftarrow z\}$
mult_com_pr: Prove mult_com from $*1 \times *2, *1 \times *2 \{x \leftarrow y, y \leftarrow x\}$
pos_product_pr: Prove pos_product from mult_non_neg
mult_riident_proof: Prove mult_riident from $*1 \times *2 \{y \leftarrow 1\}$
mult_liident_proof: Prove mult_liident from $*1 \times *2 \{x \leftarrow 1, y \leftarrow x\}$
distrib_proof: Prove distrib from
 $*1 \times *2 \{x \leftarrow x + y, y \leftarrow z\}$,
 $*1 \times *2 \{y \leftarrow z\}$,
 $*1 \times *2 \{x \leftarrow y, y \leftarrow z\}$
mult_ldistrib_proof: Prove mult_ldistrib from
 $*1 \times *2 \{y \leftarrow y + z, x \leftarrow x\}, *1 \times *2, *1 \times *2 \{y \leftarrow z\}$
mult_ldistrib_minus_proof: Prove mult_ldistrib_minus from
 $*1 \times *2 \{y \leftarrow y - z, x \leftarrow x\}, *1 \times *2, *1 \times *2 \{y \leftarrow z\}$

End

absmod: Module

Using multiplication

Exporting all

Theory

$x, y, z, x_1, y_1, z_1, x_2, y_2, z_2$: Var number
 $|*1|$: Definition function[number \rightarrow number] =
 $(\lambda x : (\text{if } x < 0 \text{ then } -x \text{ else } x \text{ end if}))$

abs_main: Lemma $|x| < z \supset (x < z \wedge -x < z)$
abs_leq_0: Lemma $|x - y| \leq z \supset (x - y) \leq z$
abs_diff: Lemma $|x - y| < z \supset ((x - y) < z \wedge (y - x) < z)$
abs_leq: Lemma $|x| \leq z \supset (x \leq z \wedge -x \leq z)$
abs_bnd: Lemma $0 \leq z \wedge 0 \leq x \wedge x \leq z \wedge 0 \leq y \wedge y \leq z \supset |x - y| \leq z$
abs_1_bnd: Lemma $|x - y| \leq z \supset x \leq y + z$
abs_2_bnd: Lemma $|x - y| \leq z \supset x \geq y - z$
abs_3_bnd: Lemma $x \leq y + z \wedge x \geq y - z \supset |x - y| \leq z$
abs_drift: Lemma $|x - y| \leq z \wedge |x_1 - x| \leq z_1 \supset |x_1 - y| \leq z + z_1$
abs_com: Lemma $|x - y| = |y - x|$
abs_drift_2: Lemma
 $|x - y| \leq z \wedge |x_1 - x| \leq z_1 \wedge |y_1 - y| \leq z_2 \supset |x_1 - y_1| \leq z + z_1 + z_2$
abs_geq: Lemma $x \geq y \wedge y \geq 0 \supset |x| \geq |y|$
abs_ge0: Lemma $x \geq 0 \supset |x| = x$
abs_plus: Lemma $|x + y| \leq |x| + |y|$
abs_diff_3: Lemma $x - y \leq z \wedge y - x \leq z \supset |x - y| \leq z$
abs_eq: Lemma $|x - y| = |y - x|$

Proof

abs_plus_pr: Prove abs_plus from $|*1| \{x \leftarrow x + y\}, |*1|, |*1| \{x \leftarrow y\}$
abs_diff_3_pr: Prove abs_diff_3 from $|*1| \{x \leftarrow x - y\}$
abs_ge0_proof: Prove abs_ge0 from $|*1|$
abs_geq_proof: Prove abs_geq from $|*1|, |*1| \{x \leftarrow y\}$
abs_drift_2_proof: Prove abs_drift_2 from
abs_drift,
abs_drift $\{x \leftarrow y, y \leftarrow y_1, z \leftarrow z_2, z_1 \leftarrow z + z_1\}$,
abs_com $\{x \leftarrow y_1\}$
abs_com_proof: Prove abs_com from $|*1| \{x \leftarrow (x - y)\}, |*1| \{x \leftarrow (y - x)\}$
abs_drift_proof: Prove abs_drift from
abs_1_bnd,
abs_1_bnd $\{x \leftarrow x_1, y \leftarrow x, z \leftarrow z_1\}$,
abs_2_bnd,
abs_2_bnd $\{x \leftarrow x_1, y \leftarrow x, z \leftarrow z_1\}$,
abs_3_bnd $\{x \leftarrow x_1, z \leftarrow z + z_1\}$
abs_3_bnd_proof: Prove abs_3_bnd from $|*1| \{x \leftarrow (x - y)\}$
abs_main_proof: Prove abs_main from $|*1|$
abs_leq_0_proof: Prove abs_leq_0 from $|*1| \{x \leftarrow x - y\}$

```

abs.diff-proof: Prove abs.diff from |* 1| {x ← (x - y)}
abs.leq-proof: Prove abs.leq from |* 1|
abs.bnd-proof: Prove abs.bnd from |* 1| {x ← (x - y)}
abs.1_bnd-proof: Prove abs.1_bnd from |* 1| {x ← (x - y)}
abs.2_bnd-proof: Prove abs.2_bnd from |* 1| {x ← (x - y)}

End absmod

rcp_defs_tcc: Module
Using rcp_defs
Exporting all with rcp_defs

Theory

p: Var naturalnumber
m: Var proc_plus
a: Var proc_plus
prop: Var function[proc_plus → boolean]
d1: Var proc_plus
b: Var proc_plus (* Existence TCC generated for processors *)
processors_TCC1: Formula (∃ p : (p > 0) ∧ (p ≤ nrep))
proc_plus_TCC1: Formula (∃ p : (p ≥ 0) ∧ (p ≤ nrep))
processors_induction_TCC1: Formula ((0 ≥ 0) ∧ (0 ≤ nrep))
processors_induction_TCC2: Formula
  (m < nrep ∧ prop(m)) ∧ (prop(0)) ⊃ ((m + 1 ≥ 0) ∧ (m + 1 ≤ nrep))
p_reachability_TCC1: Formula
  ( if a = 0 then 0 else a - 1 end if ≥ 0)
  ∧ (( if a = 0 then 0 else a - 1 end if ≥ 0)
    ∧ ( if a = 0 then 0 else a - 1 end if ≤ nrep))

Proof

processors_TCC1_PROOF: Prove processors_TCC1
proc_plus_TCC1_PROOF: Prove proc_plus_TCC1
processors_induction_TCC1_PROOF: Prove processors_induction_TCC1
processors_induction_TCC2_PROOF: Prove processors_induction_TCC2
p_reachability_TCC1_PROOF: Prove p_reachability_TCC1

End rcp_defs_tcc

```

Appendix B

LaTeX-printed Supplementary Specification Listings

rcp_defs: Module

(* This rcp_defs module differs slightly from the original. Several definitions have been moved to new modules; the originals have been commented out. *)

Exporting all

Theory

```
p: Var nat
inputs: Type (* type of external sensor input *)
outputs: Type (* actuator output type *)
nrep: nat (* number of replicated processors *)
recovery_period: nat (* number of healthy frames required to recover
from transient fault plus one *)
recovery_period_ax: Axiom recovery_period > 2

processors_exist_ax: Axiom nrep > 0

processors: Type from nat with (λ p : (p > 0) ∧ (p ≤ nrep))
phases: Type = (compute, broadcast, vote, sync)
ph: Var phases
next_phase: function[phases → phases] =
  (λ ph : if ph = compute
    then broadcast
    elsif ph = broadcast then vote elsif ph = vote then sync else compute
    end if)
prev_phase: function[phases → phases] =
  (λ ph : if ph = compute
    then sync
    elsif ph = broadcast
    then compute
    elsif ph = vote then broadcast else vote
    end if)
proc_plus: Type from nat with (λ p : (p ≥ 0) ∧ (p ≤ nrep))
k, m, a, n, b: Var proc_plus
prop: Var function[proc_plus → bool]
lessp: function[proc_plus, proc_plus → bool] == (λ m, n : m < n)

processors_induction: Lemma
  (∀ prop : prop(0) ∧ (∀ m : m < nrep ∧ prop(m) ⊃ prop(m + 1))
    ⊃ (∀ n : prop(n)))
```

Proof

Using noetherian[proc_plus, lessp]

reachability: Lemma $a \neq 0 \Leftrightarrow (\exists b : a = b + 1)$

p_processors_induction: Prove processors_induction {m ← b@P2} from
general_induction {p ← prop, d ← n, d₂ ← m},
reachability {a ← d₁@P1}

p_well_founded: Prove well_founded {measure ← (λ k → nat : k)}

p_reachability: Prove reachability {b ← if a = 0 then 0 else a - 1 end if}

End

task_model: Module

(* This module introduces an interpretation for a basic task-oriented style of computation state. It is common to both the continuous voting and cyclic voting interpretations. *)

Using rcp_defs, sets[processors], cardinality[processors], nat_inductions

Exporting all with rcp_defs, sets[processors], cardinality[processors]

Theory

p, i, j : Var processors

k, l, q : Var nat

u : Var inputs

A : Var set[processors] (* Basic definitions for schedules *)

maj_condition: function[set[processors] → bool] =

($\lambda A : 2 * \text{card}(A) > \text{card}(\text{fullset}[\text{processors}])$)

schedule_length: nat (* Number of frames in schedule cycle *)

schedule_length.ax: Axiom schedule_length > 0

control_state: Type from nat with ($\lambda k : k < \text{schedule_length}$)

K, L : Var control_state

mod_plus: function[control_state, control_state → control_state] =

($\lambda K, L \rightarrow \text{control_state} :$

if $K + L \geq \text{schedule_length}$

then $K + L - \text{schedule_length}$

else $K + L$

end if)

mod_minus: function[control_state, control_state → control_state] =

($\lambda K, L \rightarrow \text{control_state} :$

if $K \geq L$ then $K - L$ else $\text{schedule_length} - L + K$ end if)

num_cells: nat

num_cells.ax: Axiom num_cells > 0

cell: Type from nat with ($\lambda k : k < \text{num_cells}$)

cell_state: Type

cell_array: Type = array [cell] of cell_state

c, d, e : Var cell

H : Var nat

C, D : Var cell_array

(* Task schedule concepts. Each cell occupies a unique place in the schedule, being computed only once per schedule cycle. *)

cell_frame: function[cell → control_state] (* scheduled frame of cell *)

cell_subframe: function[cell → nat] (* scheduled subframe of cell *)

sched_cell: function[control_state, nat → cell] (* cell of frame, subframe *)

num_subframes: function[control_state → nat] (* subframes for this frame *)

(* Well-formedness axioms constraining these functions *)

cell_frame_ax: Axiom $c = \text{sched_cell}(K, k) \supset \text{cell_frame}(c) = K$

cell_subframe_ax: Axiom $c = \text{sched_cell}(K, k) \supset \text{cell_subframe}(c) = k$

```

sched_cell_ax: Axiom
   $K = \text{cell\_frame}(c) \wedge k = \text{cell\_subframe}(c) \supset \text{sched\_cell}(K, k) = c$ 

num_subframes_ax: Axiom
   $K = \text{cell\_frame}(c) \supset \text{cell\_subframe}(c) < \text{num\_subframes}(K)$ 

(* Processor state definition *)
Pstate: Type = Record control : control_state,
                      cells : cell_array
                      end record

null_cell_array: cell_array (* default value *)
initial_proc_state: Pstate (* assumes each processor begins identically *)
MB: Type is Pstate
MBvec: Type = array [processors] of MB
MBmatrix: Type = array [processors] of MBvec
w: Var MBvec
h: Var MBmatrix
us, ps, X, Y: Var Pstate

cell_array_equal: Axiom  $(\forall c : C(c) = D(c)) \supset C = D$ 

Pstate_equal: Axiom  $(X.\text{control} = Y.\text{control} \wedge X.\text{cells} = Y.\text{cells}) \supset X = Y$ 

(* Interpretations for task-related functions *)
succ: function[control_state  $\rightarrow$  control_state] =
   $(\lambda K \rightarrow \text{control\_state} :$ 
    if  $K + 1 < \text{schedule\_length}$  then  $K + 1$  else 0 end if)
fk: function[Pstate  $\rightarrow$  control_state] ==  $(\lambda \text{ps} : \text{ps}.\text{control})$ 
fi: function[Pstate, cell  $\rightarrow$  cell_state] ==  $(\lambda \text{ps}, c : \text{ps}.\text{cells}(c))$ 
(* Functions modeling task execution *)

exec_task: function[inputs, control_state, cell_array, nat  $\rightarrow$  cell_state]

exec_measure: function[inputs, control_state, cell_array, nat  $\rightarrow$  nat] ==
   $(\lambda u, K, C, k : k)$ 
exec: Recursive function[inputs, control_state, cell_array, nat
   $\rightarrow$  cell_array] =
   $(\lambda u, K, C, k :$ 
    if  $k = 0$ 
    then  $C$ 
    else  $\text{exec}(u, K, C, k - 1)$ 
    with  $[(\text{sched\_cell}(K, k - 1)) :=$ 
       $\text{exec\_task}(u, K, \text{exec}(u, K, C, k - 1), k - 1)]$ 
    end if)
    by  $\text{exec\_measure}$ 

fc: function[inputs, Pstate  $\rightarrow$  Pstate] =
   $(\lambda u, \text{ps} : \text{ps}$  with  $[(\text{control}) := \text{succ}(\text{ps}.\text{control}),$ 
     $(\text{cells}) := \text{exec}(u, \text{ps}.\text{control},$ 
     $\text{ps}.\text{cells}, \text{num\_subframes}(\text{ps}.\text{control}))])$ 

fa: function[Pstate  $\rightarrow$  outputs] (* actuator output *)

(* Axioms to be satisfied by the generic application *)

succ_ax: Formula  $f_k(f_c(u, \text{ps})) = \text{succ}(f_k(\text{ps}))$ 

```

components_equal: Formula $f_k(X) = f_k(Y) \wedge (\forall c : f_i(X, c) = f_i(Y, c)) \supset X = Y$

(* Support lemmas *)

succ_le_plus: Lemma $\text{succ}(K) \leq K + 1$

mod_minus_zero: Lemma $\text{mod_minus}(K, L) = 0 \Leftrightarrow K = L$

mod_minus_succ: Lemma $\text{mod_minus}(\text{succ}(K), L) = \text{succ}(\text{mod_minus}(K, L))$

mod_minus_plus: Lemma $\text{succ}(K) \neq L \supset \text{mod_minus}(\text{succ}(K), L) = \text{mod_minus}(K, L) + 1$

exec_element: Lemma

```
exec(u, K, C, num_subframes(K))(c)
= if cell.frame(c) = K
  then exec_task(u, K, exec(u, K, C, cell_subframe(c)), cell_subframe(c))
  else C(c)
end if
```

Proof

p_succ_ax: Prove succ_ax from f_c

p_components_equal: Prove components_equal $\{c \leftarrow c@p1\}$ from
cell_array_equal $\{C \leftarrow X.\text{cells}, D \leftarrow Y.\text{cells}\}$, Pstate_equal

p_succ_le_plus: Prove succ_le_plus from succ

p_mod_minus_zero: Prove mod_minus_zero from mod_minus

p_mod_minus_succ: Prove mod_minus_succ from
mod_minus $\{K \leftarrow \text{succ}(K)\}$, mod_minus, succ $\{K \leftarrow \text{mod_minus}(K, L)\}$, succ

p_mod_minus_plus: Prove mod_minus_plus from
mod_minus $\{K \leftarrow \text{succ}(K)\}$, mod_minus, succ

exe_prop: function[inputs, control_state, cell_array, cell, nat
→ function[nat → bool]] =

```
(λ u, K, C, c, k :
  (λ q : cell_subframe(c) = k ∧ q ≤ num_subframes(K)
    ⊃ exec(u, K, C, q)(c)
    = if k < q ∧ cell.frame(c) = K
      then exec_task(u, K, exec(u, K, C, k), k)
      else C(c)
    end if))
```

exe_base: Lemma $\text{exe_prop}(u, K, C, c, k)(0)$

exe_ind_1: Lemma $\text{exe_prop}(u, K, C, c, k)(q) \wedge \text{cell_subframe}(c) = q$
⊃ $\text{exe_prop}(u, K, C, c, k)(q + 1)$

exe_ind_2: Lemma $\text{exe_prop}(u, K, C, c, k)(q) \wedge \text{cell_subframe}(c) \neq q$
⊃ $\text{exe_prop}(u, K, C, c, k)(q + 1)$

p_exe_base: Prove exe_base from exe_prop $\{q \leftarrow 0\}$, exec $\{k \leftarrow 0\}$

p_exe_ind_1: Prove exe_ind_1 from

```
exe_prop {q ← q},
exe_prop {q ← q + 1},
exec {k ← q + 1},
sched_cell_ax {k ← q},
cell_frame_ax {k ← q},
num_subframes_ax
```

p_exe_ind_2: Prove exe_ind_2 from
 exe_prop {q ← q},
 exe_prop {q ← q + 1},
 exec {k ← q + 1},
 cell_frame_ax {k ← q},
 cell_subframe_ax {k ← q},
 num_subframes_ax

p_exec_element: Prove exec_element from
 nat_induction
 {p ← exe_prop(u, K, C, c, cell_subframe(c)),
 n₂ ← num_subframes(K)},
 exe_prop {q ← num_subframes(K), k ← cell_subframe(c)},
 exe_base {k ← cell_subframe(c)},
 exe_ind_1 {q ← n₁@p1, k ← cell_subframe(c)},
 exe_ind_2 {q ← n₁@p1, k ← cell_subframe(c)},
 num_subframes_ax

End

cont_voting: Module

(* Following is the interpretation for the continuous voting scheme. *)

Using task_model, nat_inductions

Exporting all with task_model

Theory

us, ps, X, Y: Var Pstate
 p, i, j: Var processors
 k, l, q: Var nat
 u: Var inputs
 w: Var MBvec
 h: Var MBmatrix
 A: Var set[processors]
 c, d, e: Var cell
 cs: Var cell_state
 K: Var control_state
 H: Var nat (* Majority functions *)
 k_maj: function[MBvec → control_state]
 k_maj_ax: Axiom (∃ A :
 maj_condition(A) ∧ (∀ p : p ∈ A ⊃ w(p).control = K))
 ⊃ k_maj(w) = K
 t_maj: function[MBvec, cell → cell_state]
 t_maj_ax: Axiom (∃ A :
 maj_condition(A) ∧ (∀ p : p ∈ A ⊃ ((w(p)).cellsc) = cs))
 ⊃ t_maj(w, c) = cs
 cell_measure: function[MBvec, nat → nat] == (λ w, k : k)

cell_maj: Recursive function[MBvec, nat \rightarrow cell_array] =
 ($\lambda w, k$: if $k = 0 \vee k > \text{num_cells}$
 then null_cell_array
 else cell_maj($w, k - 1$)
 with [($k - 1$) := t_maj($w, k - 1$)]
 end if) by cell_measure

(* Interpretations for voting-related functions *)

f_s : function[Pstate \rightarrow MB] == (λps : ps)

f_v : function[Pstate, MBvec \rightarrow Pstate] =
 ($\lambda ps, w$: ps with [(control) := k_maj(w), (cells) :=
 cell_maj($w, \text{num_cells}$)]])

rec: function[cell, control_state, nat \rightarrow bool] == ($\lambda c, K, H$: $H > 2$)

dep: function[cell, cell, control_state \rightarrow bool] == ($\lambda c, d, K$: false)

recovery_period_value: Axiom recovery_period = 3

(* Definitions derived from uninterpreted functions *)

dep_agree: function[cell, control_state, Pstate, Pstate \rightarrow bool] =
 ($\lambda c, K, X, Y$: ($\forall d$: $\text{dep}(c, d, K) \supset f_t(X, d) = f_t(Y, d)$))

w_condition: function[set[processors], MBvec, Pstate \rightarrow bool] =
 ($\lambda A, w, ps$: ($\forall p$: $p \in A \supset w(p) = f_s(ps)$))

(* Axioms to be satisfied by the generic application *)

full_recovery: Formula $H \geq \text{recovery_period} \supset \text{rec}(c, K, H)$

initial_recovery: Formula $\text{rec}(c, K, H) \supset H > 2$

dep_recovery: Formula $\text{rec}(c, \text{succ}(K), H + 1) \wedge \text{dep}(c, d, K) \supset \text{rec}(d, K, H)$

control_recovered: Formula

maj_condition(A) \wedge ($\forall p$: $p \in A \supset w(p) = f_s(ps) \supset f_k(f_v(Y, w)) = f_k(ps)$)

cell_recovered: Formula

maj_condition(A)
 \wedge ($\forall p$: $p \in A \supset w(p) = f_s(f_c(u, ps))$)
 $\wedge f_k(X) = K \wedge f_k(ps) = K \wedge \text{dep_agree}(c, K, X, ps)$
 $\supset f_t(f_v(f_c(u, X), w), c) = f_t(f_c(u, ps), c)$

vote_maj: Formula

maj_condition(A) \wedge ($\forall p$: $p \in A \supset w(p) = f_s(ps) \supset f_v(ps, w) = ps$)

(* Support lemmas *)

cell_maj_element: Lemma cell_maj($w, \text{num_cells}$)(c) = t_maj(w, c)

f_v_components: Lemma $f_k(f_v(ps, w)) = k_maj(w) \wedge f_t(f_v(ps, w), c) = t_maj(w, c)$

Proof

p.full.recovery: **Prove full_recovery from recovery_period_value**

p.initial.recovery: **Prove initial_recovery**

p.dep.recovery: **Prove dep_recovery**

p.control.recovered: **Prove control_recovered** { $p \leftarrow p@p1$ } from
k.maj.ax { $K \leftarrow ps.control$ }, f_v { $ps \leftarrow Y, w \leftarrow w$ }

p.cell.recovered: **Prove cell_recovered** { $p \leftarrow p@p1$ } from
t.maj.ax { $cs \leftarrow ((f_c(u, ps)).cellsc)$ },
 f_c { $ps \leftarrow X$ },
 f_c ,
 f_v { $ps \leftarrow f_c(u, X), w \leftarrow w$ },
cell.maj.element

p.vote.maj: **Prove vote_maj** { $p \leftarrow p@p4$ } from
components.equal { $X \leftarrow f_v(ps, w), Y \leftarrow ps$ },
k.maj.ax { $K \leftarrow ps.control$ },
t.maj.ax { $cs \leftarrow ps.cellsc@p1, c \leftarrow c@p1$ },
w.condition,
w.condition { $p \leftarrow p@p2$ },
w.condition { $p \leftarrow p@p3$ },
f.v.components { $c \leftarrow c@p1$ }

cme.prop: function[MBvec, cell \rightarrow function[nat \rightarrow bool]] =
($\lambda w, c : (\lambda q :$
 cell.maj(w, q)(c)
 = if $c < q \wedge q \leq \text{num_cells}$
 then t.maj(w, c)
 else null.cell.array(c)
 end if)

cme.base: **Lemma** cme.prop(w, c)(0)

cme.ind.1: **Lemma** cme.prop(w, c)(q) $\wedge c = q \supset$ cme.prop(w, c)($q + 1$)

cme.ind.2: **Lemma** cme.prop(w, c)(q) $\wedge c \neq q \supset$ cme.prop(w, c)($q + 1$)

p.cme.base: **Prove** cme.base from cme.prop { $q \leftarrow 0$ }, cell.maj { $k \leftarrow 0$ }

p.cme.ind.1: **Prove** cme.ind.1 from
cme.prop { $q \leftarrow q$ }, cme.prop { $q \leftarrow q + 1$ }, cell.maj { $k \leftarrow q + 1$ }

p.cme.ind.2: **Prove** cme.ind.2 from
cme.prop { $q \leftarrow q$ }, cme.prop { $q \leftarrow q + 1$ }, cell.maj { $k \leftarrow q + 1$ }

p.cell.maj.element: **Prove** cell.maj.element from
nat.induction { $p \leftarrow \text{cme.prop}(w, c), n_2 \leftarrow \text{num_cells}$ },
cme.prop { $q \leftarrow \text{num_cells}$ },
cme.base,
cme.ind.1 { $q \leftarrow n_1@p1$ },
cme.ind.2 { $q \leftarrow n_1@p1$ }

p.f.v.components: **Prove** f.v.components from f_v , cell.maj.element

End

cyclic_voting: Module

(* Following is the interpretation for the cyclic voting scheme. *)

Using task_model, nat_inductions

Exporting all with task_model

Theory

```
us, ps, X, Y: Var Pstate
p, i, j: Var processors
k, l, q: Var nat
u: Var inputs
w: Var MBvec
h: Var MBmatrix
A: Var set[processors]
c, d, e: Var cell
cs: Var cell_state
K, L: Var control_state
H: Var nat
C, D: Var cell_array
cell_fn: Type is function[cell → cell_state]
cfn: Var cell_fn (* Majority functions *)
k_maj: function[MBvec → control_state]

k_maj_ax: Axiom (∃ A :
  maj_condition(A) ∧ (∀ p : p ∈ A ⊃ w(p).control = K))
  ⊃ k_maj(w) = K

t_maj: function[MBvec, cell → cell_state]

t_maj_ax: Axiom (∃ A :
  maj_condition(A) ∧ (∀ p : p ∈ A ⊃ ((w(p)).cellsc) = cs))
  ⊃ t_maj(w, c) = cs

cell_measure: function[cell_fn, control_state, cell_array, nat → nat] ==
  (λ cfn, K, C, k : k
    if k = 0 ∨ k > num_cells
    then C
    elseif K = succ(cell_frame(k - 1))
    then cell_apply(cfn, K, C, k - 1)
    with [(k - 1) := cfn(k - 1)]
    else cell_apply(cfn, K, C, k - 1)
    end if)
  by cell_measure
```

(* Interpretations for voting-related functions *)

```
f_s: function[Pstate → MB] =
  (λ ps : ps with [(control) := ps.control, (cells) :=
    cell_apply((λ c : ps.cells(c),
      ps.control,
      null_cell_array,
      num_cells))])
```

```

f_v: function[Pstate, MBvec → Pstate] =
  (λ ps, w : ps with [(control) := k.maj(w),
                      (cells) := cell_apply((λ c : t.maj(w, c)),
                                             ps.control,
                                             ps.cells,
                                             num_cells)])

```

```

rec: function[cell, control_state, nat → bool] =
  (λ c, K, H : H
   > 1 + ( if K = cell_frame(c)
           then schedule_length
           else mod_minus(K, cell_frame(c))
         end if))

```

```

dep: function[cell, cell, control_state → bool] =
  (λ c, d, K : cell_frame(c) ≠ K ∧ c = d)

```

recovery_period_value: **Axiom** recovery_period = schedule_length + 2

(* Definitions derived from uninterpreted functions *)

```

dep_agree: function[cell, control_state, Pstate, Pstate → bool] =
  (λ c, K, X, Y : (∀ d : dep(c, d, K) ⊃ f_i(X, d) = f_i(Y, d)))

```

```

w_condition: function[set[processors], MBvec, Pstate → bool] =
  (λ A, w, ps : (∀ p : p ∈ A ⊃ w(p) = f_s(ps)))

```

(* Axioms to be satisfied by the generic application *)

full_recovery: **Formula** $H \geq \text{recovery_period} \supset \text{rec}(c, K, H)$

initial_recovery: **Formula** $\text{rec}(c, K, H) \supset H > 2$

dep_recovery: **Formula** $\text{rec}(c, \text{succ}(K), H + 1) \wedge \text{dep}(c, d, K) \supset \text{rec}(d, K, H)$

control_recovered: **Formula**

maj_condition(A) $\wedge (\forall p : p \in A \supset w(p) = f_s(ps)) \supset f_k(f_v(Y, w)) = f_k(ps)$

cell_recovered: **Formula**

maj_condition(A)
 $\wedge (\forall p : p \in A \supset w(p) = f_s(f_c(u, ps)))$
 $\wedge f_k(X) = K \wedge f_k(ps) = K \wedge \text{dep_agree}(c, K, X, ps)$
 $\supset f_i(f_v(f_c(u, X), w), c) = f_i(f_c(u, ps), c)$

vote_maj: **Formula**

maj_condition(A) $\wedge (\forall p : p \in A \supset w(p) = f_s(ps)) \supset f_v(ps, w) = ps$

(* Support lemmas *)

cell_apply_element: **Lemma**

cell_apply(cfn, K, C, num_cells)(c)
= if K = succ(cell_frame(c)) then cfn(c) else C(c) end if

f.s.components: Lemma

$$\begin{aligned} K &= \text{ps.control} \supset f_k(f_s(\text{ps})) = K \\ &\wedge f_t(f_s(\text{ps}), c) \\ &= \text{if succ}(\text{cell_frame}(c)) = K \\ &\quad \text{then ps.cells}(c) \\ &\quad \text{else null_cell_array}(c) \\ &\quad \text{end if} \end{aligned}$$

f.v.components: Lemma

$$\begin{aligned} f_k(f_v(\text{ps}, w)) &= \text{k.maj}(w) \\ &\wedge f_t(f_v(\text{ps}, w), c) \\ &= \text{if succ}(\text{cell_frame}(c)) = \text{ps.control} \\ &\quad \text{then t.maj}(w, c) \\ &\quad \text{else ps.cells}(c) \\ &\quad \text{end if} \end{aligned}$$

f.c.uncomputed_cells: Lemma

$$\text{cell_frame}(c) \neq X.\text{control} \supset f_c(u, X).\text{cells}(c) = X.\text{cells}(c)$$

Proof

p.full.recovery: Prove full.recovery from

rec,
recovery_period_value,
control_state_invariant
{control_state_var \leftarrow mod_minus(K , cell_frame($c@p1$))}

p.initial.recovery: Prove initial.recovery from

rec,
schedule_length_ax,
mod_minus_zero { $L \leftarrow$ cell_frame($c@p1$)},
nat_invariant {nat_var \leftarrow mod_minus(K , cell_frame($c@p1$))}

p.dep.recovery: Prove dep.recovery from

rec { $K \leftarrow$ succ(K), $H \leftarrow H + 1$ },
dep,
rec { $c \leftarrow d$ },
control_state_invariant {control_state_var \leftarrow mod_minus(K , cell_frame(c))},
mod_minus_plus { $L \leftarrow$ cell_frame(c)}

p.control.recovered: Prove control.recovered { $p \leftarrow p@p1$ } from

k_maj_ax { $K \leftarrow$ ps.control}, f_v { $\text{ps} \leftarrow Y$, $w \leftarrow w$ }, f_s

p.cell.recovered: Prove cell.recovered { $p \leftarrow p@p1$ } from

t_maj_ax { $\text{cs} \leftarrow ((f_s(f_c(u, \text{ps}))).\text{cellsc})$ },
dep_agree { $Y \leftarrow \text{ps}$, $d \leftarrow c$ },
dep { $d \leftarrow c$ },
f_s_components { $\text{ps} \leftarrow f_c(u, \text{ps})$, $K \leftarrow (f_c(u, X)).\text{control}$ },
f_c_uncomputed_cells { $X \leftarrow \text{ps}$ },
f_c_uncomputed_cells,
 f_c { $\text{ps} \leftarrow X$ },
 f_c ,
f_v_components { $\text{ps} \leftarrow f_c(u, X)$ }

p_vote_maj: Prove `vote_maj {p ← p@p4}` from
`components_equal {X ← fv(ps, w), Y ← ps}`,
`k_maj_ax {K ← ps.control}`,
`t_maj_ax {cs ← ps.cellsc@p1}, c ← c@p1`,
`w_condition`,
`w_condition {p ← p@p2}`,
`w_condition {p ← p@p3}`,
`fs`,
`cell_apply_element`
`{cfn ← (λ c : ps.cells(c)),`
`c ← c@p1,`
`K ← ps.control,`
`C ← null_cell_array}`,
`fv_components {c ← c@p1}`

cae_prop: function[`cell_fn, control_state, cell_array, cell`
`→ function[nat → bool]] =`
`(λ cfn, K, C, c :`
`(λ q : cell_apply(cfn, K, C, q)(c)`
`= if c < q ∧ q ≤ num_cells ∧ K = succ(cell.frame(c))`
`then cfn(c)`
`else C(c)`
`end if))`

cae_base: Lemma `cae_prop(cfn, K, C, c)(0)`

cae_ind.1: Lemma `cae_prop(cfn, K, C, c)(q) ∧ c = q`
`⊃ cae_prop(cfn, K, C, c)(q + 1)`

cae_ind.2: Lemma `cae_prop(cfn, K, C, c)(q) ∧ c ≠ q`
`⊃ cae_prop(cfn, K, C, c)(q + 1)`

p_cae_base: Prove `cae_base` from `cae_prop {q ← 0}`, `cell_apply {k ← 0}`

p_cae_ind.1: Prove `cae_ind.1` from
`cae_prop {q ← q}`, `cae_prop {q ← q + 1}`, `cell_apply {k ← q + 1}`

p_cae_ind.2: Prove `cae_ind.2` from
`cae_prop {q ← q}`, `cae_prop {q ← q + 1}`, `cell_apply {k ← q + 1}`

p_cell_apply_element: Prove `cell_apply_element` from
`nat_induction {p ← cae_prop(cfn, K, C, c), n2 ← num_cells}`,
`cae_prop {q ← num_cells}`,
`cae_base`,
`cae_ind.1 {q ← n1@p1}`,
`cae_ind.2 {q ← n1@p1}`

p_f_s_components: Prove `fs_components` from
`fs`,
`cell_apply_element`
`{cfn ← (λ c : ps.cells(c)),`
`K ← ps.control,`
`C ← null_cell_array}`

p_f_v_components: Prove `fv_components` from
`fv`,
`cell_apply_element`
`{cfn ← (λ c : t_maj(w, c)),`
`K ← ps.control,`
`C ← ps.cells}`

p.f.c.uncomputed.cells: Prove f.c.uncomputed.cells from
f_c {ps ← X}, exec_element {C ← X.cells, K ← X.control}

End

Appendix C

Results of Proof Chain Analysis

The following pages were obtained from Ehdm using the proof-chain analyzer command (M-x apcs) applied to the module top.

Terse proof chains for module top

Use of the formula

RS_to_US.frame_commutates
requires the following TCCs to be proven
RS_to_US_tcc.reachable_in_n_TCC1
RS_to_US_tcc.reachable_in_n_TCC2

Formula RS_to_US_tcc.reachable_in_n_TCC2 is a termination TCC for
DA_to_DS.reachable_in_n

Proof of

RS_to_US_tcc.reachable_in_n_TCC2
must not use
DA_to_DS.reachable_in_n

Use of the formula

rcp_defs.recovery_period_ax
requires the following TCCs to be proven
rcp_defs_tcc.processors_TCC1
rcp_defs_tcc.proc_plus_TCC1
rcp_defs_tcc.processors_induction_TCC1
rcp_defs_tcc.processors_induction_TCC2
rcp_defs_tcc.p_reachability_TCC1

Use of the formula

cardinality[rcp_defs.processors].card_empty
requires the following assumptions to be discharged
cardinality[rcp_defs.processors].finite

Use of the formula

DS_to_RS.frame_commutates
requires the following TCCs to be proven
DS_to_RS_tcc.ss_update_TCC1
DS_to_RS_tcc.ss_update_TCC2
DS_to_RS_tcc.ss_update_TCC3
DS_to_RS_tcc.MBmatrix_cons_TCC1

Formula DS_to_RS_tcc.ss_update_TCC3 is a termination TCC
for DA_to_DS.ss_update

Proof of

DS_to_RS_tcc.ss_update_TCC3
must not use
DA_to_DS.ss_update

Formula DS_to_RS_tcc.MBmatrix_cons_TCC1 is a termination TCC for
DA_to_DS.MBmatrix_cons

Proof of

DS_to_RS_tcc.MBmatrix_cons_TCC1
must not use
DA_to_DS.MBmatrix_cons

Use of the formula
noetherian[rcp_defs.proc_plus, rcp_defs.less].general_induction
requires the following assumptions to be discharged
noetherian[rcp_defs.proc_plus, rcp_defs.less].well_founded

Use of the formula
DS_support_proof.sl13_prop
requires the following TCCs to be proven
DS_support_proof_tcc.p_sl13_base_TCC1
DS_support_proof_tcc.p_sl13_ind_TCC1
DS_support_proof_tcc.p_support_13_TCC1

Use of the formula
DS_map_proof.ml1_prop
requires the following TCCs to be proven
DS_map_proof_tcc.p_ml1_base_TCC1
DS_map_proof_tcc.p_ml1_ind_TCC1
DS_map_proof_tcc.p_map_1_TCC1

Use of the formula
DA_to_DS.phase_commutates
requires the following TCCs to be proven
DA_to_DS_tcc.ss_update_TCC1
DA_to_DS_tcc.ss_update_TCC2
DA_to_DS_tcc.ss_update_TCC3
DA_to_DS_tcc.MBmatrix_cons_TCC1
DA_to_DS_tcc.reachable_in_n_TCC1
DA_to_DS_tcc.reachable_in_n_TCC2

Formula DA_to_DS_tcc.ss_update_TCC3 is a termination TCC
for DA_to_DS.ss_update
Proof of
DA_to_DS_tcc.ss_update_TCC3
must not use
DA_to_DS.ss_update

Formula DA_to_DS_tcc.MBmatrix_cons_TCC1 is a termination TCC for
DA_to_DS.MBmatrix_cons
Proof of
DA_to_DS_tcc.MBmatrix_cons_TCC1
must not use
DA_to_DS.MBmatrix_cons

Formula DA_to_DS_tcc.reachable_in_n_TCC2 is a termination TCC for
DA_to_DS.reachable_in_n
Proof of
DA_to_DS_tcc.reachable_in_n_TCC2
must not use
DA_to_DS.reachable_in_n

Use of the formula
DA_map_proof.ml1_prop
requires the following TCCs to be proven
DA_map_proof_tcc.p_ml1_base_TCC1
DA_map_proof_tcc.p_ml1_ind_TCC1
DA_map_proof_tcc.p_map_1_TCC1

Use of the formula

DA_broadcast_prf.rtp4
requires the following TCCs to be proven
DA_broadcast_prf_tcc.p_br8_TCC1

Use of the formula

clkmod.in_R_interval
requires the following TCCs to be proven
clkmod_tcc.half_TCC1
clkmod_tcc.Rho_TCC1
clkmod_tcc.Corr_TCC1
clkmod_tcc.num_good_clocks_TCC1
clkmod_tcc.num_good_clocks_TCC2
clkmod_tcc.num_good_clocks_TCC3
clkmod_tcc.num_good_clocks_TCC4
clkmod_tcc.num_good_clocks_TCC5
clkmod_tcc.C6_TCC1

Formula clkmod_tcc.Corr_TCC1 is a termination TCC for clkmod.Corr
Proof of

clkmod_tcc.Corr_TCC1
must not use
clkmod.Corr

Formula clkmod_tcc.num_good_clocks_TCC4 is a termination TCC for
clkmod.num_good_clocks

Proof of

clkmod_tcc.num_good_clocks_TCC4
must not use
clkmod.num_good_clocks

Formula clkmod_tcc.num_good_clocks_TCC5 is a termination TCC for
clkmod.num_good_clocks

Proof of

clkmod_tcc.num_good_clocks_TCC5
must not use
clkmod.num_good_clocks

Use of the formula

DA_invariants.state_invariant
requires the following TCCs to be proven
DA_invariants_tcc.Corr_lem_TCC1

Use of the formula

DA_support_proof.sl15_prop
requires the following TCCs to be proven
DA_support_proof_tcc.p_sl13_base_TCC1
DA_support_proof_tcc.p_sl13_ind_TCC1
DA_support_proof_tcc.p_support_13_TCC1
DA_support_proof_tcc.p_sl15_ind_TCC1

SUMMARY

The proof chain is complete

The axioms and assumptions at the base are:

DA.all_durations
DA.broadcast_duration

DA.broadcast_duration2
 DA.pos_durations
 RS.majority.maj_ax
 cardinality[EXPR].card_empty
 clkmod.C0
 clkmod.C1
 clkmod.C2
 clkmod.Theorem_1
 clkmod.Theorem_2
 clkmod.adj_always_pos
 generic_FT.cell_recovered
 generic_FT.components_equal
 generic_FT.control_recovered
 generic_FT.dep_recovery
 generic_FT.full_recovery
 generic_FT.initial_recovery
 generic_FT.succ_ax
 generic_FT.vote_maj
 multiplication.mult_non_neg
 nat_inductions.nat_induction
 noetherian[EXPR, EXPR].general_induction
 rcp_defs.processors_exist_ax
 rcp_defs.recovery_period_ax
 sets[EXPR].extensionality
 Total: 26

The definitions and type-constraints are:

DA.N_da
 DA.N_da_broadcast
 DA.N_da_compute
 DA.N_da_sync
 DA.N_da_vote
 DA.broadcast_received
 DA.clock_advanced
 DA.da_rt
 DA.enough_hardware
 DA.initial_da
 DA.maj_working
 DA.working_proc
 DA.working_set
 DA_invariants.cum_delta_val
 DA_invariants.lclock_eq
 DA_invariants.lclock_val
 DA_invariants.nf_clks
 DA_invariants.state_invariant
 DA_invariants.state_invariant_to_n
 DA_lemmas.hide1
 DA_map_proof.m11_prop
 DA_map_proof.m12_prop
 DA_map_proof.m14_prop
 DA_support_proof.s15_prop
 DA_to_DS.DAmap
 DA_to_DS.reachable
 DA_to_DS.reachable_in_n
 DA_to_DS.ss_update
 DS.N_ds
 DS.N_ds_broadcast
 DS.N_ds_compute

DS.N_ds_sync
 DS.N_ds_vote
 DS.allowable_faults
 DS.broadcast_received
 DS.frame_N_ds
 DS.initial_ds
 DS.maj_working
 DS.working_proc
 DS.working_set
 DS_lemmas.half_frame_N_ds
 DS_lemmas.quarter_frame_N_ds
 DS_map_proof.m11_prop
 DS_map_proof.m12_prop
 DS_support_proof.s113_prop
 DS_to_RS.DSmap
 DS_to_RS.MBmatrix_cons
 DS_to_RS.good_values_sent
 DS_to_RS.is_new_proc_state
 DS_to_RS.ss_update
 DS_to_RS.voted_final_state
 RS.N_rs
 RS.allowable_faults
 RS.good_values_sent
 RS.initial_rs
 RS.maj_working
 RS.voted_final_state
 RS.working_proc
 RS.working_set
 RS_invariants.state_invariant
 RS_invariants.state_invariant_to_n
 RS_lemmas.cell_recovery
 RS_lemmas.control_recovery
 RS_lemmas.state_recovery
 RS_lemmas.working_majority
 RS_majority.maj_exists
 RS_to_US.RSmap
 RS_to_US.reachable
 RS_to_US.reachable_in_n
 US.N_us
 US.initial_us
 absmod.abs
 clkmod.Corr
 clkmod.S1
 clkmod.S1C
 clkmod.S2
 clkmod.T_sup
 clkmod.enough_clocks
 clkmod.goodclock
 clkmod.in_R_interval
 clkmod.nonfaulty_clock
 clkmod.num_good_clocks
 clkmod.rt
 generic_FT.dep_agree
 generic_FT.maj_condition
 multiplication.mult
 rcp_defs.distinct_phases
 rcp_defs.member_phases
 rcp_defs.next_phase


```
rcp_defs.prev_phase
sets[EXPR].empty
Total: 91
```

The formulae used are:

```
DA_broadcast_prf.br1
DA_broadcast_prf.br1a
DA_broadcast_prf.br2
DA_broadcast_prf.br3
DA_broadcast_prf.br3_aa
DA_broadcast_prf.br4
DA_broadcast_prf.br5
DA_broadcast_prf.br6
DA_broadcast_prf.br7
DA_broadcast_prf.br8
DA_broadcast_prf.br9
DA_broadcast_prf.int5
DA_broadcast_prf.rtp0
DA_broadcast_prf.rtp0a
DA_broadcast_prf.rtp1
DA_broadcast_prf.rtp2
DA_broadcast_prf.rtp3
DA_broadcast_prf.rtp4
DA_broadcast_prf.rtp4a
DA_broadcast_prf.rtp4b
DA_broadcast_prf.rtp5
DA_broadcast_prf.rtp6
DA_broadcast_prf.rtp7
DA_broadcast_prf.tcc.p_br8_TCC1
DA_intervals.br_int
DA_intervals.int0
DA_intervals.int1
DA_intervals.int1a
DA_intervals.int2
DA_intervals.int2a
DA_intervals.int3
DA_intervals.int4
DA_invariants.base_state_ind
DA_invariants.cdi_l2a
DA_invariants.clkval_inv
DA_invariants.clkval_inv_l1
DA_invariants.clkval_inv_l2
DA_invariants.cum_delta_inv
DA_invariants.cum_delta_inv_l1
DA_invariants.cum_delta_inv_l2
DA_invariants.cum_delta_inv_l4
DA_invariants.da_rt_lem
DA_invariants.enough_inv
DA_invariants.enough_inv_l1
DA_invariants.enough_inv_l2
DA_invariants.ind_state_ind
DA_invariants.lclock_inv
DA_invariants.lclock_inv_l1
DA_invariants.lclock_inv_l2
DA_invariants.lclock_inv_l2b
DA_invariants.lclock_inv_l2c
DA_invariants.lclock_inv_l3
DA_invariants.lclock_inv_l4
```

DA_invariants.nfclk_inv
 DA_invariants.nfclk_inv_l1
 DA_invariants.nfclk_inv_l2
 DA_invariants.rtl1
 DA_invariants.state_induction
 DA_invariants_tcc.Corr_lem_TCC1
 DA_lemmas.ELT
 DA_lemmas.com_broadcast_1
 DA_lemmas.com_broadcast_2
 DA_lemmas.com_broadcast_3
 DA_lemmas.com_broadcast_4
 DA_lemmas.com_broadcast_5
 DA_lemmas.com_sync_1
 DA_lemmas.com_sync_2
 DA_lemmas.com_sync_3
 DA_lemmas.com_sync_4
 DA_lemmas.com_vote_1
 DA_lemmas.com_vote_2
 DA_lemmas.com_vote_3
 DA_lemmas.com_vote_4
 DA_lemmas.earliest_later_time
 DA_lemmas.elt_a
 DA_lemmas.map_1
 DA_lemmas.map_2
 DA_lemmas.map_3
 DA_lemmas.map_4
 DA_lemmas.map_7
 DA_lemmas.phase_com_broadcast
 DA_lemmas.phase_com_compute
 DA_lemmas.phase_com_lx1
 DA_lemmas.phase_com_lx2
 DA_lemmas.phase_com_lx4
 DA_lemmas.phase_com_lx7
 DA_lemmas.phase_com_sync
 DA_lemmas.phase_com_vote
 DA_lemmas.support_1
 DA_lemmas.support_14
 DA_lemmas.support_15
 DA_lemmas.support_16
 DA_map_proof.m11_base
 DA_map_proof.m11_ind
 DA_map_proof.m12_base
 DA_map_proof.m12_ind
 DA_map_proof.m14_base
 DA_map_proof.m14_ind
 DA_map_proof_tcc.p_map_1_TCC1
 DA_map_proof_tcc.p_m11_base_TCC1
 DA_map_proof_tcc.p_m11_ind_TCC1
 DA_support_proof.s115_base
 DA_support_proof.s115_ind
 DA_support_proof_tcc.p_s113_base_TCC1
 DA_support_proof_tcc.p_s113_ind_TCC1
 DA_support_proof_tcc.p_s115_ind_TCC1
 DA_support_proof_tcc.p_support_13_TCC1
 DA_to_DS.initial_maps
 DA_to_DS.phase_commutes
 DA_to_DS_tcc.MBmatrix_cons_TCC1
 DA_to_DS_tcc.reachable_in_n_TCC1

DA_to_DS_tcc.reachable_in_n_TCC2
DA_to_DS_tcc.ss_update_TCC1
DA_to_DS_tcc.ss_update_TCC2
DA_to_DS_tcc.ss_update_TCC3
DS_lemmas.fc_A_1a
DS_lemmas.fc_A_1b
DS_lemmas.fc_A_1c
DS_lemmas.fc_A_1d
DS_lemmas.fc_A_1e
DS_lemmas.fc_A_1f
DS_lemmas.fc_A_2a
DS_lemmas.fc_A_2b
DS_lemmas.fc_A_2c
DS_lemmas.fc_A_2d
DS_lemmas.fc_A_3a
DS_lemmas.fc_A_3b
DS_lemmas.fc_A_3c
DS_lemmas.fc_A_3d
DS_lemmas.map_1
DS_lemmas.map_2
DS_lemmas.map_3
DS_lemmas.map_4
DS_lemmas.map_5
DS_lemmas.map_7
DS_lemmas.support_1
DS_lemmas.support_10
DS_lemmas.support_11
DS_lemmas.support_13
DS_lemmas.support_4
DS_lemmas.support_5
DS_lemmas.support_6
DS_lemmas.support_7
DS_lemmas.support_8
DS_lemmas.support_9
DS_map_proof.m11_base
DS_map_proof.m11_ind
DS_map_proof.m12_base
DS_map_proof.m12_ind
DS_map_proof_tcc.p_map_1_TCC1
DS_map_proof_tcc.p_m11_base_TCC1
DS_map_proof_tcc.p_m11_ind_TCC1
DS_support_proof.s113_base
DS_support_proof.s113_ind
DS_support_proof_tcc.p_s113_base_TCC1
DS_support_proof_tcc.p_s113_ind_TCC1
DS_support_proof_tcc.p_support_13_TCC1
DS_to_RS.fc_A
DS_to_RS.fc_B
DS_to_RS.fr_com_1
DS_to_RS.fr_com_2
DS_to_RS.frame_commutates
DS_to_RS.initial_maps
DS_to_RS_tcc.MBmatrix_cons_TCC1
DS_to_RS_tcc.ss_update_TCC1
DS_to_RS_tcc.ss_update_TCC2
DS_to_RS_tcc.ss_update_TCC3
RS_invariants.base_state_ind
RS_invariants.ind_state_ind

RS_invariants.maj_working_inv
RS_invariants.maj_working_inv_11
RS_invariants.maj_working_inv_12
RS_invariants.state_induction
RS_invariants.state_rec_inv
RS_invariants.state_rec_inv_11
RS_invariants.state_rec_inv_12
RS_invariants.state_rec_inv_13
RS_invariants.state_rec_inv_14
RS_invariants.state_rec_inv_15
RS_lemmas.consensus_prop
RS_lemmas.initial_maj
RS_lemmas.initial_maj_cond
RS_lemmas.initial_working
RS_lemmas.maj_sent
RS_lemmas.rec_maj_exists
RS_lemmas.rec_maj_f_c
RS_lemmas.working_set_healthy
RS_to_US.frame_commutates
RS_to_US.initial_maps
RS_to_US_tcc.reachable_in_n_TCC1
RS_to_US_tcc.reachable_in_n_TCC2
absmod.abs_ge0
absmod.abs_leq
absmod.abs_main
cardinality[rcp_defs.processors].finite
clkmod.sync_thm
clkmod_tcc.C6_TCC1
clkmod_tcc.Corr_TCC1
clkmod_tcc.Rho_TCC1
clkmod_tcc.half_TCC1
clkmod_tcc.num_good_clocks_TCC1
clkmod_tcc.num_good_clocks_TCC2
clkmod_tcc.num_good_clocks_TCC3
clkmod_tcc.num_good_clocks_TCC4
clkmod_tcc.num_good_clocks_TCC5
clkprop.GOAL
clkprop.ft10
clkprop.ft11
clkprop.ft12
clkprop.ft2
clkprop.ft3
clkprop.ft4
clkprop.ft5
clkprop.ft6
clkprop.ft7
clkprop.ft8
clkprop.ft8a
clkprop.ft9
clkprop.nfc_a
clkprop.nfc_lem
generic_FT.card_fullset
generic_FT.nat_nit
generic_FT.proc_extensionality
multiplication.distrib_minus
multiplication.mult_com
multiplication.mult_ldistrib_minus
multiplication.mult_leq

```
noetherian[rcp_defs.proc_plus, rcp_defs.lessp].well_founded
rcp_defs.processors_induction
rcp_defs.reachability
rcp_defs.tcc.p_reachability_TCC1
rcp_defs.tcc.proc_plus_TCC1
rcp_defs.tcc.processors_TCC1
rcp_defs.tcc.processors_induction_TCC1
rcp_defs.tcc.processors_induction_TCC2
Total: 235
```

The completed proofs are:

```
DA_broadcast_prf.p_br1
DA_broadcast_prf.p_bria
DA_broadcast_prf.p_br2
DA_broadcast_prf.p_br3
DA_broadcast_prf.p_br3_aa
DA_broadcast_prf.p_br4
DA_broadcast_prf.p_br5
DA_broadcast_prf.p_br6
DA_broadcast_prf.p_br7
DA_broadcast_prf.p_br8
DA_broadcast_prf.p_br9
DA_broadcast_prf.p_com_broadcast_5
DA_broadcast_prf.p_rtp0
DA_broadcast_prf.p_rtp0a
DA_broadcast_prf.p_rtp1
DA_broadcast_prf.p_rtp2
DA_broadcast_prf.p_rtp3
DA_broadcast_prf.p_rtp4
DA_broadcast_prf.p_rtp4a
DA_broadcast_prf.p_rtp4b
DA_broadcast_prf.p_rtp5
DA_broadcast_prf.p_rtp6
DA_broadcast_prf.p_rtp7
DA_broadcast_prf.tcc.p_br8_TCC1_PROOF
DA_intervals.p_br_int
DA_intervals.p_int0
DA_intervals.p_int1
DA_intervals.p_int1a
DA_intervals.p_int2
DA_intervals.p_int2a
DA_intervals.p_int3
DA_intervals.p_int4
DA_intervals.p_int5
DA_invariants.p_base_state_ind
DA_invariants.p_cdi_l2a
DA_invariants.p_clkval_inv
DA_invariants.p_clkval_inv_l1
DA_invariants.p_clkval_inv_l2
DA_invariants.p_cum_delta_inv
DA_invariants.p_cum_delta_inv_l1
DA_invariants.p_cum_delta_inv_l2
DA_invariants.p_cum_delta_inv_l4
DA_invariants.p_da_rt_lem
DA_invariants.p_enough_inv
DA_invariants.p_enough_inv_l1
DA_invariants.p_enough_inv_l2
DA_invariants.p_ind_state_ind
```

DA_invariants.p_lclock_inv
 DA_invariants.p_lclock_inv_l1
 DA_invariants.p_lclock_inv_l2
 DA_invariants.p_lclock_inv_l2b
 DA_invariants.p_lclock_inv_l2c
 DA_invariants.p_lclock_inv_l3
 DA_invariants.p_lclock_inv_l4
 DA_invariants.p_nfclk_inv
 DA_invariants.p_nfclk_inv_l1
 DA_invariants.p_nfclk_inv_l2
 DA_invariants.p_rtl1
 DA_invariants.p_state_induction
 DA_invariants_tcc.Corr_lem_TCC1_PROOF
 DA_map_proof.p_map_1
 DA_map_proof.p_map_2
 DA_map_proof.p_map_3
 DA_map_proof.p_map_4
 DA_map_proof.p_map_7
 DA_map_proof.p_m1_base
 DA_map_proof.p_m1_ind
 DA_map_proof.p_m2_base
 DA_map_proof.p_m2_ind
 DA_map_proof.p_m4_base
 DA_map_proof.p_m4_ind
 DA_map_proof_tcc.p_map_1_TCC1_PROOF
 DA_map_proof_tcc.p_m1_base_TCC1_PROOF
 DA_map_proof_tcc.p_m1_ind_TCC1_PROOF
 DA_support_proof.p_sl15_base
 DA_support_proof.p_sl15_ind
 DA_support_proof.p_support_1
 DA_support_proof.p_support_14
 DA_support_proof.p_support_15
 DA_support_proof.p_support_16
 DA_support_proof_tcc.p_sl13_base_TCC1_PROOF
 DA_support_proof_tcc.p_sl13_ind_TCC1_PROOF
 DA_support_proof_tcc.p_support_13_TCC1_PROOF
 DA_tcc_proof.C6_TCC1_PROOF
 DA_tcc_proof.Rho_TCC1_PROOF
 DA_tcc_proof.p_sl15_ind_TCC1_PROOF
 DA_to_DS_tcc.MBmatrix_cons_TCC1_PROOF
 DA_to_DS_tcc.reachable_in_n_TCC1_PROOF
 DA_to_DS_tcc.reachable_in_n_TCC2_PROOF
 DA_to_DS_tcc.ss_update_TCC1_PROOF
 DA_to_DS_tcc.ss_update_TCC2_PROOF
 DA_to_DS_tcc.ss_update_TCC3_PROOF
 DA_top_proof.p_ELT
 DA_top_proof.p_com_broadcast_1
 DA_top_proof.p_com_broadcast_2
 DA_top_proof.p_com_broadcast_3
 DA_top_proof.p_com_broadcast_4
 DA_top_proof.p_com_sync_1
 DA_top_proof.p_com_sync_2
 DA_top_proof.p_com_sync_3
 DA_top_proof.p_com_sync_4
 DA_top_proof.p_com_vote_1
 DA_top_proof.p_com_vote_2
 DA_top_proof.p_com_vote_3
 DA_top_proof.p_com_vote_4

DA_top_proof.p_earliest_later_time
DA_top_proof.p_elt_a
DA_top_proof.p_initial_maps
DA_top_proof.p_phase_com_broadcast
DA_top_proof.p_phase_com_compute
DA_top_proof.p_phase_com_lx1
DA_top_proof.p_phase_com_lx2
DA_top_proof.p_phase_com_lx4
DA_top_proof.p_phase_com_lx7
DA_top_proof.p_phase_com_sync
DA_top_proof.p_phase_com_vote
DA_top_proof.p_phase_commutates
DS_map_proof.p_map_1
DS_map_proof.p_map_2
DS_map_proof.p_map_3
DS_map_proof.p_map_4
DS_map_proof.p_map_5
DS_map_proof.p_map_7
DS_map_proof.p_ml1_base
DS_map_proof.p_ml1_ind
DS_map_proof.p_ml2_base
DS_map_proof.p_ml2_ind
DS_map_proof_tcc.p_map_1_TCC1_PROOF
DS_map_proof_tcc.p_ml1_base_TCC1_PROOF
DS_map_proof_tcc.p_ml1_ind_TCC1_PROOF
DS_support_proof.p_sl13_base
DS_support_proof.p_sl13_ind
DS_support_proof.p_support_1
DS_support_proof.p_support_10
DS_support_proof.p_support_11
DS_support_proof.p_support_13
DS_support_proof.p_support_4
DS_support_proof.p_support_5
DS_support_proof.p_support_6
DS_support_proof.p_support_7
DS_support_proof.p_support_8
DS_support_proof.p_support_9
DS_support_proof_tcc.p_sl13_base_TCC1_PROOF
DS_support_proof_tcc.p_sl13_ind_TCC1_PROOF
DS_support_proof_tcc.p_support_13_TCC1_PROOF
DS_to_RS_tcc.MBmatrix_cons_TCC1_PROOF
DS_to_RS_tcc.ss_update_TCC1_PROOF
DS_to_RS_tcc.ss_update_TCC2_PROOF
DS_to_RS_tcc.ss_update_TCC3_PROOF
DS_top_proof.p_fc_A
DS_top_proof.p_fc_A_1a
DS_top_proof.p_fc_A_1b
DS_top_proof.p_fc_A_1c
DS_top_proof.p_fc_A_1d
DS_top_proof.p_fc_A_1e
DS_top_proof.p_fc_A_1f
DS_top_proof.p_fc_A_2a
DS_top_proof.p_fc_A_2b
DS_top_proof.p_fc_A_2c
DS_top_proof.p_fc_A_2d
DS_top_proof.p_fc_A_3a
DS_top_proof.p_fc_A_3b
DS_top_proof.p_fc_A_3c

DS_top_proof.p_fc_A_3d
 DS_top_proof.p_fc_B
 DS_top_proof.p_fr_com_1
 DS_top_proof.p_fr_com_2
 DS_top_proof.p_frame_commutates
 DS_top_proof.p_initial_maps
 RS_invariants.p_base_state_ind
 RS_invariants.p_ind_state_ind
 RS_invariants.p_maj_working_inv
 RS_invariants.p_maj_working_inv_l1
 RS_invariants.p_maj_working_inv_l2
 RS_invariants.p_state_induction
 RS_invariants.p_state_rec_inv
 RS_invariants.p_state_rec_inv_l1
 RS_invariants.p_state_rec_inv_l2
 RS_invariants.p_state_rec_inv_l3
 RS_invariants.p_state_rec_inv_l4
 RS_invariants.p_state_rec_inv_l5
 RS_tcc_proof.proc_plus_TCC1_PROOF
 RS_tcc_proof.processors_TCC1_PROOF
 RS_to_US_tcc.reachable_in_n_TCC1_PROOF
 RS_to_US_tcc.reachable_in_n_TCC2_PROOF
 RS_top_proof.p_consensus_prop
 RS_top_proof.p_frame_commutates
 RS_top_proof.p_initial_maj
 RS_top_proof.p_initial_maj_cond
 RS_top_proof.p_initial_maps
 RS_top_proof.p_initial_working
 RS_top_proof.p_maj_sent
 RS_top_proof.p_rec_maj_exists
 RS_top_proof.p_rec_maj_f_c
 RS_top_proof.p_working_set_healthy
 absmod.abs_ge0_proof
 absmod.abs_leq_proof
 absmod.abs_main_proof
 clkmod.p_sync_thm
 clkmod_tcc.Corr_TCC1_PROOF
 clkmod_tcc.half_TCC1_PROOF
 clkmod_tcc.num_good_clocks_TCC1_PROOF
 clkmod_tcc.num_good_clocks_TCC2_PROOF
 clkmod_tcc.num_good_clocks_TCC3_PROOF
 clkmod_tcc.num_good_clocks_TCC4_PROOF
 clkmod_tcc.num_good_clocks_TCC5_PROOF
 clkprop.p_GOAL
 clkprop.p_ft10
 clkprop.p_ft11
 clkprop.p_ft12
 clkprop.p_ft2
 clkprop.p_ft3
 clkprop.p_ft4
 clkprop.p_ft5
 clkprop.p_ft6
 clkprop.p_ft7
 clkprop.p_ft8
 clkprop.p_ft8a
 clkprop.p_ft9
 clkprop.p_nfc_a
 clkprop.p_nfc_lem

generic_FT.disharge_finite
generic_FT.p_card_fullset
generic_FT.p_nat_nit
generic_FT.p_proc_extensionality
multiplication.distrib_minus_pr
multiplication.mult_com_pr
multiplication.mult_ldistrib_minus_proof
multiplication.mult_leq_pr
rcp_defs.p_processors_induction
rcp_defs.p_reachability
rcp_defs.p_well_founded
rcp_defs_tcc.p_reachability_TCC1_PROOF
rcp_defs_tcc.processors_induction_TCC1_PROOF
rcp_defs_tcc.processors_induction_TCC2_PROOF
top.p_DA_initial_maps
top.p_DA_phase_commutates
top.p_DS_frame_commutates
top.p_DS_initial_maps
top.p_RS_frame_commutates
top.p_RS_initial_maps
Total: 241

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE January 1992	3. REPORT TYPE AND DATES COVERED Technical Memorandum	
4. TITLE AND SUBTITLE Formal Design and Verification of a Reliable Computing Platform for Real-Time Control, Phase 2 Results		5. FUNDING NUMBERS WU 505-64-10-05	
6. AUTHOR(S) Ricky W. Butler and Benedetto L. Di Vito*			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NASA Langley Research Center Hampton, VA 23665-52255		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001		10. SPONSORING / MONITORING AGENCY REPORT NUMBER NASA TM-104196	
11. SUPPLEMENTARY NOTES *VIGYAN, Inc., Hampton, VA			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 62		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) In this paper the design and formal verification of the Reliable Computing Platform (RCP), a fault-tolerant computing system for digital flight control applications, is presented. The RCP utilizes N-Multiply Redundant (NMR) style redundancy to mask faults and internal majority voting to flush the effects of transient faults. The system is formally specified and verified using the Ehdn verification system. A major goal of this work is to provide the system with significant capability to withstand the effects of High Intensity Radiated Fields (HIRF).			
14. SUBJECT TERMS Fault tolerance, formal methods, ultrareliability, transient faults, HIRF, software verification		15. NUMBER OF PAGES 160	16. PRICE CODE A08
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT