Introduction to the Theory of Nested Transactions

Nancy Lynch
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA. 02139

Michael Merritt
A. T. and T. Bell Laboratories
Murray Hill, NJ 07974-2070

July 7, 1986

ABSTRACT

A new formal model is presented, for studying concurrency and resiliency properties for nested transactions. The model is used to state and prove correctness of a well-known locking algorithm.

Keywords: nested transactions, transactions, concurrency control, resiliency, databases, serializability, orphans

©1986 Massachusetts Institute of Technology, Cambridge, MA. 02139

This work was supported in part by the Office of Naval Research under Contract N00014-85-K-0168, by the Office of Army Research under Contract DAAG29-84-K-0058, by the National Science Foundation under Grant DCR-83-02391, and by the Defense Advanced Research Projects Agency (DARPA) under Grant N00014-83-K-0125.

Introduction to the Theory of Nested Transactions

Nancy Lynch
Massachusetts Institute of Technology
Cambridge, Mass.

Michael Merritt
A. T. and T. Bell Laboratories
Murray Hill, New Jersey

Abstract: A new formal model is presented, for studying concurrency and resiliency properties for nested transactions. The model is used to state and prove correctness of a well-known locking algorithm.

1. Introduction

This paper develops the foundation for a general theory of nested transactions. We present a simple formal model for studying concurrency and resiliency in a nested environment. This model has distinct advantages over the many alternatives, the greatest of which is the unification of a subject replete with formalisms, correctness conditions and proof techniques. The authors are presently engaged in an ambitious project to recast the substantial amount of work in nested transactions within this single intuitive framework. These pages contain the preliminary results of that project – a description of the model, and its use in stating and proving correctness conditions for two variations of a well-known algorithm.

The model is based on *I/O automata*, a simple formalization of communicating automata. It is not complex – it is easily presented in a few pages, and easy to understand, given a minimal background in automata theory. Each nested transaction and data object is modelled by a separate I/O automaton. These automata, the system *primitives*, issue requests to and receive replies from some *scheduler*, which is simply another I/O automaton. Simple syntactic constraints on the interactions of these automata ensure, for example, that no transaction requests the creation of the same child more than once. One scheduler, in this case the "serial scheduler", interacts with the transactions and objects in a particularly constrained way. The "serial schedules" of the primitives and the serial scheduler are the basis of our correctness conditions. Specifically, alternative schedulers are required to ensure that nested transaction automata individually have local schedules which they could have in a serial schedule. In essence, each scheduler must "fool" the transactions into believing that the system is executing in conjunction with the serial scheduler.

In the past ten years, an important and substantial body of work has appeared on the design and analysis of algorithms for implementing concurrency control and resiliency in database transaction systems

[EGLT,RLS,BG,KS,Gr,LaS, etc.]. Among this has been a number of results dealing with nested transactions [R,Mo,LiS,LHJLSW,AM,BBGLS,BBG, etc.]. The present work does not replace these other contributions, but augments them by providing a unifying and mathematically tractable framework for posing and exploring a variety of questions. This previous work uses behavioral specifications of nested transactions, focusing on what nested transactions do, rather than what they are. By answering the question "What is a nested transaction?", I/O automata provide a powerful tool for understanding and reasoning about them.

Some unification is vitally important to further development in this field. The plethora and complexity of existing formalizations is a challenge to the most seasoned researcher. More critically, it belies the argument that nested transactions provide a clean and intuitive tool for organizing distributed databases and more general distributed applications. It is particularly important to provide an intuitive and precise description of nested transactions themselves, as in typical systems, these are the components which the application programmer must implement!

The remainder of this paper is organized as follows. The I/O automaton model is described in Section 2. The rest of the paper contains an extended example, which establishes correctness properties for two related lock-based concurrent schedulers.

Section 3 contains simple definitions for naming nested transactions and objects, and for specifying the operations (interactions) of these components. Simple syntactic restrictions on the orders of these operations are presented, and then a particular system of I/O automata is presented, describing the interactions of nested transactions and objects with a serial scheduler. The interface between the serial scheduler and the transactions provides a basis for the specification of correctness conditions for alternative schedulers. These schedulers would presumably be more efficient than the serial scheduler. The strongest correctness condition, "serial correctness," requires that *all* non-access transactions see serial behavior at their interface with the system. The second condition, "correctness for T_0," only requires that this serial interface be maintained at the interface of the system and the external world. These interfaces also provide simple descriptions of the environment in which nested transactions can be assumed to execute. A particular contribution is the clear and concise semantics of ABORT operations which arises naturally from this formalization. The section closes with a collection of lemmas describing useful properties of serial systems.

Next, a lock-based concurrent system is presented. Section 4 contains a description of a special type of object, called a "resilient object", which is used in the concurrent system. Section 5 describes the remainder of the concurrent system, the "concurrent scheduler." This concurrent scheduler includes "lock manager" modules for all the objects; lock managers coordinate concurrent accesses.

Section 6 defines a system which is closely related to the concurrent system, the "weak concurrent system." This system preserves serial correctness for those transactions whose ancestors do not abort (i.e., those that are not "orphans"). Since the root of the transaction tree, T_0, has no ancestor, weak concurrent systems are correct for T_0. Section 7 contains complete proofs of correctness of the concurrent and weak concurrent sytems; concurrent systems are serially correct, and weak concurrent systems are correct for T_0. The stronger condition is obtained for concurrent systems as a corollary to a result about weak concurrent systems.

It is interesting that the concurrent system algorithms are described in complete detail (essentially, in "pseudocode"), yet significant formal claims about their behavior can be stated clearly and easily. Although the full presentation involves a large number of lemmas, the ideas described by the lemmas are quite simple and intuitive. We think it is remarkable that these interesting properties of concurrent systems can be proved with complete rigor, in full detail, in so short a development. Despite the detailed level of presentation, the underlying model is general enough that the results apply to a wide range of implementations.

The style of the correctness proof is also noteworthy. It is a constructive proof, in that for each step of the weak concurrent system and each non-orphan transaction, an execution of the serial system is explicitly constructed. The transaction's local "view" in the constructed execution is identical to that in the original weak concurrent execution, establishing the correctness of the weak concurrent system. One may think of the weak concurrent system as maintaining consistent, parallel "world views" within which concurrent siblings execute. As siblings return to their parent, these parallel worlds are "merged" to form a single consistent view. The locking policy prevents collisions between different views at the shared data. This intuition is strongly supported and clarified by the correctness proof, which constructs the parallel views as different serial schedules consistent with each sibling's local history. Lemmas illustrate how these serial schedules can be merged as siblings return or abort to their parent.

Section 8 contains a discussion of the relationship of this work to previous results, and Section 9 contains an indication of the work that lies ahead.

2. Basic Model

In this section, we present the basic I/O automaton model, which is used to describe all components of our systems. This model consists of rather standard, possibly infinite-state, nondeterministic automata that have operation names associated with their state transitions. Communication among automata is described by identifying their operations. This model is very similar to models used by Milner, Hoare [Mi,Ho] and others. There are a few differences: first, we find it important to classify operations of any automaton or system of automata as either "input" or "output" operations, of that automaton or system, and we treat these two cases differently. Also, we allow identification of arbitrary numbers of operations from different automata, rather

than just pairwise identification as considered in [Mi].

This paper is not intended to develop the basic model. For the general theory of I/O automata, including a unified treatment of finite and infinite behavior, we refer the reader to [LT]. In the present treatment of concurrent transaction systems, we only prove properties of finite behavior, so we only require a simple special case of the general model.

2.1. I/O Automata

All components in our systems, transactions, objects and schedulers, will be modelled by *I/O automata.* An I/O automaton $\mathcal{A}$ has components *states*($\mathcal{A}$), *start*($\mathcal{A}$), *out*($\mathcal{A}$), *in*($\mathcal{A}$), and *steps*($\mathcal{A}$). Here, *states*($\mathcal{A}$) is a set of states, of which a subset *start*($\mathcal{A}$) is designated as the set of start states. The next two components are disjoint sets: *out*($\mathcal{A}$) is the set of *output operations*, and *in*($\mathcal{A}$) is the set of *input operations.* The union of these two sets is the set of *operations* of the automaton. Finally, *steps*($\mathcal{A}$) is the transition relation of $\mathcal{A}$, which is a set of triples of the form (s',π,s), where s' and s are states, and π is an operation. This triple means that in state s', the automaton can atomically do operation π and change to state s. An element of the transition relation is called a *step* of $\mathcal{A}$.

The output operations are intended to model the actions that are triggered by the automaton itself, while the input operations model the actions that are triggered by the environment of the automaton. Our partitioning of operations into input and output indicates that each operation is only triggered in one place. We require the following condition.

Input Condition: For each input operation π and each state s', there exist a state s and a step (s',π,s).

This condition says that an I/O automaton must be prepared to receive any input operation at any time. This condition makes intuitive sense if we think of the input operations as being triggered externally. (In this paper, this condition serves mainly as a technical convenience, but in [LT], where infinite behavior is considered, it is critical.)

An *execution* of $\mathcal{A}$ is an alternating sequence $s_0,\pi_1, s_1,\pi_2,\ldots$ of states and operations of $\mathcal{A}$; the sequence may be infinite, but if it is finite, it ends with a state. Furthermore, s_0 is in start($\mathcal{A}$), and each triple (s',π,s) which occurs as a consecutive subsequence is a step of $\mathcal{A}$. From any execution, we can extract the *schedule*, which is the subsequence of the execution consisting of operations only. Because transitions to different states may have the same operation, different executions may have the same schedule.

Lemma 1: If α is a schedule of I/O automaton $\mathcal{A}$, then every prefix of α is a schedule of $\mathcal{A}$.

If S is any set of schedules (or property of schedules), then $\mathcal{A}$ is said to *preserve* S provided that the following holds. If $\alpha = \alpha'\pi$ is any schedule of $\mathcal{A}$, where π is an output operation, and α' is in S, then α is in S. That is, the automaton is not the first to violate the property described by S.

2.2. Composition of Automata

We describe systems as consisting of interacting components, each of which is an I/O automaton. It is convenient and natural to view systems as I/O automata, also. Thus, we define a composition operation for I/O automata, to yield a new I/O automaton.

A set of I/O automata may be composed to create a *system* $\mathcal{S}$, if all of the output operations are disjoint. (Thus, every output operation in $\mathcal{S}$ will be triggered by exactly one component.) The system $\mathcal{S}$ is itself an I/O automaton. A state of the composed automaton is a tuple of states, one for each component, and the start states are tuples consisting of start states of the components. The set of *operations* of $\mathcal{S}$, *ops*($\mathcal{S}$), is exactly the union of the sets of operations of the component automata. The set of *output operations* of $\mathcal{S}$, *out*($\mathcal{S}$), is likewise the union of the sets of output operations of the component automata. Finally, the set of *input operations* of $\mathcal{S}$, *in*($\mathcal{S}$), is *ops*($\mathcal{S}$) - *out*($\mathcal{S}$), the set of operations of $\mathcal{S}$ that are not output operations of $\mathcal{S}$. The output operations of a system are intended to be exactly those that are triggered by components of the system, while the input operations of a system are those that are triggered by the system's environment.

The triple (s',π,s) is in the transition relation of $\mathcal{S}$ if and only if for each component automaton $\mathcal{A}$, one of the following two conditions holds. Either π is an operation of $\mathcal{A}$, and the projection of the step onto $\mathcal{A}$ is a step of $\mathcal{A}$, or else π is not an operation of $\mathcal{A}$, and the states corresponding to $\mathcal{A}$ in the two tuples s' and s are identical. Thus, each operation of the composed automaton is an operation of a subset of the component automata. During an operation π of $\mathcal{S}$, each of the components which has operation π carries out the operation, while the remainder stay in the same state. Again, the operation π is an output operation of the composition if it is the output operation of a component – otherwise, π is an input operation of the composition.[1]

Lemma 2: The composition of I/O automata is an I/O automaton.

The next lemma allows us to compose automata in any order.

Lemma 3: Up to isomorphism, composition of I/O automata is associative and commutative.

---

[1] Note that our model has chosen a particular convention for identifying operations of different components in a system: we simply identify those with the same name. This convention is simple, and sufficient for what we do in this paper. However, when this work is extended to more complicated systems, it may be expedient to generalize the convention for identifying operations, to permit reuse of the same operation name internally to different components. This will require introducing a renaming operator for operations, or else defining composition with respect to a designated equivalence relation on operations. We leave this for later work.

An *execution* of a system is defined to be an execution of the automaton composed of the individual automata of the system. If α is a schedule of a system with component $\mathcal{A}$, then we denote by $\alpha|\mathcal{A}$ the subsequence of α containing all the operations of $\mathcal{A}$. Clearly, $\alpha|\mathcal{A}$ is a schedule of $\mathcal{A}$.

Lemma 4: Let α' be a schedule of a system $\mathcal{S}$, and let $\alpha = \alpha'\pi$, where π is an output operation of component $\mathcal{A}$. If $\alpha|\mathcal{A}$ is a schedule of $\mathcal{A}$, then α is a schedule of $\mathcal{S}$.

Proof: Since $\alpha|\mathcal{A}$ is a schedule of $\mathcal{A}$, there is an execution β of $\mathcal{A}$ with schedule $\alpha|\mathcal{A}$. Let β' be the execution of $\mathcal{A}$ consisting of all but the last step of β. Similarly, since α' is a schedule of $\mathcal{S}$, there is an execution γ of $\mathcal{S}$ with schedule α'. It is possible that $\mathcal{A}$ has an execution in γ which is different from β', since different executions may have the same schedule. But it is easy to show, by induction on the length of γ, that there is another execution γ' of $\mathcal{S}$ in which component $\mathcal{A}$ has execution β', and which is otherwise identical to γ. The schedule of γ' is α'. Since π is not an output operation of any other component, π is defined from the state reached at the end of γ', so that $\alpha = \alpha'\pi$ is a schedule of $\mathcal{S}$. ∎

3. Serial Systems

In this paper, we define three kinds of systems: "serial systems" and two types of "concurrent systems". Serial systems describe serial execution of transactions. Serial systems are defined for the purpose of providing a correctness condition for other systems: that the schedules of the other systems should "look like" schedules of the serial system to the transactions. As with serial schedules of single-level transaction systems, our serial schedules are too inefficient to use in practice. Thus, we define systems which allow concurrency, and which permit the abort of transactions after they have performed some work. We then prove that the schedules permitted by concurrent systems are correct.

In this section, we define "serial systems". Serial systems consist of "transactions" and "basic objects" communicating with a "serial scheduler". Transactions and basic objects describe user programs and data, respectively. The serial scheduler controls communication between the other components, and thereby defines the allowable orders in which the transactions may take steps. All three types of system components are modelled as I/O automata.

We begin by defining a structure which describes the nesting of transactions. Namely, a *system type* is a four-tuple $(\mathcal{T}, \text{parent}, O, V)$, where $\mathcal{T}$, the set of transaction names, is organized into a tree by the mapping $\text{parent}: \mathcal{T} \rightarrow \mathcal{T}$, with T_0 as the root. In referring to this tree, we use traditional terminology, such as child, leaf, least common ancestor (lca), ancestor and descendant. (A transaction is its own ancestor and descendant.) The leaves of this tree are called *accesses*. The set O denotes the set of objects; formally, O is a partition of the set of accesses, where each element of the partition contains the accesses to a particular object. The set V is a set of *values*, to be used as return values of transactions.

The tree structure can be thought of as a predefined naming scheme for all possible transactions that might

ever be invoked. In any particular execution, however, only some of these transactions will actually take steps. We imagine that the tree structure is known in advance by all components of a system. The tree will, in general, be an infinite structure.

The classical transactions of concurrency control theory (without nesting) appear in our model as the children of a "mythical" transaction, T_0, the root of the transaction tree. (In work on nested transactions, such as ARGUS [LiS,LHJLSW], the children of T_0 are often called "top-level" transactions.) It is very convenient to introduce the new root transaction to model the environment in which the rest of the transaction system runs. Transaction T_0 has operations that describe the invocation and return of the classical transactions. It is natural to reason about T_0 in much the same way as about all of the other transactions, although it is distinguished from the other transactions by having no parent transaction. Since committing and aborting are operations which take place at the parent of each transaction (see below), T_0 can neither commit nor abort. Thus, a commit or abort of a top-level transaction to T_0 is an irreversible step.

The internal nodes of the tree model transactions whose function is to create and manage subtransactions, but not to access data directly. The only transactions which actually access data are the leaves of the transaction tree, and thus they are distinguished as "accesses". The partition $\mathcal{O}$ simply identifies those transactions which access the same object.

A serial system of a given system type is the composition of a set of I/O automata. This set contains a transaction for each *internal* (i.e. non-leaf, non-access) node of the transaction tree, a basic object for each element of $\mathcal{O}$ and a serial scheduler. These automata are described below. (If X is a basic object associated with an element $\mathcal{X}$ of the partition $\mathcal{O}$, and T is an access in $\mathcal{X}$, we write $T \in accesses(X)$ and say that "T is an access to X".)

3.1. Transactions

This paper differs from earlier work such as [Ly,Go,Wel] in that we model the transactions explicitly, as I/O automata. In modelling transactions, we consider it very important not to constrain them unnecessarily; thus, we do not want to require that they be expressible as programs in any particular high-level programming language. Modelling the transactions as I/O automata allows us to state exactly the properties that are needed, without introducing unnecessary restrictions or complicated semantics.

A non-access *transaction* T is modelled as an I/O automaton, with the following operations.

Input operations:
CREATE(T)
COMMIT(T',v), for T' $\in$ children(T) and v $\in$ V
ABORT(T'), for T' $\in$ children(T)

Output operations:
REQUEST–CREATE(T'), for T' ∈ children(T)
REQUEST–COMMIT(T,v), for v ∈ V

The CREATE input operation "wakes up" the transaction. The REQUEST–CREATE output operation is a request by T to create a particular child transaction.[2] The COMMIT input operation reports to T the successful completion of one of its children, and returns a value recording the results of that child's execution. The ABORT input operation reports to T the unsuccessful completion of one of its children, without returning any other information. We call COMMIT(T',v), for any v, and ABORT(T') *return* operations for transaction T'. The REQUEST–COMMIT operation is an announcement by T that it has finished its work, and includes a value recording the results of that work.

It is convenient to use two separate operations, REQUEST–CREATE and CREATE, to describe what takes place when a subtransaction is activated. The REQUEST–CREATE is an operation of the transaction's parent, while the actual CREATE takes place at the subtransaction itself. In actual systems such as ARGUS, this separation does occur, and the distinction will be important in our results and proofs. Similar remarks hold for the REQUEST–COMMIT and COMMIT operations.[3] We leave the executions of particular transaction automata largely unspecified; the choice of which children to create, and what value to return, will depend on the particular implementation. For the purposes of the schedulers studied here, the transactions (and in large part, the objects) are "black boxes." Nevertheless, it is convenient to assume that schedules of transaction automata obey certain syntactic constraints. Thus, transaction automata are required to preserve well-formedness, as defined below.

We recursively define *well-formedness* for sequences of operations of transaction T. Namely, the empty schedule is well-formed. Also, if $\alpha = \alpha'\pi$ is a sequence of operations of T, where π is a single operation, then α is well-formed provided that α' is well-formed, and the following hold.

- If π is CREATE(T), then
  (i) there is no CREATE(T) in α'.

- If π is COMMIT(T',v) or ABORT(T') for a child T' of T, then

---

[2] Note that there is no provision for T to pass information to its child in this request. In a programming language, T might be permitted to pass parameter values to a subtransaction. Although this may be a convenient descriptive aid, it is not necessary to include in it the underlying formal model. Instead, we consider transactions that have different input parameters to be different transactions.

[3] Note that we do not include a REQUEST–ABORT operation for a transaction: we do not model the situation in which a transaction decides that its own existence is a mistake. Rather, we assign decisions to abort transactions to another component of the system, the scheduler. In practice, the scheduler must have some power to decide to abort transactions, as when it detects deadlocks or failures. In ARGUS, transactions are permitted to request to abort; we regard this request simply as a "hint" to the scheduler, to restrict its allowable executions in a particular way. This operation could be made explicit, constraining the scheduler to abort the requesting transaction, without substantively changing the model or results.

(i) REQUEST–CREATE(T') appears in α' and
(ii) there is no return operation for T' in α'.

- If π is REQUEST–CREATE(T') for a child T' of T, then
  (i) there is no REQUEST–CREATE(T') in α'
  (ii) there is no REQUEST–COMMIT(T) in α' and
  (iii) CREATE(T) appears in α'.

- If π is a REQUEST–COMMIT for T, then
  (i) there is no REQUEST–COMMIT for T in α' and
  (ii) CREATE(T) appears in α'.

These restrictions are very basic; they simply say that a transaction does not get created more than once, does not receive repeated notification of the fates of its children, does not receive conflicting information about the fates of its children, and does not receive information about the fate of any child whose creation it has not requested; also, a transaction does not perform any output operations before it has been created or after it has requested to commit, and does not request the creation of the same child more than once. Except for these minimal conditions, there are no restrictions on allowable transaction behavior. For example, the model allows a transaction to request to commit without discovering the fate of all subtransactions whose creation it has requested. Also, a transaction can request creation of new subtransactions at any time, without regard to its state of knowledge about subtransactions whose creation it has previously requested. Particular programming languages may choose to impose additional restrictions on transaction behavior. (An example is ARGUS, which suspends activity in transactions until subtransactions complete.) However, our results do not require such restrictions.

The following easy lemma summarizes the properties of well-formed sequences of transaction operations.

Lemma 5: Let α be a well-formed sequence of operations of transaction T. Then the following conditions hold.

1. The first operation of α is a CREATE(T) operation, and there are no other CREATE operations.

2. If a REQUEST–COMMIT operation occurs in α, then there are no later output operations in α.

3. There is at most one REQUEST–CREATE(T') operation for each child T' of T, in α.

4. Every return operation in α has a preceding REQUEST–CREATE operation in α for the same child transaction.

3.2. Basic Objects

Recall that I/O automata are associated with non-access transactions only. Since access transactions model abstract operations on shared data objects, we associate a single I/O automaton with each object, rather than one for each access. The operations for each object are just the CREATE and REQUEST−COMMIT operations for all the corresponding access transactions. Although we give these operations the same names as the operations of non-access transactions, it is helpful to think of the operations of access transactions in other terms also: a CREATE corresponds to an invocation of an operation on the object, while a REQUEST−COMMIT corresponds to a response by the object to an invocation. Actually, these CREATE and REQUEST−COMMIT operations generalize the usual invocations and responses in that our operations carry with them a designation of the position of the access in the transaction tree. We depart from the traditional notational distinction between creation of subtransactions and invocations on objects, since the common terminology for access and non-access transactions is of great benefit in unifying the statements and proofs of our results. Thus, a *basic object* X is modelled as an automaton, with the following operations.

Input operations:
CREATE(T), for T in accesses(X)

Output operations:
REQUEST−COMMIT(T,v), for T in accesses(X)

The CREATE operation is an invocation of an access to the object, while the REQUEST−COMMIT is a return of a value in response to such an invocation.

As with transactions, while specific objects are left largely unspecified, it is convenient to require that schedules of basic objects satisfy certain syntactic conditions. Thus, each basic object is required to preserve well-formedness, defined below.

Let α be a sequence of operations of basic object X. Then an access T to X is said to be *pending* in α provided that there is a CREATE(T), but no REQUEST−COMMIT for T, in α. We define *well-formedness* for sequences of operations of basic objects recursively. Namely, the empty schedule is well-formed. Also, if $\alpha = \alpha'\pi$ is a sequence of operations of basic object X, where π is a single operation, then α is well-formed provided that α' is well-formed, and the following hold.

- If π is CREATE(T), then
  (i) there is no CREATE(T) in α', and
  (ii) there are no pending accesses in α'.

- If π is REQUEST−COMMIT for T, then
  (i) there is no REQUEST−COMMIT for T in α', and
  (ii) CREATE(T) appears in α'.

These restrictions simply say that the same access does not get created more than once, nor does a creation of a new access occur at a basic object before the previous access has completed (i.e. requested to commit); also, a basic object does not respond more than once to any access, and only responds to accesses that have previously been created.

The following easy lemma summarizes the properties of well-formed sequences of basic object operations.

Lemma 6: Let α be a well-formed sequence of operations of basic object X. Then α consists of alternating CREATE and REQUEST−COMMIT operations, starting with a CREATE, and with each consecutive (CREATE,REQUEST−COMMIT) pair having a common transaction.

3.3. Serial Scheduler

The third kind of component in a serial system is the serial scheduler. The serial scheduler is also modelled as an automaton. The transactions and basic objects have been specified to be any I/O automata whose operations and behavior satisfy simple syntactic restrictions. The serial scheduler, however, is a fully specified automaton, particular to each system type. It runs transactions according to a depth-first traversal of the transaction tree. The serial scheduler can choose nondeterministically to abort any transaction after its parent has requested its creation, as long as the transaction has not actually been created. In the context of this scheduler, the "semantics" of an ABORT(T) operation are that transaction T was never created. The operations of the serial scheduler are as follows.

Input Operations:
REQUEST−CREATE(T)
REQUEST−COMMIT(T,v)

Output Operations:
CREATE(T)
COMMIT(T,v)
ABORT(T)

The REQUEST−CREATE and REQUEST−COMMIT inputs are intended to be identified with the corresponding outputs of transaction and object automata, and correspondingly for the CREATE, COMMIT and ABORT output operations. Each state s of the serial scheduler consists of four sets: create−requested(s), created(s), commit−requested(s), and returned(s). The set commit−requested(s) is a set of (transaction,value) pairs. The others are sets of transactions. There is exactly one initial state, in which the set create−requested is $\{T_0\}$, and the other sets are empty.

The transition relation consists of exactly those triples (s',π,s) satisfying the pre- and postconditions below, where π is the indicated operation. For brevity, we include in the postconditions only those conditions on the state s which may change with the operation. If a component of s is not mentioned in the postcondition, (such

as returned(s) in the postcondition for REQUEST−CREATE(T)), it is implicit that the set is the same in s' and s (that returned(s') = returned(s), in this example). Note that here, as elsewhere, we have tried to specify the component as nondeterministically as possible, in order to achieve the greatest possible generality for our results.

- REQUEST−CREATE(T)
  Postcondition:
  create−requested(s) = create−requested(s') ∪ {T}

- REQUEST−COMMIT(T,v)
  Postcondition:
  commit−requested(s) = commit−requested(s') ∪ {(T,v)}

- CREATE(T)
  Precondition:
  T ∈ create−requested(s') - created(s')
  siblings(T) ∩ created(s') ⊆ returned(s')
  Postcondition:
  created(s) = created(s') ∪ {T}

- COMMIT(T,v)
  Precondition:
  (T,v) ∈ commit−requested(s')
  T ∉ returned(s')
  children(T) ∩ create−requested(s') ⊆ returned(s')
  Postcondition:
  returned(s) = returned(s') ∪ {T}

- ABORT(T)
  Precondition:
  T ∈ create−requested(s') - created(s')
  siblings(T) ∩ created(s') ⊆ returned(s')
  Postcondition:
  created(s) = created(s') ∪ {T}
  returned(s) = returned(s') ∪ {T}

The input operations, REQUEST−CREATE and REQUEST−COMMIT, simply result in the request being recorded. A CREATE operation can only occur if a corresponding REQUEST−CREATE has occurred and the CREATE has not already occurred. The second precondiition on the CREATE operation says that the serial scheduler does not create a transaction until all its previously created sibling transactions have returned. That is, siblings are run sequentially. The precondition on the COMMIT operation says that the scheduler does not allow a transaction to commit to its parent until its children have returned. The precondition on the ABORT operation says that the scheduler does not abort a transaction while there is activity going on on behalf of any of its siblings. That is, aborted transactions are run sequentially with respect to their siblings. The next lemma relates a schedule of the serial scheduler to the state which results

from applying that schedule.

Lemma 7: Let α be a schedule of the serial scheduler, and let s be a state which can result from applying α to the initial state. Then the following conditions are true.

1. T is in create−requested(s) exactly if $T = T_0$ or α contains a REQUEST−CREATE(T) operation.

2. T is in created(s) exactly if α contains either a CREATE(T) or ABORT(T) operation.

3. (T,v) is in commit−requested(s) exactly if α contains a REQUEST−COMMIT(T,v) operation.

4. T is in returned(s) exactly if α contains a return operation for T.

3.4. Serial Systems and Serial Schedules

In this subsection, we define serial systems precisely and provide some useful terminology for talking about them.

The composition of transactions with basic objects and the serial scheduler for a given system type is called a *serial system.* Define the *serial operations* to be those operations which occur in the serial system: REQUEST−CREATES, REQUEST−COMMITS, CREATES, COMMITS and ABORTS. The schedules of a serial system are called *serial schedules.* The non-access transactions and basic objects are called the system *primitives.* (Recall that each basic object is an automaton corresponding to a set of access transactions. Thus, individual access transactions are not considered to be primitives.)

Recall that the operations of the basic objects have the same syntax as transaction operations. It is convenient to refer to CREATE(T) and REQUEST−COMMIT(T), when T is an access to basic object X, both as operations of transaction T and of object X. To avoid confusion, it is important to remember that there is no transaction automaton associated with any access operation.

For any serial operation π, we define *transaction*(π) to be the transaction at which the operation occurs. (For CREATE(T) operations and REQUEST−COMMIT operations for T, the transaction is T, while for REQUEST−CREATE(T) operations, and COMMIT and ABORT operations for T, the transaction is parent(T).) For a sequence α of serial operations, transaction(α) is the set of transactions of the operations in α.

Two sequences of serial operations, α and α', are said to be *equivalent* provided that they consist of the same operations, and $\alpha|P = \alpha'|P$ for each primitive P. Obviously, this yields an equivalence relation on sequences of serial operations.

We let $\alpha|T$ denote the subsequence of α consisting of operations whose transaction is T, even if T is an access. (This is an extension of the previous definition of $\alpha|T$, as accesses are not component automata of the serial system.)

Let α be a sequence of serial operations. We say that a transaction T is *live* in α provided that a CREATE(T), but no COMMIT(T,v) or ABORT(T), occurs in α. We say that transaction T' is *visible* to T in α provided that for each ancestor T'' of T' which is a proper descendant of lca(T,T'), some COMMIT(T'',v) occurs in α. (In particular, any ancestor of T is visible to T in α.) For sequence α and transaction T, let *visible*(α,*T)* be the subsequence of α consisting of operations whose transactions are visible to T in α. (These include access transactions T'.) We say that transaction T *sees everything* in α provided that visible(α,T) = α.

This is the same definition of visibility as appears, in a different model, in [Ly]. Visibility captures an intuitive notion suggested by the name: the transactions visible to a transaction T in α are those whose effects T is permitted to "see" in α. If transaction T' is visible to transaction T in α, it means that descendants of T may have passed to T information about T', obtained by accessing objects that were previously accessed by descendants of T'.

If α is a sequence of operations, not necessarily all serial, then define serial(α) to be the subsequence of α consisting of the serial operations. We say that T is *live* in α provided that it is live in serial(α). We say that T' is *visible* to T in α if T' is visible to T in serial(α), and define visible(α,T) to be visible(serial(α),T). Also, T *sees everything* in α provided that T sees everything in serial(α). Similarly, define transaction(α) = transaction(serial(α)).

A sequence α of serial operations is said to be *well-formed* if its projection at every primitive is well-formed.

3.5. Correctness Condition

We use serial schedules as the basis of our correctness definitions. Namely, we say that a sequence of operations is *serially correct for a primitive P* provided that its projection on P is identical to the projection on P of some serial schedule. We say that any sequence of operations is *serially correct* if it is serially correct for every non-access transaction. That is, α "looks like" a serial schedule to every non-access transaction.

In the remainder of this paper, we define two systems: concurrent systems and weak concurrent systems. We show that schedules of concurrent systems are serially correct, and that schedules of weak concurrent systems are serially correct for T_0.

Thus, we use the serial scheduler as a way of describing desirable behavior, just as serial schedules describe

desirable behavior in more classical concurrency control settings (those without nesting). Then serial correctness plays the role in our theory that serializability plays in classical settings.

Motivation for our use of serial schedules to define correctness derives from the simple behavior of the serial scheduler, which determines the sequence of interactions between the primitives. Each transaction T is created only after parent(T) requests it, no siblings of T are created until T has returned, T is not committed until each of its requested children has itself returned, and T is not aborted until each of its created siblings has returned. The result is a depth-first traversal of the transaction tree, with requests flowing down and responses flowing up. We believe this depth-first traversal to be a natural notion of correctness which corresponds precisely to the intuition of how nested transaction systems ought to behave. Furthermore, it is a natural generalization of serializability, the correctness condition generally chosen for classical transaction systems.

Serial correctness is a condition which guarantees to implementors of transactions that their code will encounter only situations which can arise in serial executions. Correctness for T_0 is a natural alternative, which guarantees only that the external world will encounter only situations which can arise in serial executions. This condition permits less constrained implementations, in that schedulers in such systems need not insure that orphans see consistent data. On the other hand, in such systems the authors of transactions must insure that their programs behave well even if they see inconsistencies. (For example, orphans that see inconsistent data should not consume too many system resources, garble data beyond repair, dispense drugs or initiate military hostilities.) We hope this work will provide a tool for exploring the inherent costs of different correctness conditions such as these.

Note that our correctness conditions are defined at the transaction interface only, and do not constrain the object interface. We believe that this makes the conditions more meaningful to users, and more likely to suffice for a large variety of algorithms, which may use a variety of back-out, locking or version schemes to implement objects. Previous work has focussed on correctness conditions at the object interface [EGLT, etc.]. While we believe that object interface conditions are important, their proper role in the theory is not to serve as the basic correctness condition. Rather, they are useful as intermediate conditions for proving correctness of particular implementations: such conditions can be shown to be sufficient, in combination with an appropriate scheduler, to ensure our correctness condition at the transaction interface. This observation is an important unifying contribution of our work. Our current research is focussing on demonstrating the usefulness of this approach, for a variety of object interface correctness conditions.

The serial correctness condition says that a schedule α must look like a serial schedule to each non-access transaction; this allows for the possibility that α might look like *different* serial schedules to different non-

access transactions. This condition may at first seem to be too weak. It may seem that we should require that all transactions see a projection of the *same* serial schedule. But this stronger condition is not satisfied by most of the known concurrency control algorithms. It is true that stronger conditions than ours can sometimes be proved, but such conditions are more complicated to state, and it is not yet clear which such conditions are most interesting.

The serial correctness condition is really not as weak as it may seem at first because T_0, the root transaction, is included among the transactions to which α must appear serial. As discussed above, transaction T_0 can be thought of as modelling the environment in which the rest of the transaction system runs. Its REQUEST–CREATE operations correspond to the invocation of top-level transactions, while its COMMIT and ABORT operations correspond to return values and external effects of those transactions. Since α's projection on T_0 must be serial, the environment of the transaction system will see only results that could arise in a serial execution. Indeed, this is the justification of the correctness condition for the weak concurrent system, whose schedules are shown to be correct for T_0, but not necessarily for any other transaction.

It is possible to use a different serial scheduler as a basis for correctness conditions. For example, the scheduler might delay creating one sibling until another *requests* to return, rather than until it actually returns to the parent [We2]. Such a scheduler would provide less information to the parent about the actual order in which its children are executed, and consequently provide more freedom for concurrent schedulers to schedule various events. Timestamp-based systems such as [R] may support this weaker correctness condition, rather than the one described above, but this remains to be studied.

Our approach is really a general technique for studying operating system algorithms. A simple, intuitive and inefficient algorithm (automaton) is used to specify a "contract" between the users and implementor of an operating system. The user is guaranteed that applications (transactions, in our work) which are correct when run with the simple algorithm will also be correct when run with the actual operating system, which presumably will be more efficient. On the other hand, the implementor also has a formal and intuitive specification of the user interface.

3.6. Properties of Serial Systems

In this subsection, we prove a number of lemmas about the behavior of serial systems. They are collected here for reference later in this paper and in future work. Most of the lemmas describe properties that are quite easy to understand and believe, and the corresponding proofs are very straightforward. In the last paragraph of this subsection, there are some specialized lemmas that are somewhat more difficult. These are used in the proof of the main theorem in Section 7.

3.6.1. Fundamental Properties of Visibility

The first few lemmas give fundamental properties of visibility in sequences of serial operations. In this paragraph, we do not even require that the sequences be schedules of serial systems, but only that they be sequences of serial operations. The proofs of these lemmas are straightforward from the definitions.

Lemma 8: Let α be a sequence of serial operations, and T, T' and T'' transactions.

1. If T' is a descendant of T, then T is visible to T' in α.
2. T'' is visible to T in α if and only if T'' is visible to lca(T,T') in α.
3. If T''' is visible to T'' in α and T'' is visible to T in α, then T''' is visible to T in α.
4. If T'' is a descendant of T and T''' is visible to T in α, then T''' is visible to T'' in α.
5. If T'' is a descendant of T and T'' is visible to T''' in α, then T is visible to T''' in α.
6. If T'' is a proper descendant of T, T''' is visible to T'' in α, but T''' is not visible to T in α, then T''' is a descendant of the child of T which is an ancestor of T''.

Lemma 9: Let α and β be sequences of serial operations, with β a subsequence of α.

1. If transaction T is visible to transaction T' in β, then it is visible to transaction T' in α.
2. If operation π is in visible(β,T), then it is in visible(α,T).

Lemma 10: Let α, α', β and β' be sequences of serial operations, and let T and T' be transactions.

1. If α is equivalent to α', and T'' is visible to T in α, then T'' is visible to T in α'.
2. If α is equivalent to α', then visible(α,T) is equivalent to visible(α',T).
3. If β is equivalent to β', then $\alpha \cdot \beta = \alpha \cdot \beta'$.
4. If α is equivalent to α', and β is equivalent to β', then $\alpha \cdot \beta$ is equivalent to $\alpha' \cdot \beta'$.
5. If β = visible(α,T), then T sees everything in β.
6. If β is equivalent to visible(α,T), then T sees everything in β.
7. If β = visible(α,T) and T'' is visible to T in α, then visible(β,T'') = visible(α,T'').
8. If β is equivalent to visible(α,T), β' is equivalent to visible(α,T''), and T'' is visible to T in α, then β' is equivalent to visible(β,T'').

Lemma 11: Let α be a sequence of serial operations, and let T and T' be transactions. Then visible(α,T)|T' is equal to α|T' if T' is visible to T in α, and is equal to the empty string otherwise.

Lemma 12: Let $\alpha\pi$ be a sequence of serial operations, where π is a single operation. Let T be a transaction and assume that transaction(π) is visible to T in $\alpha\pi$. Assume that π is not a COMMIT operation. Then visible($\alpha\pi$,T) = visible(α,T)π.

3.6.2. Operations in Serial Schedules

The lemmas in this paragraph describe the kinds and orders of operations that can occur in well-formed serial schedules. In the next paragraph, we show that all serial schedules are well-formed, so that all these properties actually follow just from the fact that the schedules are serial.

Lemma 13: Let α be a well-formed serial schedule, and let $T \neq T_0$ be a transaction.

1. If α contains any operation with transaction T, then α contains a REQUEST–CREATE(T).

2. If α contains a COMMIT for T, then α contains a REQUEST–COMMIT for T, a CREATE(T) and a REQUEST–CREATE(T).

3. If α contains an ABORT(T), then α contains a REQUEST–CREATE(T).

Proof: Straightforward from well-formedness and the scheduler preconditions. ∎

Lemma 14: Let α be a well-formed serial schedule, and T a transaction. Assume that some descendant of T is in transaction(α). Then the following hold.

1. CREATE(T) occurs in α.

2. If $T \neq T_0$, then REQUEST–CREATE(T) occurs in α.

Proof: 1. By induction on the length of α. The basis is easy. Let $\alpha = \alpha'\pi$, where π is a single operation, and assume that the result holds for α'. Let T" = transaction(π), and let T be any ancestor of T". We must show that CREATE(T) occurs in α.

Because α is well-formed, CREATE(T") occurs in α. If T = T", we are done. Otherwise, Lemma 13 implies that REQUEST–CREATE(T") occurs in α. This occurs at parent(T"), which is a descendant of T. The inductive hypothesis then implies that α contains a CREATE(T).

2. By part 1. and Lemma 13. ∎

Lemma 15: Let α be a serial schedule, and let T be a transaction. Then α cannot contain both a CREATE(T) and an ABORT(T) operation.

Proof: By the scheduler preconditions. ∎

Lemma 16: Let α be a well-formed serial schedule, and let T be a transaction. If ABORT(T) occurs in α, then α contains no operations whose transactions are descendants of T.

Proof: Assume the contrary. Then Lemma 14 implies that a CREATE(T) operation occurs in α. But Lemma 15 yields a contradiction. ∎

Lemma 17: Let α be a well-formed serial schedule, and let $T \neq T_0$ be a transaction.

1. If α contains a REQUEST–CREATE(T), but does not contain a return operation for T, then parent(T) is live in α.

2. If T is live in α, then parent(T) is live in α.

3. If α contains a REQUEST–CREATE(T) but does not contain a CREATE(T) or an ABORT(T), then parent(T) is live in α.

Proof:

1. Well-formedness implies that the REQUEST−CREATE(T) is preceded in α by a CREATE(parent(T)). Suppose that parent(T) is not live in α. Then a return operation for parent(T) occurs in α. By Lemma 15, ABORT(parent(T)) cannot appear in α. Thus, a COMMIT operation for parent(T) must appear in α. This COMMIT operation for parent(T) must be preceded by a REQUEST−COMMIT for parent(T), by the scheduler preconditions. By well-formedness, the REQUEST−COMMIT for parent(T) must follow the REQUEST−CREATE(T) operation, so that the COMMIT for parent(T) follows the REQUEST−CREATE(T) operation. Then by the scheduler preconditions for the COMMIT operation, there must be a return operation for T in α, a contradiction.

2. Since T is live in α, CREATE(T) occurs in α and so Lemma 13 implies that REQUEST−CREATE(T) occurs in α. The result then follows from part 1.

3. Since there is no CREATE(T) in α, there can be no REQUEST−COMMIT for T, by well-formedness. Then there can be no COMMIT for T, by the scheduler preconditions. The result follows from part 1.

∎

Lemma 18: Let α be a well-formed serial schedule, and let T be a transaction.

1. If α contains a REQUEST−CREATE(T) but does not contain a return operation for T, then any proper ancestor of T is live in α.

2. If T is live in α, then any ancestor of T is live in α.

3. If α contains a REQUEST−CREATE(T) but does not contain a CREATE(T) or an ABORT(T), then any proper ancestor of T is live in α.

Proof: By repeated use of Lemma 17. ∎

Lemma 19: Let α be a well-formed serial schedule, and let T and T' be transactions with T' a descendant of T. Assume that there is a COMMIT operation for T in α.

1. If a REQUEST−CREATE(T') occurs in α, then there is a return operation for T' in α.

2. If T' is in transaction(α), then there is a COMMIT operation for T' in α.

Proof:

1. By Lemma 18.

2. Lemma 13 implies that REQUEST−CREATE(T') occurs in α. Part 1 then implies that there is a return operation for T' in α. Since T' is in transaction(α), Lemma 16 implies that there cannot be an ABORT(T') in α. Thus, there is a COMMIT for T' in α.

∎

Lemma 20: Let α be a well-formed serial schedule.
If a return operation for T is in α, then it follows all operations in α whose transaction is T.

Proof: Lemma 16 implies the result if an ABORT(T) occurs in α. So assume that a COMMIT for T occurs in α. This must be preceded by a REQUEST−COMMIT for T, by scheduler preconditions. Well-formedness implies that the REQUEST−COMMIT is preceded by a

CREATE(T), and is not followed by any output operations of T. Thus, the only operations of T that could follow the REQUEST–COMMIT are return operations for children of T. Let T' be a child of T for which a return operation occurs in α. By scheduler preconditions, there is only one return operation for T' in α. By Lemma 13, α also contains a REQUEST–CREATE(T'). Since this is an output operation of T, it precedes the REQUEST–COMMIT for T, and hence precedes the COMMIT for T. Then the scheduler preconditions imply that the return operation for T' precedes the COMMIT for T. ∎

Lemma 21: Let α be a well-formed serial schedule.
If a return operation for T is in α, then it follows all operations in α whose transactions are descendants of T.

Proof: Since a return operation for T occurs in α, we have $T \neq T_0$. Let T' be a descendant of T, and assume for the sake of obtaining a contradiction that an operation π with transaction(π) = T' occurs after the return for T in α. Let α' be the prefix of α preceding π.

Lemma 16 implies the result if an ABORT(T) occurs in α. So assume that a COMMIT for T occurs in α. By Lemma 13, α' contains a REQUEST–CREATE(T') operation. Then Lemma 19 implies that α' contains a return operation for T'. But then the well-formed schedule $\alpha'\pi$ contains a return for T' followed by an operation of T', which contradicts Lemma 20. ∎

Lemma 22: Let α be a well-formed serial schedule. If T is a pending access in $\alpha|X$, then T is live in α.

Proof: If T is a pending access in $\alpha|X$, then a CREATE(T) occurs in α, but no REQUEST–COMMIT for T occurs in α. Thus, by the scheduler preconditions, no COMMIT for T can occur in α. ∎

Lemma 23: Let α be a well-formed serial schedule, and let T and T' be distinct transactions live in α. Then the following are true.

1. T and T' are not siblings.

2. Either T is an ancestor of T' or vice versa.

Proof:

1. Assume the contrary. Assume without loss of generality that CREATE(T) precedes CREATE(T') in α. Then the scheduler preconditions for the CREATE(T') operation imply that a return operation for T occurs in α. This contradicts the assumption that T is live in α.

2. By part 1 and Lemma 18.

∎

3.6.3. Well-Formedness

Now we show that all serial schedules are well-formed. It follows that all the properties proved in the previous paragraph for well-formed serial schedules are actually true for all serial schedules. Subsequently, we will use these properties without explicitly mentioning well-formedness.

Lemma 24: Let α be a serial schedule. Then α is well-formed.

Proof: By induction on the length of schedules. The base, length = 0, is trivial. Suppose that $\alpha\pi$ is a serial schedule, and assume that α is well-formed. If π is an output of a primitive P, then

$\alpha\pi|P$ is well-formed because P preserves well-formedness, and so $\alpha\pi$ is well-formed. So assume that π is an input to a primitive P. It suffices to show that $\alpha\pi|P$ is well-formed. There are four cases.

(1) π is CREATE(T) and T is a non-access transaction.
The scheduler preconditions insure that CREATE(T) does not appear in α.

(2) π is COMMIT(T,v) for some transaction T and value v.
Then π is an input to transaction parent(T) = T'. The scheduler preconditions imply that α contains a REQUEST–COMMIT(T,v), and so Lemma 13 implies that α contains a REQUEST–CREATE(T). Also, the scheduler preconditions imply that no return operation for T occurs in α.

(3) π is ABORT(T) for some transaction T.
Then π is an input to transaction parent(T) = T'. The scheduler preconditions imply that α contains a REQUEST–CREATE(T), but no return operation for T.

(4) π is CREATE(T) and T is an access to basic object X.
By the scheduler preconditions, no CREATE(T) or ABORT(T) appears in α, but a REQUEST–CREATE(T) appears in α. Assume for the sake of deriving a contradiction that T' is a pending access in $\alpha|X$. Then Lemma 22 implies that T' is live in α. Also, Lemma 17 implies that parent(T) is live in α. Then Lemma 23 implies that one of T' or parent(T) is an ancestor of the other; since T and T' are both leaves of the transaction tree, the only possibility is that parent(T) is a proper ancestor of T'. Let T'' be the sibling of T which is an ancestor of T'. Then T'' is live in α, by Lemma 18. That is, there is a CREATE(T''), but no COMMIT for T'' in α. But this contradicts the scheduler preconditions for π. Therefore, there is no pending access in $\alpha|X$. ∎

3.6.4. Visibility and Serial Schedules

In this paragraph, we prove interesting lemmas about visibility in serial schedules.

Lemma 25: Let α be a serial schedule, and π an operation in α. Then transaction(π) is visible in α to some transaction which is live in α.

Proof: Let T = transaction(π). Since α is not empty, T_0 is live in α. Let T' be the least ancestor of T which is live in α. The proof is by induction on the distance from T' to T. If T = T', the result is trivial. So assume that T ≠ T'. Then COMMIT(T) is in α, and so T is visible to parent(T) in α. Lemma 13 implies that α contains a REQUEST–CREATE(T) operation, which occurs at parent(T). Then the inductive hypothesis implies that parent(T) is visible to T'. Then T is visible to T' by Lemma 8. ∎

Lemma 26:

1. Let α be a serial schedule, T a transaction and X an object. Then visible(α,T)|X is a prefix of $\alpha|X$.

2. Let α be a serial schedule, T a transaction and P a primitive. Then visible(α,T)|P is a prefix of $\alpha|P$.

Proof: 1. Let π and φ be operations in $\alpha|X$, with π preceding φ, and φ an operation in visible(α,T). Let α' be the prefix of α preceding φ. Let T' = transaction(φ) and T'' = transaction(π). Since φ is either a CREATE or a REQUEST–COMMIT for T', well-formedness

of α implies that T' is live in $\alpha'\varphi$. Thus, by Lemma 23, the only live transactions in $\alpha'\varphi$ are ancestors of T'. By Lemma 25, T'' is visible to an ancestor of T' in $\alpha'\varphi$, and hence in α. By Lemma 8, T'' is visible to T' in α. But T' is visible to T in α, by assumption. Lemma 8 then implies that T'' is visible to T in α, which gives the result.

2. Immediate from Lemma 11 and part 1. ∎

Lemma 27: Let α be a nonempty serial schedule. Let π be the last operation in α which is an output of the serial scheduler. Then transaction(π) sees everything in α.

Proof: Let T = transaction(π). We first show that T is live in α. Either π is a CREATE(T) or else it is a return operation for a child T' of T. In the latter case, Lemma 14 implies that CREATE(T) also occurs in α. Thus, in either case, CREATE(T) occurs in α. Now, if a return operation for T occurs in α, Lemma 21 implies that it follows π, which is impossible. Thus, no return operation for T occurs in α. It follows that T is live in α.

Then Lemma 23 implies that the only other transactions that are live in α must be ancestors or descendants of T. We claim that no proper descendants of T are live in α. So assume for the sake of obtaining a contradiction that U is a proper descendant of T which is live in α. Then U is a descendant of a child V of T, and V is live in α, by Lemma 18. Let α' be the prefix of α preceding π. There are three cases.

(1) π is CREATE(T).
Then Lemma 14 yields a contradiction.

(2) π is a COMMIT operation for T', a child of T.
Then T' ≠ V, since T' is not live in α. But T' and V are both live in α', which contradicts Lemma 23.

(3) π is an ABORT(T'), for child T' of T.
Then T' ≠ V, since T' is not live in α. But V is live in α'. But then the scheduler preconditions for π are not satisfied, a contradiction.

Thus, no descendants are live in α, so the only transactions that are live in α are ancestors of T. Now let φ be any operation in α. Lemma 25 implies that transaction(φ) is visible in α to some ancestor of T, and hence to T. ∎

Lemma 28: Let α be a serial schedule, and T a transaction. Then visible(α,T) is a serial schedule.

Proof: We proceed by induction on the length of α. The basis, length 0, is trivial. Let $\alpha = \alpha'\pi$, where π is a single operation. Fix transaction T, and let T' = transaction(π). If T' is not visible to T in α, then visible(α,T) = visible(α',T), and the result is true by inductive hypothesis. So assume that T' is visible to T in α.

If π is an output operation of a primitive P, then visible(α,T)|P is a prefix of α|P, by Lemma 26, and thus is a schedule of P. By the inductive hypothesis, visible(α',T) is a serial schedule. Also, visible(α,T) = visible(α',T)π by Lemma 12. Then Lemma 4 shows that visible(α,T) is a serial schedule.

On the other hand, if π is an output operation of the scheduler, then Lemma 27 implies that T' sees everything in α. But since T' is visible to T in α, it follows that T sees everything in α. Thus,

visible(α,T) = α, a serial schedule. ∎

3.6.5. Reordering and Combining Serial Schedules

In this paragraph, we describe ways in which serial schedules can be modified and combined to yield other serial schedules. These lemmas are used in the proof of the main theorem, in Section 7.

Lemma 29: Let α and α' be two equivalent serial schedules. If β is a sequence of serial operations such that $\alpha\beta$ is a serial schedule, then $\alpha'\beta$ is a serial schedule, and is equivalent to $\alpha\beta$.

Proof: Equivalence is trivial. The fact that $\alpha'\beta$ is a serial schedule follows because the preconditions of the serial scheduler depend only upon the presence of previous operations, not their order. ∎

The next lemma says that any serial schedule can be transformed by moving all the operations visible to any particular transaction to the beginning of the schedule, and the result is another serial schedule. This lemma can be thought of as describing a kind of "canonical form" for a serial schedule, with respect to a particular transaction.

Lemma 30: Let α be a serial schedule, and T any transaction. Let β = visible(α,T). Then $\beta(\alpha - \beta)$ is equivalent to α and is serial.

Proof: Let $\alpha' = \beta(\alpha - \beta)$. If P is any primitive, then Lemma 26 implies that $\beta|P$ is a prefix of $\alpha|P$. Thus, α' is equivalent to α.

To show that α' is serial, we proceed by induction on its prefixes. By Lemma 28, β is serial, so we can use β as the basis. Let $\gamma\pi$ be a prefix of α', where π is a serial operation in $\alpha - \beta$ and γ is a serial schedule. If π is an output operation of a primitive P, then $\gamma\pi|P$ is a prefix of $\alpha'|P$, = $\alpha|P$ by equivalence, which is a schedule of P. Then Lemma 4 shows that $\gamma\pi$ is a serial schedule. So assume that π is an output operation of the serial scheduler.

Let s be the state of the serial scheduler after γ. Let $\gamma'\pi$ be the prefix of α ending in π, and let s' be the state of the serial scheduler after γ'. Then π is enabled in s'. We must show that π is enabled in s. This suffices, by Lemma 4.

Since every operation in γ' is also in γ, it follows that each component set of s' is a subset of the corresponding set of s. There are three cases.

(1) π is CREATE(T") for some transaction T".
Then transaction(π) = T", and T" is not visible to T in α. Then T" ∈ create−requested(s') ⊆ create−requested(s). Also, it is easy to show that T" ∉ created(s). Now let U be in siblings(T") ∩ created(s). If U ∈ created(s'), then U ∈ returned(s') since π is enabled in s', ⊆ returned(s), as needed. So suppose that U ∉ created(s'). Then CREATE(U) occurs in β, so U is visible to T in α.

Since α contains both CREATE(T") and CREATE(U), Lemma 23 implies that α must contain a COMMIT for at least one of T" or U. If α contains a COMMIT for U, then it occurs in β, so U ∈ returned(s). On the other hand, if α contains a COMMIT for T", then T" is visible to U in α, so Lemma 8 implies that T" is visible to T in α, a contradiction.

(2) π is COMMIT(T",v) for some transaction T" and value v.

Then transaction(π) is parent(T"), which is not visible to T in α. Then (T",v) is in commit-requested(s') $\subseteq$ commit-requested(s). Also, it is easy to show that T" is not in returned(s). Now let U be in children(T") $\cap$ create-requested(s). Then there is a REQUEST-CREATE(U) in γ. This REQUEST-CREATE(U) occurs at T", which cannot be visible to T in α since parent(T") is not visible to T in α. Thus, the REQUEST-CREATE(U) does not occur in β, so it occurs in γ'. Since π is enabled in s', we have U $\in$ returned(s') $\subseteq$ returned(s).

(3) π is ABORT(T") for some transaction T".
Then transaction(π) = parent(T"), and parent(T") is not visible to T in α. Then T" $\in$ create-requested(s') $\subseteq$ create-requested(s). Also, it is easy to show that T" $\notin$ created(s). Now let U $\in$ siblings(T") $\cap$ created(s). Then CREATE(U) occurs in γ. But CREATE(U) occurs at U, and U cannot be visible to T in α since parent(U) = parent(T") is not visible to T in α. Therefore, CREATE(U) does not occur in β, so it occurs in γ'. Then U is in siblings(T") $\cap$ created(s') $\subseteq$ returned(s') $\subseteq$ returned(s). ∎

The following lemma is an easy consequence of the preceding one.

Lemma 31: Let α be a schedule of serial operations, and let T and T" be two transactions with T" visible to T in α. Let β and β' be serial schedules, such that β is equivalent to visible(α,T) and β' is equivalent to visible(α,T"). Then β" = β'(β - β') is equivalent to β and serial.

Proof: Let γ = visible(β,T"). Then γ is serial by Lemma 28. Lemma 30 implies that $\gamma(\beta - \gamma)$ is equivalent to β and serial. Lemma 10 implies that β' is equivalent to γ, and thus that $\beta - \gamma = \beta - \beta'$. Then Lemma 29 implies that β" is equivalent to $\gamma(\beta - \gamma)$ and serial. Thus, β" is equivalent to β and serial. ∎

The next two lemmas are used in the proof of Theorem 68. Each describes a way of "cutting and pasting" two serial schedules to yield a new serial schedule.

Lemma 32: Let $\alpha\beta_1$COMMIT(T",u) and $\alpha\beta_2$ be two serial schedules and T, T" and T"' three transactions such that the following conditions hold:
(1) T" is a child of T"' and T is a descendant of T"' but not of T",
(2) T" sees everything in $\alpha\beta_1$,
(3) T sees everything in $\alpha\beta_2$,
(4) α = visible($\alpha\beta_1$,T"') = visible($\alpha\beta_2$,T"') and
(5) no basic object has operations in both β_1 and β_2.
Then $\alpha\beta_1$COMMIT(T",u)β_2 is a serial schedule.

Proof: Note first that if T = T"', then β_2 is empty and the result is trivial. So assume that T $\neq$ T"'. Then T is a descendant of a child U of T"', and U $\neq$ T".

Any operation in $\alpha\beta_1$ whose transaction is not a descendant of T", must be in visible($\alpha\beta_1$,T") by Lemma 8. Similarly, any operation in $\alpha\beta_2$ whose transaction is not a descendant of U, must be in visible($\alpha\beta_2$,T"). Thus, β_1 and β_2 contain only operations at descendants of T" and U, respectively. Since T" and U are distinct siblings, and by assumption no objects have operations in both β_1 and β_2, it follows that no primitive has an operation occurring in both β_1 and β_2.

We proceed by induction on prefixes of $\alpha\beta_1$COMMIT(T",u)β_2. Let $\alpha'\varphi$ be a prefix of $\alpha\beta_1$COMMIT(T",u)β_2, with α' a serial schedule and φ a serial operation. We use $\alpha'\varphi$ = $\alpha\beta_1$COMMIT(T",u) as the basis, since $\alpha\beta_1$COMMIT(T",u) is a serial schedule by assumption. So

assume that $\alpha' = \alpha\beta_1 COMMIT(T'',u)\beta'$ for some sequence β'. There are two cases, depending on whether φ is an output of a primitive or of the serial scheduler.

Suppose that φ is an output operation of a primitive P. Then $\beta_1 COMMIT(T'',v)$ contains no operations at P. Thus, $\alpha'\varphi|P = \alpha\beta'\varphi|P$, which is a prefix of $\alpha\beta_2|P$, which is a schedule of P since $\alpha\beta_2$ is a serial schedule. Thus, $\alpha'\varphi|P$ is a schedule of P. The result follows by Lemma 4.

So suppose φ is an output of the serial scheduler. Then transaction(φ) = V for some descendant V of U. Let s be the state of the serial scheduler after α', and let s' be the state of the serial scheduler after $\alpha\beta'$. Then the following relationships hold between s and s'.

1. $V \in$ create−requested(s') - created(s') iff $V \in$ create−requested(s) - created(s)

2. children(V) $\cap$ create−requested(s') $\subseteq$ returned(s') iff children(V) $\cap$ create−requested(s) $\subseteq$ returned(s)

3. $(V,v) \in$ commit−requested(s') iff $(V,v) \in$ commit−requested(s)

4. $V \notin$ returned(s') iff $V \notin$ returned(s)

5. siblings(V) $\cap$ created(s') $\subseteq$ returned(s') iff siblings(V) $\cap$ created(s) $\subseteq$ returned(s)

Since the operations in β_1 are all at descendants of T'', and those of β_2 are all at descendants of U, the first four biconditionals are immediate from Lemma 7. If V is a proper descendant of U, the last biconditional also follows. It remains to show that siblings(U) $\cap$ created(s') $\subseteq$ returned(s') iff siblings(U) $\cap$ created(s) $\subseteq$ returned(s). But any sibling of U created in $\alpha\beta'$ is created in α', and the only sibling of U created in α' and not $\alpha\beta'$ is T'', and $T'' \in$ returned(s). Thus, the claims are true.

Since φ is enabled in s', the claims above imply that φ is also enabled in s. The result follows. ∎

Lemma 33: Let $\alpha ABORT(T'')$ and $\alpha\beta$ be two serial schedules, and let T, T'' and T''' be transactions, such that the following conditions hold:
(1) T'' is a child of T''' and T is a descendant of T''' but not of T'',
(2) T sees everything in $\alpha\beta$, and
(3) α = visible(α,T''') = visible($\alpha\beta$,T''').
Then $\alpha ABORT(T'')\beta$ is a serial schedule.

Proof: Similar to, but somewhat simpler than, the proof of Lemma 32. ∎

4. Resilient Objects

Having stated our correctness conditions, we are now ready to begin describing implementations and proving that they meet the requirements. This section and the next are devoted to the description of a concurrent system which permits the abort of transactions that have performed steps. An important component of a concurrent system is a new kind of object called a "resilient object," which we introduce in this section. A resilient object is similar to a basic object, but it has the additional capability to undo operations of transactions that it discovers have aborted.

Resilient objects have no capabilities for managing concurrency: rather, they assume that concurrency control is handled externally (by lock manager components of the scheduler). This section defines resilient objects and presents some of their properties. It also digresses slightly from the main development by describing and proving correct a particular implementation of resilient objects, which are constructed by keeping multiple copies of corresponding basic objects. The resilient object manages these copies as versions of the data object. Upon learning of an abort, the appropriate stored version is used in place of the current version.

4.1. Definitions

Resilient object R(X) mimics the behavior of basic object X, but has two additional input operations, INFORM−COMMIT−AT(X)OF(T) and INFORM−ABORT−AT(X)OF(T), for every transaction T. Upon receiving an INFORM−ABORT−AT(X)OF(T), R(X) erases any effects of accesses which are descendants of T. This property is made formal as the "Resiliency Condition" below.

R(X) has the following operations, which we call *R(X)-operations.*

Input Operations:
CREATE(T), T an access to X
INFORM−COMMIT−AT(X)OF(T)
INFORM−ABORT−AT(X)OF(T)

Output Operations:
REQUEST−COMMIT(T,v), T an access to X

In order to describe well-formedness for resilient objects, we require a technical definition for the set of transactions which are *active* after a sequence of R(X)-operations. Roughly speaking, the transactions which are active are those on whose behalf the object has carried out some activity, but whose fate the object does not know.

The definition is recursive on the length of the sequence of R(X) operations. Namely, only T_0 is active after the empty sequence. Let $\alpha = \beta\pi$, where π is a single operation, and let A and B denote the sets of active transactions after α and β, respectively. If π is CREATE(T), then A = B ∪ {T}. If π is a REQUEST−COMMIT for T, then A = B. If π is INFORM−COMMIT−AT(X)OF(T), and if T is in B, then A = (B - {T}) ∪ {parent(T)}; if T is not in B, then A = B. If π is INFORM−ABORT−AT(X)OF(T), then A = B - descendants(T).

Now we define *well-formedness* for sequences of R(X) operations. Again, the definition is recursive. Namely, the empty schedule is well-formed. Also, if $\alpha = \alpha'\pi$ is a sequence of R(X)-operations, then α is

well-formed provided that α' is well-formed, and the following hold.

- If π is CREATE(T), then
  (i) there is no CREATE(T) in α';
  (ii) all the transactions which are active after α' are ancestors of T.

- If π is a REQUEST−COMMIT for T, then
  (i) there is no REQUEST−COMMIT for T in α', and
  (ii) T is active after α'.

- If π is INFORM−COMMIT−AT(X)OF(T), then
  (i) there is no INFORM−ABORT−AT(X)OF(T) in α', and
  (ii) if T is an access to X, then a REQUEST−COMMIT for T occurs in α'.

- If π is INFORM−ABORT−AT(X)OF(T), then
  (i) there is no INFORM−COMMIT−AT(X)OF(T) in α'.

An immediate consequence of these definitions is that the transactions active after any well-formed sequence of R(X)-operations α are a subset of the ancestors of a single active transaction, which we denote least(α).

For α a sequence of R(X)-operations, define *undo(α)* recursively as follows. Define undo(λ) = λ, where λ is the empty sequence. Let $\alpha = \beta\pi$, where π is a single operation. If π is a serial operation (a CREATE or a REQUEST−COMMIT), then undo(α) = undo(β)π. If π is INFORM−COMMIT−AT(X)OF(T), then undo(α) = undo(β). If π is INFORM−ABORT−AT(X)OF(T), then undo(α) is the result of eliminating, from undo(β), all operations whose transactions are descendants of T. Note that undo(α) contains only serial operations.

Let α be any sequence of R(X)-operations, and let π be an operation in α of the form INFORM−ABORT−AT(X)OF(T). Then the *scope* of π in α is the subsequence γ of α consisting of operations eliminated by π.

Resiliency Condition

Resilient object R(X) *satisfies the resiliency condition* if for every well-formed schedule α of R(X), undo(α) is a schedule of basic object X.

We require that resilient object R(X) preserve well-formedness and satisfy the resiliency condition.

The resiliency condition is the correctness condition required by the concurrent schedulers at the object interface. The well-formedness requirement is a syntactic restriction, and the condition that undo(α) be a schedule of basic object X expresses the required semantic relationship between the resilient object and the

basic object it incorporates. The important property which must be preserved is that the correctness condition at the resilient objects, together with the behavior of the concurrent scheduler, assures correctness at the transaction boundaries.

4.2. Properties of Resilient Objects

This subsection contains a collection of simple lemmas about the properties of well-formed sequences of R(X) operations.

Lemma 34: Let $\alpha\pi$ be a well-formed sequence of R(X) operations, with π a single operation. The following are true.

1. If π is a serial operation, then transaction(π) is active after $\alpha\pi$.

2. If T is an access active after a prefix of α but not after α, then T is not active after $\alpha\pi$.

3. If π is a REQUEST−COMMIT for T, then CREATE(T) is the last serial operation in α.

Proof:

1. Immediate from the definition of active and well-formedness.

2. Because T has no descendants, it can only become active when a CREATE(T) operation occurs, which can only happen once in a well-formed schedule.

3. Suppose the last serial operation in α is φ, with $\varphi \neq$ CREATE(T). Let transaction(φ) = T'. By well-formedness, T ≠ T'. Also by well-formedness, T is active in α, so that CREATE(T) must occur in α, and so precedes φ. By part (1), T is active following CREATE(T) and after π, and T' is active following φ. But T cannot be active when φ occurs, by well-formedness, contradicting part (2) of this lemma.

∎

Lemma 35: Let α be a well-formed sequence of R(X) operations. Let T and T' be accesses to X, with T ≠ T', and let π and φ be serial operations with transactions T and T', respectively. If π precedes φ in α, then between π and φ, there is either an INFORM−ABORT−AT(X) for some ancestor of T, or else there are INFORM−COMMIT−AT(X)OF(U) operations for all ancestors U of T which are not ancestors of T', occurring in order from lowest to highest in the transaction tree ordering.

Proof: By part 3 of Lemma 34 and well-formedness, we may assume that φ = CREATE(T'). Lemma 34 implies that T is active immediately after π. By well-formedness, before CREATE(T') can occur, it must be that all transactions which are active are ancestors of T'. There are only two ways in which this can happen. One possibility is that R(X) first receives INFORM−COMMITS for all ancestors of T up to lca(T,T'), in order from lowest to highest in the transaction tree ordering. The other possibility is that R(X) first receives an INFORM−ABORT for an ancestor of T. ∎

Lemma 36: Let $\alpha\pi$ be a well-formed sequence of R(X) operations, with π = INFORM−ABORT−AT(X)OF(T). Then undo($\alpha\pi$) is a prefix of undo(α).

Proof: Suppose not. Then there is a subsequence $\varphi\psi$ of two operations in undo(α), such that ψ is in undo($\alpha\pi$) and φ is not. Clearly, φ and ψ are serial operations, transaction(φ) is a descendant

of T and transaction(ψ) is not. Since φ is not in the scope of an INFORM–ABORT in α, by Lemma 35, there is an INFORM–COMMIT between φ and ψ for every proper descendant of lca(transaction(φ),transaction(ψ)) that is an ancestor of transaction(φ), including T. This contradicts the well-formedness of $\alpha\pi$. ∎

Lemma 37: Let α be a well-formed sequence of R(X) operations, and let T be any transaction active in α, other than T_0. Then undo(α) contains an operation φ at a descendant T" of T, which is followed in α by an INFORM–COMMIT for every ancestor of T" which is a proper descendant of T.

Proof: The proof is by induction on α, with a trivial basis. Let $\alpha = \alpha'\pi$, such that the lemma is true for α' and that π is a single operation. Let T be a transaction active after α. There are four cases.

Suppose π is CREATE(T"). Then undo(α) = undo(α')π. If T ≠ T", the result is immediate by the induction hypothesis, since T is active after α'. If T = T", then the lemma follows, with $\pi = \varphi$.

If π is a REQUEST–COMMIT for a transaction T", then undo(α) = undo(α')π and the same transactions are active in α and α'. The result is immediate.

Suppose π is an INFORM–COMMIT for a transaction T". Then undo(α) = undo(α')π. If T is active after α', the result is immediate. If T is not active after α', it follows that T = parent(T"). The result is immediate from the induction hypothesis.

Suppose π is an INFORM–ABORT for a transaction U. Since T is active after α, it was active after α' and U is not an ancestor of T. Let φ be the transaction of transaction T" which follows from the inductive hypothesis applied to T and α'. Since α is well-formed and α' contains INFORM–COMMITs for every ancestor of T" up to T, U is not an ancestor of T". It follows that φ is in undo(α) and the result holds. ∎

Lemma 38: Let α be a well-formed sequence of R(X) operations, and let least(α) = T. If undo(α) is nonempty, then it ends in an operation of a descendant of T.

Proof: If T = T_0, the result is trivial, so assume otherwise. By the previous lemma, undo(α) contains an operation φ at a descendant of T. Without loss of generality, assume that φ is the last operation in undo(α) at a descendant of T. If any other operation π followed φ in undo(α), by Lemma 35 α would contain INFORM–COMMITs for every ancestor of transaction(φ) up to lca(transaction(φ),transaction(π)), which includes T. Then T is not active in α, a contradiction. ∎

Lemma 39: Let $\alpha\pi$ be a well-formed sequence of R(X) operations, with π = INFORM–ABORT–AT(X)OF(T). If T is not an ancestor of least(α), then undo($\alpha\pi$) = undo(α).

Proof: Suppose that T is not an ancestor of least(α) and that undo($\alpha\pi$) ≠ undo(α). Then undo(α) contains a serial operation φ at a descendant T" of T. By Lemma 38, φ is followed in undo(α) by an operation at a descendant of least(α). By Lemma 35, α contains an INFORM–COMMIT for every ancestor least(α) up to lca(least(α),T"), which includes T, contradicting the well-formedness of $\alpha\pi$.

We are now able to show that the undo operator preserves well-formedness.

Lemma 40: If α is a well-formed sequence of R(X)-operations, then undo(α) is a well-formed sequence of X-operations.

Proof: The proof is by induction on the length of α. The basis is trivial. Assume $\alpha = \alpha'\pi$, where π is a single operation, and undo(α') is a well-formed sequence of X-operations. If π is an INFORM-ABORT or INFORM−COMMIT, undo(α) is a prefix of undo(α'), by Lemma 36, and the result is immediate.

If π is CREATE(T), then undo(α) = undo(α')π. By the well-formedness of α, CREATE(T) does not appear in α', and so not in undo(α'). Hence, (i) is satisfied. To see (ii), assume that CREATE(T') occurs in undo(α'), for access T'. Then Lemma 35 implies that INFORM−COMMIT−AT(X)OF(T') occurs after CREATE(T') in α. Then well-formedness (the precondition for INFORM−COMMIT−AT(X)OF(T')) implies that a REQUEST−COMMIT for T' occurs in α', and well-formedness also implies that the REQUEST−COMMIT for T' follows the CREATE(T'). Therefore, the REQUEST−COMMIT occurs in undo(α'), and so T' is not pending in undo(α'). Thus, (ii) is satisfied.

If π is a REQUEST−COMMIT for T, then again undo(α) = undo(α')π, and by the well-formedness of α, (i) no REQUEST−COMMIT for T appears in α', and so not in undo(α'), and (ii) T is active after α', and it follows that CREATE(T') occurs in undo(α'). ∎

4.3. Construction of a Resilient Object

In this subsection, we describe a construction of a resilient object R(X) from a basic object X.

Recall that a resilient object X is distinguished from a basic object in that it has INFORM−ABORT and INFORM−COMMIT operations, a different definition of well-formedness, and satisfies the resiliency condition. The resilient object R(X) is constructed from the states, transition function and operation labels of the basic object X. The resilient object R(X) maintains a collection of "copies of X" (i.e. remembers states of X), one for each active transaction, with a particular current copy (corresponding to the least active transaction) to which CREATE operations are sent. When R(X) receives an INFORM−ABORT, the appropriate stored copy becomes the current copy, thereby erasing the effects of the operations in the scope of the INFORM−ABORT.

The state of R(X) consists of a pair (act,map), where act is a set of "active" transactions, and map is a function from act to states of basic object X. In the well-formed executions of R(X) (defined below), act will always be a subset of the set of ancestors of one particular transaction in act, called least(act). (In case act has no least member (which, again, will not arise in executions with well-formed schedules), define least(act) arbitrarily.) The copy for least(act) is considered to be current. The initial states of R(X) are those in which act = $\{T_0\}$ and map(T_0) is an initial state of the basic object X. In the following specification of the operations of R(X), let (act',map') be the state of R(X) prior to the operation, and (act,map) be the state of R(X) after the operation.

- CREATE(T), T an access to X:
  Postcondition:

act = act' ∪ {T}
map(U) = map'(U) for all U ∈ act - {T}
map(T) = s, where (map'(least(act')),CREATE(T),s) is in the transition relation of X

- INFORM−ABORT−AT(X)OF(T):
  Postcondition:
  act = act' - {descendants(T)}
  map(U) = map'(U) for all U ∈ act

- INFORM−COMMIT−AT(X)OF(T):
  Postcondition:
  if T ∈ act' then
  begin
  act = (act' - {T}) ∪ {parent(T)}
  map(U) = map'(U) for U ∈ act - {parent(T)}
  map(parent(T)) = map'(T)
  end
  if T ∉ act' then act = act' and map = map'

- REQUEST−COMMIT(T,v):
  Precondition:
  least(act') = T
  (map'(T),REQUEST−COMMIT(T,v),s) is in the transition relation of X
  Postcondition:
  act = act'
  map(U) = map'(U) for all U ∈ act - {T}
  map(T) = s

Now we prove that this implementation is a correct resilient object.

Lemma 41: Let α be a well-formed schedule of R(X) which can leave R(X) in state (act,map). Then act coincides with the set of transactions which are active after α.

Proof: The proof is by induction on the length of α. The basis is trivial. Let $\alpha = \alpha'\pi$, where π is a single operation. There are four cases, depending on the type of operation π. Each is immediate from the definition of active and the implementation of R(X). ∎

Lemma 42: Let α be a well-formed schedule of R(X) which can leave R(X) in state (act,map). Then the following conditions hold.

- undo(α) is a schedule of basic object X which can leave X in state map(least(act)), and
- if T" is any transaction other than T_0, and αINFORM−ABORT−AT(X)OF(T")) is well-formed, then undo(αINFORM−ABORT−AT(X)OF(T")) is a schedule of basic object X which can leave X in state map(U), where U is the least element of act which is not a descendant of T".

Proof: First, observe that if T" is not an ancestor of least(act), and αINFORM−ABORT−AT(X)OF(T") is well-formed, then Lemmas 41 and 39 imply that undo(αINFORM−ABORT−AT(X)OF(T")) = undo(α), so the second condition follows from the first.

The proof is by induction on the length of α. In each case, we prove the first condition, then the second condition assuming that T" is an ancestor of least(act). By the observation above, this is sufficient.

The basis is trivial. Let $\alpha = \alpha'\pi$, where π is a single operation. Let (act',map') be a state of R(X) after α', such that ((act',map'),π,(act,map)) is a transition for R(X). There are four cases.

1) π = CREATE(T)
Then undo(α) = undo(α')π. By the inductive assumption, undo(α') is a schedule of X which can leave X in state map'(least(act')). By the implementation of R(X), (map'(least(act')),π,map(T)) is a transition of X, and T = least(act). Thus the first condition of the lemma is satisfied.

To see that the second condition holds, note that all active transactions after α are ancestors of T, and by well-formedness, are exactly the transactions active after α', together with T. Let φ be INFORM−ABORT−AT(X)OF(T"), where T" is an ancestor of T other than T_0, and $\alpha\varphi$ is well-formed. If T" is a proper descendant of least(act'), by Lemma 39, undo($\alpha\varphi$) = undo(α'), which is a schedule of basic object X which can leave X in state map(least(act'))), by the inductive hypothesis. If T" is an ancestor of least(act'), undo($\alpha\varphi$) = undo($\alpha'\varphi$), the least element of act which is not a descendant of T" is also the least element of act' which is not a descendant of T", and the result follows by the inductive hypothesis.

2) π = REQUEST−COMMIT(T,v)
Then undo(α) = undo(α')π. By the inductive assumption, undo(α') is a schedule of X which can leave X in state map'(least(act')). By the implementation of R(X), (map'(least(act')),π,map(T)) is a transition of X, and T = least(act). Thus the first condition of the lemma is satisfied.

To see that the second condition holds, note that the active transactions after α are all ancestors of T, and by well-formedness, are exactly the transactions active after α'. Let φ be INFORM−ABORT−AT(X)OF(T"), where T" is an ancestor of T other than T_0, and $\alpha\varphi$ is well-formed. Then undo($\alpha\varphi$) = undo($\alpha'\varphi$), which is a schedule of basic object X which can leave X in state map(least(act'))), by the inductive hypothesis. Furthermore, the least element of act which is not a descendant of T" is also the least element of act' which is not a descendant of T", and the result follows by the inductive hypothesis.

3) π = INFORM−COMMIT−AT(X)OF(T)
Then undo(α) = undo(α'). Also, map(least(act)) = map(least(act')), by definition of R(X). The first condition follows.

By the definition of R(X), least(act) is an ancestor of least(act'). Let φ be INFORM−ABORT−AT(X)OF(T"), where T" is an ancestor of least(act) other than T_0, and $\alpha\varphi$ is well-formed. Then $\alpha'\varphi$ is well-formed, and undo($\alpha\varphi$) = undo($\alpha'\varphi$). Also, since $\alpha\varphi$ is well-formed, T" ≠ T. Let U and U' be the least elements of act and act', respectively, which are not descendants of T".

If T ∉ act', or if U ≠ parent(T), then U = U' and map(U) = map'(U'), and the second condition follows from the inductive hypothesis. So assume that T ∈ act' and U = parent(T). Then since T" ≠ T, it follows that U' = T. Then map'(U') = map(U), and the second condition again follows from the inductive hypothesis.

4) π = INFORM−ABORT−AT(X)OF(T)
If T is not an ancestor of least(act'), then undo(α) = undo(α'), by Lemma 39. Furthermore, the state of R(X) is not changed, αINFORM−ABORT−AT(X)−OF(T') is well-formed only if α'INFORM−ABORT−AT(X)−OF(T') is, and the active transactions after α are exactly those active after α'. The result follows.

Suppose that T is an ancestor of least(act'). The first condition is immediate from the inductive hypothesis. Let φ be INFORM−ABORT−AT(X)OF(T'), where T' is an ancestor of least(act) other than T_0, and $\alpha\varphi$ is well-formed. Since act = act' - descendants(T), least(act), and hence T', is an ancestor of T, undo($\alpha\varphi$) = undo($\alpha'\pi\varphi$) = undo($\alpha'\varphi$), and the second condition follows as well. ∎

Theorem 43: R(X) is a resilient object.

Proof: We must show that R(X) preserves well-formedness and satisfies the resiliency condition. That R(X) satisfies the resiliency condition follows immediately from Lemma 42.

Assume that α is a well-formed schedule of R(X) and π is an output operation of R(X) enabled after an execution with schedule α. We must show that $\alpha\pi$ is a well-formed sequence of R(X)-operations.

Since π is an output, it has the form REQUEST−COMMIT(T,v) for some access T and value v. Let (act,map) be a state of R(X) after α, such that π is enabled in (act,map). Clearly, π is an output of basic object X enabled from state map(least(act)). By Lemma 42, undo(α) is a schedule of basic object X which can leave X in state map(least(act))), so undo(α)π = undo($\alpha\pi$) is a schedule of basic object X.

Since X preserves well-formedness for basic objects, and by Lemma 40 undo(α) is a well-formed sequence of X-operations, undo(α) ends with the operation φ = CREATE(T) and contains no other operations with transaction T. Let $\beta\varphi$ be the prefix of α ending in φ. Suppose first that a REQUEST−COMMIT for T occurs in α. Since α is well-formed, φ is the only CREATE(T) operation in α, and by Lemma 34, the second REQUEST−CREATE for T follows φ, and by the definition of undo, is in undo(α) if φ is, a contradiction.

It remains to show that T is active after α. By Lemma 34, T is active after $\beta\varphi$. No INFORM−COMMIT for T can occur after φ in α, since by well-formedness, there is no REQUEST−COMMIT for T in α. Also, since φ is in undo(α), no INFORM−ABORT for an ancestor of T can occur after φ in α. Thus T is still active after α. ∎

5. Concurrent Systems

As with serial schedules in classical settings, our serial schedules contain no concurrency or resiliency and thus are too inefficient to use in practice. Their importance is solely for defining correctness for transaction systems. In this section, we define a new kind of system called a *concurrent system*. The new system consists of the same transactions as in a serial system, a resilient object R(X) for every basic object X of the serial system, and a concurrent scheduler.

Concurrent systems describe computations in which transactions run concurrently and can be aborted after

they have performed some work. The concurrent scheduler has the joint responsibility of controlling concurrency and of seeing that the effects of aborted transactions (and their descendants) become undone. Concurrent systems make use of the roll-back capabilities of resilient objects to make sure that ABORT operations in concurrent systems have the same semantics (so far as the transactions can tell) as they do in serial systems.

Concurrent systems are defined in this section. In the next section, the more permissive "weak concurrent systems" are defined. In Section 7, we prove that the schedules of concurrent systems are serially correct, as a corollary of a weaker correctness property for the weak concurrent system.

5.1. Lock Managers

The scheduler we define is called the *concurrent scheduler.* It is composed of several automata: a *lock manager* for every object X, and a single *concurrent controller.* The job of the lock managers is to insure that the associated object receives no CREATES until the lock manager has received abort or commit information for all necessary preceding transactions. This lock manager models an exclusive locking protocol derived from Moss' algorithm [Mo]. The lock manager has the following operations.

Input Operations:
INTERNAL−CREATE(T), where T is an access to X
INFORM−COMMIT−AT(X)OF(T), for T any transaction
INFORM−ABORT−AT(X)OF(T), for T any transaction

Output Operations:
CREATE(T), where T is an access to X

The input operations INTERNAL−CREATE, INFORM−COMMIT and INFORM−ABORT will compose with corresponding output operations of the concurrent scheduler which we will construct in this subsection. The output CREATE operation composes with the CREATE input operation of the resilient object R(X). The lock manager receives and manages requests to access object X, using a hierarchical locking scheme. It uses information about the commit and abort of transactions to decide when to release locks.

Each state s of the lock manager consists of the following three sets of transactions: lock−holders(s), create−requested(s), and created(s). Initially, lock−holders = $\{T_0\}$, and the other sets are empty. The operations work as follows.

- INTERNAL−CREATE(T)
  Postcondition:
  create−requested(s) = create−requested(s') ∪ {T}

- INFORM−COMMIT−AT(X)OF(T)

Postcondition:
if T ∈ lock−holders(s') then lock−holders(s) = (lock−holders(s') - {T}) ∪ {parent(T)}

- INFORM−ABORT−AT(X)OF(T)
  Postcondition:
  lock−holders(s) = lock−holders(s') - descendants(T)

- CREATE(T)
  Precondition:
  T ∈ create−requested(s') - created(s')
  lock−holders(s') ⊆ ancestors(T)
  Postcondition:
  lock−holders(s) = lock−holders(s') ∪ {T}
  created(s) = created(s') ∪ {T}

Note that resilient object R(X) and the lock manager for X share the INFORM−ABORT and INFORM−COMMIT input operations. These compose with the output from the concurrent controller defined below.

Thus, the lock manager only sends a CREATE(T) operation on to the object in case all the current lock−holders are ancestors of T. When the lock manager learns about the commit of a transaction T for which it holds a lock, it releases the lock to T's parent. When the lock manager learns about the abort of a transaction T for which it holds a lock, it simply releases all locks held by that transaction and its descendants. Our model provides an exceptionally simple and clear way of describing this important algorithm.

A key property of lock managers is described by the following lemma.

Lemma 44: Let X be an object and let T and T' be accesses to X. Let U be an ancestor of T which is not an ancestor of T'. Let α be a schedule of the lock manager for X. If CREATE(T) precedes CREATE(T') in α, then between the two CREATE operations, there is either an INFORM−COMMIT−AT(X)OF(U) operation, or else an INFORM−ABORT−AT(X) for some ancestor of T.

Proof: At the time the CREATE(T) occurs, the lock manager puts T into the set of lock−holders. Before the lock manager can send in CREATE(T'), it must be that all the transactions in lock−holders are ancestors of T'. There are only two ways in which this can happen. One possibility is that the lock manager first receives INFORM−COMMITS for all ancestors of T up to lca(T,T'), including INFORM−COMMIT−AT(X)OF(U). The other possibility is that the lock manager first receives an INFORM−ABORT for an ancestor of T. ∎

5.2. The Concurrent Controller

The concurrent controller is similar to the serial scheduler, but it allows siblings to proceed concurrently. In order to manage this properly, it interacts with "concurrent objects" (lock managers and resilient objects) instead of just basic objects. The operations are as follows.

Input Operations:
REQUEST–CREATE(T)
REQUEST–COMMIT(T,v)

Output Operations:
CREATE(T), T a non-access transaction
INTERNAL–CREATE(T), T an access transaction
COMMIT(T,v)
ABORT(T)
INFORM–COMMIT–AT(X)OF(T)
INFORM–ABORT–AT(X)OF(T)

Each state s of the concurrent controller consists of five sets: create–requested(s), created(s), commit–requested(s), committed(s), and aborted(s). The set commit–requested(s) is a set of (transaction,value) pairs, and the others are sets of transactions. (As before, we will occasionally write $T \in$ commit–requested(s) for $(T,v) \in$ commit–requested(s) for some v.) All sets are initially empty except for create–requested, which is $\{T_0\}$. Define returned(s) = committed(s) $\cup$ aborted(s). The operations are as follows.

- REQUEST–CREATE(T)
  Postcondition:
  create–requested(s) = create–requested(s') $\cup$ {T}

- REQUEST–COMMIT(T,v)
  Postcondition:
  commit–requested(s) = commit–requested(s') $\cup$ {(T,v)}

- CREATE(T), T a non-access transaction
  Precondition:
  $T \in$ create–requested(s') - created(s')
  Postcondition:
  created(s) = created(s') $\cup$ {T}

- INTERNAL–CREATE(T), T an access transaction
  Precondition:
  $T \in$ create–requested(s') - created(s')
  Postcondition:
  created(s) = created(s') $\cup$ {T}

- COMMIT(T,v)
  Precondition:
  $(T,v) \in$ commit–requested(s')
  $T \notin$ returned(s')
  children(T) $\cap$ create–requested(s') $\subseteq$ returned(s')
  Postcondition:
  committed(s) = committed(s') $\cup$ {T}

- ABORT(T)
  Precondition:
  T ∈ (create-requested(s') - created(s')) ∪ (commit−requested(s') - returned(s'))
  children(T) ∩ create−requested(s') ⊆ returned(s')
  Postcondition:
  created(s) = created(s') ∪ {T}
  aborted(s) = aborted(s') ∪ {T}

- INFORM−COMMIT−AT(X)OF(T):
  Precondition:
  T ∈ committed(s')

- INFORM−ABORT−AT(X)OF(T):
  Precondition:
  T ∈ aborted(s')

The concurrent controller is closely related to the serial scheduler. In place of the serial scheduler's CREATE operations, the concurrent controller has two kinds of operations, CREATE operations and INTERNAL−CREATE operations. The former is used for interaction with non-access transactions, while the latter is used for interaction with access transactions. From the concurrent controller's viewpoint, the two operations are the same; however, our naming convention for operations requires us to assign them different names, since the INTERNAL−CREATE operations are intended to be identified with INTERNAL−CREATE operations of the lock managers (which also have CREATE operations, for interaction with the resilient objects). The precondition on the serial scheduler's CREATE operation which insures serial processing of sibling transactions, does not appear in the concurrent controller. Thus, the concurrent controller may run any number of sibling transactions concurrently, provided their parent has requested their creation.

The concurrent controller's COMMIT operation is the same as the serial scheduler's COMMIT operation (except for a minor difference in bookkeeping). The concurrent controller's ABORT operation is different, however; in addition to aborting a transaction in the way that the serial scheduler does, the concurrent controller has the additional capability to abort a transaction that has actually been created and has carried out some steps. In this particular formulation, aborts occur if the transaction was not created (as with the serial scheduler), or if the transaction has previously requested to commit, and its children have returned. Together with the requirements on the COMMIT operation, this condition insures that all transaction completion occurs bottom-up. In the weak concurrent system to be considered in Section 6, a different, "weak", concurrent controller will be used; it differs from the concurrent controller of this section precisely in not requiring ABORT operations to wait for their transactions (and subtransactions) to complete.

The concurrent controller also has two additional operations not present in the serial scheduler. These

operations allow the concurrent controller to forward necessary abort and commit information to the lock managers and resilient objects.

Lemma 45: Let α be a schedule of the concurrent scheduler, and let s be a state which can result from applying α to the initial state. Then the following conditions are true.

1. T is in create−requested(s) exactly if $T = T_0$ or α contains a REQUEST−CREATE(T) operation.

2. If T is a non-access transaction, then T is in created(s) exactly if α contains either a CREATE(T) or ABORT(T) operation.

3. If T is an access transaction, then T is in created(s) exactly if α contains either an INTERNAL−CREATE(T) or ABORT(T) operation.

4. (T,v) is in commit−requested(s) exactly if α contains a COMMIT−REQUEST(T,v) operation.

5. (T,v) is in committed(s) exactly if α contains a COMMIT(T,v) operation.

6. T is in aborted(s) exactly if α contains an ABORT(T) operation.

5.3. Concurrent Systems

The composition of transactions, resilient objects and the concurrent scheduler (lock managers and concurrent controller) is the *concurrent system*. A schedule of the concurrent system is a *concurrent schedule*, and the operations of a concurrent system are *concurrent operations*.

A sequence α of concurrent operations is *well-formed* if for every primitive P, $\alpha|P$ is well-formed (using the appropriate definition of well-formedness).

The main result of this paper is that every concurrent schedule is serially correct. This will be proved as a corollary of a stronger result, in Section 7.

5.4. Properties of Concurrent Systems

As we did for serial schedules, we now prove some useful basic properties for concurrent schedules. These lemmas describe the possible kinds and orders of operations that can occur in well-formed concurrent schedules. Later, we will see that all concurrent schedules are well-formed, so these properties actually follow just from the fact that these schedules are concurrent. All results and proofs in this subsection are straightforward.

Lemma 46: Let α be a well-formed concurrent schedule, and let $T \neq T_0$ be a transaction.

1. If α contains any operation with transaction T, then α contains a CREATE(T) and a REQUEST−CREATE(T).

2. If α contains a COMMIT for T, then α contains a REQUEST−COMMIT for T, a CREATE(T) and a REQUEST−CREATE(T).

3. If α contains an ABORT(T), then α contains a REQUEST−CREATE(T).

Lemma 47: Let α be a well-formed concurrent schedule, and T a transaction. Assume that some descendant of T is in transaction(α). Then the following hold.

1. CREATE(T) occurs in α.

2. If $T \neq T_0$, then REQUEST−CREATE(T) occurs in α.

Lemma 48: Let α be a well-formed concurrent schedule, and let $T \neq T_0$ be a transaction.

1. If α contains a REQUEST−CREATE(T), but does not contain a return operation for T, then parent(T) is live in α.

2. If T is live in α, then parent(T) is live in α.

3. If α contains a REQUEST−CREATE(T) but does not contain a CREATE(T) or ABORT(T), then parent(T) is live in α.

Proof: 1. Well-formedness implies that the REQUEST−CREATE(T) is preceded by a CREATE(parent(T)). Suppose that parent(T) is not live in α. Then a return operation for parent(T) occurs in α. In case the return operation for parent(T) is an ABORT(parent(T)), scheduler preconditions imply that the CREATE(parent(T)) must precede the ABORT(parent(T)). Then the scheduler preconditions for the return operation imply that the return for parent(T) must be preceded by a REQUEST−COMMIT for parent(T). By well-formedness, the REQUEST−COMMIT for parent(T) must follow the REQUEST−CREATE(T), so that the return for parent(T) must follow the REQUEST−CREATE(T) Then the scheduler preconditions for the return operation imply that there must be a return operation for T in α, a contradiction.

2. and 3. are as in Lemma 17. ∎

Lemma 49: Let α be a well-formed concurrent schedule, and let T be a transaction.

1. If α contains a REQUEST−CREATE(T), but does not contain a return operation for T, then all proper ancestors of T are live in α.

2. If T is live in α, then any ancestor of T is live in α.

3. If α contains a REQUEST−CREATE(T) but does not contain a CREATE(T) or ABORT(T), then all proper ancestors of T are live in α.

Lemma 50: Let α be a well-formed concurrent schedule, and let T and T' be transactions with T' a descendant of T. Assume that there is a return operation for T in α.

1. If there is a REQUEST−CREATE(T') in α, then there is a return operation for T' in α.

2. If T' is in transaction(α), then there is a return operation for T' in α.

Proof:

1. By Lemma 49.

2. By Lemma 46 and part 1.

∎

Lemma 51: Let α be a well-formed concurrent schedule. If a return operation for T is in α, then it follows all operations in α whose transaction is T.

Proof: First consider the case where T is not an access. If no CREATE(T) occurs in α, the result is immediate, so assume that CREATE(T) occurs in α. In case an ABORT(T) occurs in α, scheduler preconditions imply that the CREATE(T) must precede the ABORT(T). Then the return operation for T must be preceded by a REQUEST–COMMIT for T, by scheduler preconditions. Well-formedness implies that the REQUEST–COMMIT is preceded by CREATE(T), and is not followed by any output operations of T. Thus, the only serial operations of T that could follow the REQUEST–COMMIT are return operations of children of T.

Let T' be a child of T for which a return operation occurs in α. By scheduler preconditions, there is only one return operation for T' in α. By Lemma 46, α also contains REQUEST–CREATE(T'). Since this is an output operation of T, it precedes the REQUEST–COMMIT for T, and hence precedes the return operation for T. Then the scheduler preconditions imply that the return operation for T' precedes the return for T.

Next, consider the case where T is an access. If no INTERNAL–CREATE(T) occurs in α, the result is immediate, so assume that INTERNAL–CREATE(T) occurs in α. In case an ABORT(T) occurs in α, scheduler preconditions imply that the INTERNAL–CREATE(T) must precede the ABORT(T). Then the return operation for T must be preceded by a REQUEST–COMMIT for T, and well-formedness implies that this is in turn preceded by CREATE(T). Thus, no serial operations of T can follow the return operation for T. ∎

Lemma 52: Let α be a well-formed concurrent schedule. If a return operation for T is in α, then it follows all operations in α whose transactions are descendants of T.

Proof: Since a return operation for T occurs in α, we have $T \neq T_0$. Let T' be a descendant of T, and assume for the sake of obtaining a contradiction that a serial operation π with transaction(π) = T' occurs after the return for T in α. Let α' be the prefix of α preceding π.

By Lemma 46, α' contains a REQUEST–CREATE(T'). Then Lemma 50 implies that α' must contain a return operation for T'. But then the well-formed schedule $\alpha'\pi$ contains a return operation for T' followed by an operation of T', which contradicts Lemma 51. ∎

Weak concurrent systems are defined in the following section, and many of their properties are stated and proved. Weak concurrent systems are obtained by replacing the concurrent scheduler with a more permissive scheduler, the weak concurrent scheduler. Results in Section 7 prove that every execution of the concurrent system is also an execution of the weak concurrent system. Thus, additional interesting properties of concurrent system behavior follow immediately from the corresponding properties of weak concurrent system behavior, proven in that section.

6. Weak Concurrent Systems

In this section, we define "weak concurrent systems", which are exactly the same as concurrent systems, except that they have a more permissive controller, the "weak concurrent controller". The weak concurrent controller reports aborts to a transaction's parent while there is still activity going on in the aborted transaction's subtree. In this paper, weak concurrent systems are used primarily to provide an intermediate step in proving the correctness of concurrent systems: proving a weaker condition for weak concurrent systems allows us to infer the stronger correctness condition for concurrent systems. However, weak concurrent systems are also of interest in themselves. In a distributed implementation of a nested transaction system, performance considerations may make it important for the system to allow a transaction to abort without waiting for activity in the transaction's subtree to subside. In this case, a weak concurrent system might be an appropriate choice, even though the correctness conditions which they satisfy are weaker. Weak concurrent systems also appears to have further technical use, for example in providing simple explanations of the ideas used in "orphan detection" algorithms [HLMW].

6.1. The Weak Concurrent Controller

In this subsection, we define the weak concurrent controller. As we have already said, it is identical to the concurrent controller except that it has a more permissive ABORT operation. For convenience, we describe the controller here in its entirety. It has the same operations as the concurrent controller:

Input Operations:
REQUEST−CREATE(T)
REQUEST−COMMIT(T,v)

Output Operations:
CREATE(T), T a non-access transaction
INTERNAL−CREATE(T), T an access transaction
COMMIT(T,v)
ABORT(T)
INFORM−COMMIT−AT(X)OF(T)
INFORM−ABORT−AT(X)OF(T)

Each state s of the concurrent controller consists of five sets: create−requested(s), created(s), commit−requested(s), committed(s), and aborted(s). The set commit−requested(s) is a set of (transaction,value) pairs, and the others are sets of transactions. (As before, we will occasionally write $T \in$ commit−requested(s) for $(T,v) \in$ commit−requested(s) for some v.) All are empty initially except for create−requested, which is $\{T_0\}$. Define returned(s) = committed(s) $\cup$ aborted(s). The operations are as follows.

- REQUEST−CREATE(T)
  Postcondition:

create−requested(s) = create−requested(s') ∪ {T}

- REQUEST−COMMIT(T,v)
  Postcondition:
  commit−requested(s) = commit−requested(s') ∪ {(T,v)}

- CREATE(T), T a non-access transaction
  Precondition:
  T ∈ create−requested(s') - created(s')
  Postcondition:
  created(s) = created(s') ∪ {T}

- INTERNAL−CREATE(T), T an access transaction
  Precondition:
  T ∈ create−requested(s') - created(s')
  Postcondition:
  created(s) = created(s') ∪ {T}

- COMMIT(T,v)
  Precondition:
  (T,v) ∈ commit−requested(s')
  T ∉ returned(s')
  children(T) ∩ create−requested(s') ⊆ returned(s')
  Postcondition:
  committed(s) = committed(s') ∪ {T}

- ABORT(T)
  Precondition:
  T ∈ create-requested(s') - returned(s')
  Postcondition:
  created(s) = created(s') ∪ {T}
  aborted(s) = aborted(s') ∪ {T}

- INFORM−COMMIT−AT(X)OF(T):
  Precondition:
  T ∈ committed(s')

- INFORM−ABORT−AT(X)OF(T):
  Precondition:
  T ∈ aborted(s')

Thus, the weak concurrent controller is permitted to abort any transaction that has had its creation requested, and which has not yet returned.

Lemma 53: Let α be a schedule of the concurrent scheduler, and let s be a state which can result from applying α to the initial state. Then the following conditions are true.

1. T is in create−requested(s) exactly if $T = T_0$ or α contains a REQUEST−CREATE(T) operation.

2. If T is a non-access transaction, then T is in created(s) exactly if α contains either a CREATE(T) or ABORT(T) operation.

3. If T is an access transaction, then T is in created(s) exactly if α contains either an INTERNAL–CREATE(T) or ABORT(T) operation.

4. (T,v) is in commit–requested(s) exactly if α contains a COMMIT–REQUEST(T,v) operation.

5. (T,v) is in committed(s) exactly if α contains a COMMIT(T,v) operation.

6. T is in aborted(s) exactly if α contains an ABORT(T) operation.

6.2. Weak Concurrent Systems

The composition of transactions, resilient objects and the weak concurrent scheduler (lock managers and weak concurrent controller) is the *weak concurrent system.* A schedule of the weak concurrent system is a *weak concurrent schedule.*

Weak concurrent systems exhibit nice behavior to transactions except possibly to those which are descendants of aborted transactions. Thus, we say that a transaction T is an *orphan* in any sequence α of operations provided that an ancestor of T is aborted in α. In many of the properties we prove for weak concurrent systems, we will have to specify that the transactions involved are not orphans.

6.3. Properties of Weak Concurrent Systems

As we did for serial and concurrent schedules, we here prove a number of useful basic properties for weak concurrent schedules. As before, most of these properties are simple to state and prove.

6.3.1. Operations in Weak Concurrent Schedules

As before, we include a collection of lemmas describing the possible kinds and orders of operations that can occur in well-formed weak concurrent schedules. These lemmas are analogous to some in Section 5, and have similar proofs; the main difference is that we must take proper care with orphans. As before, we go on to show that all weak concurrent schedules are well-formed, so these properties actually follow just from the fact that these schedules are weak concurrent.

Lemma 54: Let α be a well-formed weak concurrent schedule, and let $T \neq T_0$ be a transaction.

1. If α contains any operation with transaction T, then α contains a CREATE(T), and a REQUEST–CREATE(T).

2. If α contains a COMMIT for T, then α contains a REQUEST–COMMIT for T, a CREATE(T) and a REQUEST–CREATE(T).

3. If α contains an ABORT(T), then α contains a REQUEST–CREATE(T).

Lemma 55: Let α be a well-formed weak concurrent schedule, and T a transaction. Assume that some descendant of T is in transaction(α). Then the following hold.

1. CREATE(T) occurs in α.

2. If $T \neq T_0$, then REQUEST–CREATE(T) occurs in α.

Lemma 56: Let α be a well-formed weak concurrent schedule, and let $T \neq T_0$.

1. If α contains a REQUEST–CREATE(T), but does not contain a return operation for T, then parent(T) is not committed in α.

2. If T is live in α, then parent(T) is not committed in α.

3. If α contains a REQUEST–CREATE(T) but does not contain a CREATE(T) or ABORT(T), then parent(T) is not committed in α.

Proof: 1. Suppose a COMMIT operation for parent(T) occurs in α. Then the weak concurrent controller preconditions for the COMMIT operation imply that the COMMIT for parent(T) must be preceded by a REQUEST–COMMIT for parent(T). By well-formedness, the REQUEST–COMMIT for parent(T) must follow the REQUEST–CREATE(T), so that the COMMIT for parent(T) must follow the REQUEST–CREATE(T). Then the weak concurrent controller preconditions for the COMMIT operation imply that there must be a COMMIT operation for T in α, a contradiction.

2. and 3. are as in 3.6.2. ∎

Lemma 57: Let α be a well-formed weak concurrent schedule, and let T be a transaction which is not an orphan in α.

1. If α contains a REQUEST–CREATE(T), but does not contain a COMMIT operation for T, then all proper ancestors of T are live in α.

2. If T is live in α, then all proper ancestors of T are live in α.

3. If α contains a REQUEST–CREATE(T) but does not contain a CREATE(T), then all proper ancestors of T are live in α.

Proof: By repeated use of the previous lemma, well-formedness and the weak concurrent controller preconditions. ∎

Lemma 58: Let α be a well-formed weak concurrent schedule, and let T and T' be transactions with T' a descendant of T. Assume that T' is not an orphan in α and that there is a COMMIT operation for T in α.

1. If there is a REQUEST–CREATE(T') in α, then there is a COMMIT operation for T' in α.

2. If T' is in transaction(α), then there is a COMMIT operation for T' in α.

Proof:

1. By Lemma 57.

2. By Lemma 54 and part 1.

∎

6.3.2. Objects and Locking

In this paragraph, we give two simple lemmas about the behavior of the locking strategy.

Lemma 59: Let α be a weak concurrent schedule. Let X be an object, and let T and T' be accesses to X. Let U be an ancestor of T which is not an ancestor of T'. Assume that CREATE(T) precedes CREATE(T') in α.

1. There is either an INFORM_COMMIT_AT(X)OF(U), or else an INFORM_ABORT_AT(X) for some ancestor of T, occurring between CREATE(T) and CREATE(T') in α.

2. Either CREATE(T') is preceded by a COMMIT operation for U, and by a REQUEST_COMMIT operation for U, or else CREATE(T') is preceded by an ABORT operation for some ancestor of T.

Proof:

1. By Lemma 44.

2. By part 1 and the preconditions of the weak concurrent controller.

∎

Lemma 60: Let α be a well-formed weak concurrent schedule, and X a basic object. Then the set of active transactions after $\alpha|R(X)$ is exactly the set of lockholders in the lock manager for X after α.

Proof: By induction on the length of α. ∎

6.3.3. Well-Formedness

Here, we show that every weak concurrent schedule is well-formed. It follows that all the properties proved earlier in this section are actually true for all weak concurrent schedules. From now on, we will use these properties without explicitly mentioning well-formedness.

Lemma 61: Let α be a weak concurrent schedule. Then α is well-formed.

Proof: By induction on the length of schedules. The base, length = 0, is trivial. Suppose that $\alpha\pi$ is a weak concurrent schedule, where π is a single operation, and assume that α is well-formed. If π is an output of a primitive P, then the result is immediate, since each primitive preserves well-formedness. No INTERNAL_CREATE operation can cause a violation. So assume that π is an input to a primitive P. It suffices to show that $\alpha\pi|P$ is well-formed. There are six cases.

(1) π is CREATE(T) and T is a non-access transaction.
The controller preconditions insure that CREATE(T) does not appear in α.

(2) π is CREATE(T) and T is an access to resilient object R(X).
By the lock manager preconditions, no CREATE(T) appears in α. The lock manager preconditions and Lemma 60 imply that all the transactions which are active after α are ancestors

of T.

(3) π is COMMIT(T,v).
Then π is an input to transaction parent(T). Weak concurrent controller preconditions imply that α contains REQUEST−COMMIT(T,v), and so Lemma 54 implies that α contains REQUEST−CREATE(T). Also, weak concurrent controller preconditions insure that α does not contain a return operation for T.

(4) π is ABORT(T).
Then π is an input to transaction parent(T). Weak concurrent controller preconditions imply that α contains a REQUEST−CREATE(T). Weak concurrent controller preconditions insure that α does not contain a return operation for T.

(5) π is INFORM−COMMIT−AT(X)OF(T) at resilient object R(X).
By the preconditions of the weak controller, α contains a COMMIT for T. If INFORM−ABORT−AT(X)OF(T) occurs in α, then α also contains an ABORT for T, which contradicts weak concurrent controller preconditions. Thus, no INFORM−ABORT−AT(X)OF(T) occurs in α. Since a COMMIT for T occurs in α, weak concurrent controller preconditions imply that a REQUEST−COMMIT for T also occurs in α.

(6) π is INFORM−ABORT−AT(X)OF(T) at resilient object R(X).
By the preconditions of the weak concurrent controller, α contains ABORT(T). If INFORM−COMMIT−AT(X)OF(T) occurs in α, then α contains a COMMIT for T, which contradicts weak concurrent controller preconditions. Thus, no INFORM−COMMIT−AT(X)OF(T) occurs in α. ∎

6.3.4. Visibility and Weak Concurrent Schedules

This paragraph states and proves important properties involving visibility in weak concurrent schedules. In particular, the most important result of this paragraph is Lemma 66, which relates the portion of a weak concurrent schedule which is visible to a particular transaction, to schedules of transactions and basic objects. The first lemma shows how visibility propagates among the transactions during a weak concurrent execution.

Lemma 62: Let $\alpha\pi$ be a weak concurrent schedule, where π is a single operation.

1. If π is CREATE(T), then visible($\alpha\pi$,T) = visible(α,parent(T))π.

2. If π is COMMIT(T,v), then visible($\alpha\pi$,parent(T)) = visible(α,T)π.

3. If π is ABORT(T), then visible($\alpha\pi$,parent(T)) = visible(α,parent(T))π.

4. If π is COMMIT(T,v), and T' is a descendant of parent(T) but not T, then visible($\alpha\pi$,T') - visible($\alpha\pi$,parent(T)) = visible(α,T') - visible(α,T).

Proof: 1. By Lemma 55, π is the first serial operation in $\alpha\pi$ whose transaction is a descendant of T, and T is not visible to parent(T). Thus any transaction other than T visible to T in $\alpha\pi$ is visible to parent(T) in $\alpha\pi$. Then parent(T) is visible to T in $\alpha\pi$, and by Lemma 8, visible($\alpha\pi$,parent(T))π = visible($\alpha\pi$,T).

By the definition of visibility, any transaction visible to parent(T) in $\alpha\pi$ is visible to parent(T) in α, and visible(α,parent(T)) = visible($\alpha\pi$,parent(T)). Substituting in the equality above, we have the result.

2. By the definition of visibility, any transaction visible to parent(T) in $\alpha\pi$ is either visible to parent(T) in α, or is visible to T in α. But any transaction visible to parent(T) in α is visible to T in α, so we have that any transaction visible to parent(T) in $\alpha\pi$ is visible to T in α, and visible($\alpha\pi$,parent(T)) is a subsequence of visible(α,T)π. It follows immediately from the definition of visibility that any transaction visible to T in α is visible to parent(T) in $\alpha\pi$, so that visible(α,T) is a subsequence of visible($\alpha\pi$,parent(T)). The result is immediate.

3. Immediate from the definition of visibility.

4. Clearly, visible(α,T') is a subsequence of visible($\alpha\pi$,T'). Any operation in visible($\alpha\pi$,T') - visible(α,T') has a transaction which is a descendant of T, and so is either π or is visible to T in α, and thus is in visible(α,T)π. Thus we have visible($\alpha\pi$,T') - visible(α,T)π = visible(α,T') - visible(α,T)π. As π is not in visible(α,T'), this equals visible(α,T') - visible(α,T). By part 2, visible($\alpha\pi$,parent(T)) = visible(α,T)π, and the result follows by substitution in the first identity. ∎

Now we prove two lemmas involving visibility and the behavior of resilient objects in weak concurrent systems.

Lemma 63: Let α be a weak concurrent schedule. Let R(X) be a resilient object, and let T and T' be accesses to R(X). If T' is live and not an orphan in α and CREATE(T) occurs in α, then either T is visible to T' in α, or else CREATE(T) is in the scope of an INFORM_ABORT_AT(X)OF(U) in α|R(X).

Proof: There are two cases.

(1) CREATE(T) precedes CREATE(T') in α.
Assume T is not visible to T' in α. Then Lemma 59 implies that there is an INFORM_ABORT_AT(X) operation for some ancestor of T, occurring after CREATE(T) in α.

(2) CREATE(T') precedes CREATE(T) in α
Then Lemma 59 implies that there is either a COMMIT or an ABORT for some ancestor of T', in α. Since T' is not an orphan in α, there is a COMMIT for an ancestor of T' in α. Then Lemma 58 implies that T' is returned in α, a contradiction. ∎

Lemma 64: Let α be a weak concurrent schedule. Let R(X) be a resilient object, let T and T' be accesses to R(X), and let T'' be any transaction. Assume that T'' is not an orphan in α. If an operation π of T precedes an operation π' of T' in α, π is not in the scope of an INFORM_ABORT and T' is visible to T'' in α, then T is visible to T'' in α.

Proof: By well-formedness, CREATE(T) and CREATE(T') are operations in α, in that order. Let α' be the prefix of α ending with CREATE(T'). Then T' is live and not an orphan in α'. By Lemma 63, T is visible to T' in α', and so in α. Lemma 8 implies that T is visible to T'' in α. ∎

The following lemma is straightforward.

Lemma 65: Let α be a weak concurrent schedule, and let T be a transaction which is not an

orphan in α. Any transaction T'' visible to T in α is not an orphan in α.

Proof: If T'' is an ancestor of T, the result is immediate. Otherwise, COMMIT operations appear in α for every proper descendant of lca(T,T'') that is an ancestor of T''. By well-formedness, none of these transactions has aborted. Since the remaining ancestors of T'' are also ancestors of T, and the result follows. ∎

We are now ready to prove the key lemma of this paragraph.

Lemma 66: Let α be a weak concurrent schedule, let T be live and not an orphan in α, and let P be a resilient primitive.

1. If P is a transaction T'', then visible(α,T)|T'' is a prefix of α|T'' and a schedule of T''.

2. If P is a resilient object R(X), then visible(α,T)|R(X) is a prefix of undo(α|R(X)) and a schedule of basic object X.

Proof: 1. Immediate from Lemmas 11 and 1.

2. First, we show that any operation in visible(α,T)|R(X) also occurs in undo(α|R(X)). If π is in visible(α,T)|R(X), it means that all ancestors of transaction(π) up to lca(transaction(π),T) have committed. By assumption, T is not an orphan in α, so Lemma 65 implies that transaction(π) is not an orphan in α. Thus, by the preconditions of the weak concurrent controller there is no INFORM – ABORT for any ancestor of transaction(π) in α. Therefore, π is in undo(α|R(X)).

Now we consider any two operations π and π' of undo(α|R(X)), where π precedes π'. Assume that π' is in visible(α,T)|R(X). Let T'' = transaction(π) and T''' = transaction(π'). Then T''' is visible to T in α, and T'' is not an orphan in α by Lemma 65. Since π is in undo(α|R(X)), no INFORM – ABORT occurs at R(X) for any ancestor of T''' in α, with π in its scope. Then Lemma 64 implies that T'' is visible to T in α. Thus, π is in visible(α,T)|R(X). It follows that visible(α,T)|R(X) is a prefix of undo(α|R(X)).

By Lemma 61, α|R(X) is a well-formed schedule of resilient object R(X). Then the resiliency condition implies that undo(α|R(X)) is a schedule of basic object X. So by Lemma 1, visible(α,T)|R(X) is a schedule of basic object X. ∎

Finally, we prove that, in a weak concurrent schedule, concurrently executing transactions access disjoint sets of resilient objects.

Lemma 67: Let α be a weak concurrent schedule, with transactions T and T'' live and not orphans in α. Let T''' = lca(T,T''). Let β = visible(α,T) - visible(α,T''') and β' = visible(α,T'') - visible(α,T'''). Then no resilient object has operations in both β and β'.

Proof: The result is trivial if T is an ancestor of T'' or vice versa. So assume that lca(T,T'') is neither T nor T''. Let R(X) be a resilient object such that both β and β' contain operations of R(X). By well-formedness, we can assume without loss of generality that there are two accesses to X (not necessarily distinct) such that π = CREATE(U) and φ = CREATE(V) are in β and β', respectively, and neither U nor V is visible to lca(T,T'') in α. Also, we can assume that π appears in α no later than φ.

We have that U is visible to some ancestor of T in α, and V is visible to some ancestor of T'' in α,

and since T and T" are not orphans in α, Lemma 65 implies that no ancestor of U or V has aborted in α. Also, neither U nor V is visible to lca(T,T") in α, so it must be that U $\neq$ V. But then π precedes φ in α, and Lemma 59 implies that some ancestor of T is committed in α. Then Lemma 57 implies that T is returned in α, a contradiction. ∎

7. Simulation of Serial Systems by Concurrent Systems

In this section, we prove the main results of this paper, that concurrent schedules are serially correct, and that weak concurrent schedules are correct at T_0. Both these results follow from an interesting theorem about weak concurrent schedules, which says that the portion of any weak concurrent schedule which is visible to a live non-orphan transaction is equivalent to (i.e. looks the same at *all* primitives as) a serial schedule.

The proof of this theorem is quite interesting, as it provides considerable insight into the scheduling algorithm. The proof shows not only that a transaction's view of a weak concurrent schedule is equivalent to *some* serial schedule, but by a recursive construction, it actually produces such a schedule. It is interesting and instructive to observe how the views that different transactions have of the system execution get passed up and down the transaction tree, as CREATES, COMMITS and ABORTS occur.

Theorem 68: Let α be a weak concurrent schedule, and T any transaction which is live and not an orphan in α. Then there is a serial schedule β which is equivalent to visible(α,T).

Proof: We proceed by induction on the length of α. The basis, length 0, is trivial. Fix α of length at least 1, and assume that the claim is true for all shorter weak concurrent schedules. Let π be the last operation of α, and let $\alpha = \alpha'\pi$. Fix T which is live and not an orphan in α. We must show that there is a serial schedule β which is equivalent to visible(α,T).

If π is not a serial operation, then visible(α',T) = visible(serial(α'),T) = visible(serial(α),T) = visible(α,T), and the result is immediate by induction. So we can assume that π is a serial operation. Also, if transaction(π) is not visible to T in α, then visible(α,T) = visible(α',T), and the result is again immediate by induction. Thus, we can assume that transaction(π) is visible to T in α. Also, T is not an orphan in α'.

There are four cases.

(1) π is an output operation of a transaction or resilient object.
Then the inductive hypothesis implies the existence of a serial schedule β' which is equivalent to visible(α',T). Let $\beta = \beta'\pi$. We must show that β is equivalent to visible(α,T) and serial.

Let P be any primitive. Then $\beta|P = \beta'\pi|P$ = visible(α',T)$\pi|P$ by inductive hypothesis, = visible(α,T)|P, by Lemma 12. Therefore, β is equivalent to visible(α,T).

Let π be an output of primitive P. Then $\beta|P$ = visible(α,T)|P by equivalence, which is a schedule of P by Lemma 66. Lemma 4 implies that β is serial.

(2) π is a CREATE(T") operation.
Then transaction(π) = T", and so T" is visible to T in α. Then Lemma 55 implies that π is the first operation whose transaction is a descendant of T". Then by the definition of visibility, it must be

that T" = T. By Lemma 57, parent(T) is live in α'. Since parent(T) is not an orphan, the inductive hypothesis implies the existence of a serial schedule β' which is equivalent to visible(α',parent(T)). Let $\beta = \beta'\pi$. We must show that β is equivalent to visible(α,T) and serial.

Let P be any primitive. Then $\beta|P = \beta'\pi|P$, = visible(α',parent(T))$\pi|P$ by inductive hypothesis, = visible(α,T)|P, by Lemma 62. Thus, β is equivalent to visible(α,T).

Consider any execution of the serial system having β' as its operation sequence, and let s' be the state of the serial scheduler after β'. We show that π is enabled in s'. That is, we show that T $\in$ create−requested(s'), that T $\notin$ created(s'), and that siblings(T) $\cap$ created(s') $\subseteq$ returned(s').

Consider any execution of the weak concurrent system having α as its operation sequence, and let s be the state of the weak concurrent scheduler after α'. State s contains a component s_c for the weak concurrent controller and a component s_X for the lock manager for each object X.

If T = T_0, then T $\in$ create−requested(s') by the initial conditions. If T $\neq$ T_0, then T $\in$ create−requested(s_c) by the preconditions of the concurrent scheduler, so a REQUEST−CREATE(T) operation occurs in α'. The REQUEST−CREATE(T) operation has transaction parent(T), and so is in visible(α',parent(T)), and thus is in β'. Therefore, T $\in$ create−requested(s').

If T $\in$ created(s'), then there is either a CREATE(T) or an ABORT(T) operation in β', and hence in α'. In the former case, α would have two such operations, while in the latter case, α would have an ABORT(T) followed by a CREATE(T). Both are impossible, so T $\notin$ created(s').

Assume U $\in$ siblings(T) $\cap$ created(s'). Then there is either a CREATE(U) or an ABORT(U) operation in β'. In the latter case, U is obviously in returned(s'). So suppose CREATE(U) occurs in β', and so in visible(α',parent(T)). Since CREATE(U) occurs at U, U is visible to parent(T) = parent(U) in α'; thus, COMMIT(U,u) occurs in α', for some u. Since COMMIT(U,u) occurs at parent(T), COMMIT(U,u) is in visible(α',parent(T)), and so in β'. Thus, U $\in$ returned(s').

(3) π is a COMMIT(T",v) operation.
Then T"' = parent(T") = transaction(π) is visible to T and not an orphan in α. Also, T" is not an orphan in α', by Lemma 65. Then since α is well-formed, T" is live in α', and so by Lemma 57, T"' is live in α' and so in α. Since T"' is live and visible to T, T"' is an ancestor of T. Since T is live in α, Lemma 58 implies that T is not a descendant of T". The inductive hypothesis yields two serial schedules, β' and β", which are equivalent to visible(α',T") and visible(α',T), respectively. Let γ = visible(β',T"). Let $\beta_1 = \beta' \cdot \gamma$ and $\beta_2 = \beta'' \cdot \gamma$. We show that $\beta = \gamma\beta_1\pi\beta_2$ is equivalent to visible(α,T) and serial.

Lemma 28 implies that γ is a serial schedule.

Since T"' is visible to T" in α', Lemma 10 implies that visible(α',T"') = visible(visible(α',T"),T"'), which is equivalent to visible(β',T"') = γ; thus γ is equivalent to visible(α',T"'). Also, since T"' is visible to T in α', Lemma 10 implies that visible(α',T"') = visible(visible(α',T),T"'), which is equivalent to visible(β",T"'). Thus, γ is also equivalent to visible(β",T"').

Then by Lemma 31 (applied with serial(α') as the schedule α hypothesized in the lemma), $\gamma\beta_1$

and $\gamma\beta_2$ are serial schedules which are equivalent to β' and β'', respectively.

We have that visible(α,T''') = visible(α',T''')π by Lemma 62, which is equivalent to $\beta'\pi$, which is in turn equivalent to $\gamma\beta_1\pi$. That is, visible(α,T''') is equivalent to $\gamma\beta_1\pi$.

Since β'' is equivalent to visible(α',T) and γ is equivalent to visible(α',T'''), by Lemma 10, β_2 = β'' - γ is equivalent to visible(α',T) - visible(α',T'''), = visible(α,T) - visible(α,T''') by Lemma 62.

Thus, β is equivalent to visible(α,T''')(visible(α,T)-(visible(α,T''')). Since T''' is visible to T in α, by Lemma 8, it is easy to see that the same operations appear in this schedule as in visible(α,T). Let P be any primitive. Then visible(α,T'''))|P is a prefix of visible(α,T)|P, by Lemma 66. It follows that β|P = visible(α,T)|P, so that β is equivalent to visible(α,T).

It remains to show that β is serial. This follows from Lemma 32, provided we can show that:
(3.a) $\gamma\beta_1\pi$ is a serial schedule,
(3.b) T'' sees everything in $\gamma\beta_1$,
(3.c) T sees everything in $\gamma\beta_2$,
(3.d) γ = visible($\gamma\beta_1$,T''') = visible($\gamma\beta_2$,T''') and
(3.e) no basic object has operations in both β_1 and β_2.

(3.a) Consider any execution of the serial system having $\gamma\beta_1$ as its operation sequence, and let s' be a state of the serial scheduler after $\gamma\beta_1$. We show that π is enabled in state s'. That is, we show that (T'',v) $\in$ commit−requested(s'), that T'' $\notin$ returned(s'), and that children(T'') $\cap$ create−requested(s') $\subseteq$ returned(s').

Consider any execution of the weak concurrent system having α as its operation sequence, and let s be the state of the weak concurrent scheduler after α', with components s_c (the weak controller state), and s_X for every object X (the lock managers).

Since the weak concurrent scheduler is able to perform COMMIT(T'',v) in state s, it must be that (T'',v) is in commit−requested(s_c), and so it must be that T'' issues a REQUEST−COMMIT(T'',v) in α'. Since T'' is visible to itself, and β' is equivalent to visible(α',T''), it follows that this REQUEST−COMMIT(T'',v) operation also occurs in $\gamma\beta_1$. Therefore, (T'',v) is in commit−requested(s').

Since α is well-formed, at most one return operation for T'' appears in α; thus, T'' is not in returned(s').

Fix U $\in$ children(T'') $\cap$ create−requested(s'). Then REQUEST−CREATE(U) is performed at T'' in $\gamma\beta_1$, and hence in α', so U $\in$ create−requested(s_c). Since the weak concurrent scheduler is able to perform COMMIT(T'',v) in state s, it must be that U $\in$ returned(s_c). Therefore, a return operation for U is performed at T'', in α'. Since T'' is visible to itself, and $\gamma\beta_1$ is equivalent to visible(α',T''), this implies that the return for U also occurs at T'' in $\gamma\beta_1$. Therefore, U is in returned(s').

(3.b) Immediate from Lemma 10.

(3.c) Immediate from Lemma 10.

(3.d) We have that γ is equivalent to both visible(β',T") and visible(β",T"), and that $\gamma\beta_1$ and $\gamma\beta_2$ are equivalent to β' and β", respectively. By Lemma 10, γ is equivalent to both visible($\gamma\beta_1$,T") and visible($\gamma\beta_2$,T"). Equality follows.

(3.c) Immediate from Lemma 67.

(4) π is an ABORT(T") operation.
Then T"' = parent(T") = transaction(π) is visible to T in α, and so is not an orphan in α, by Lemma 65. Then T" is live in α', and by Lemma 57, T"' is live in α' and so in α. Since T"' is live and visible to T in α, T is a descendant of T"'. Since T is not an orphan in α, T is not a descendant of T". The inductive hypothesis yields two serial schedules, β' and β", which are equivalent to visible(α',T"') and visible(α',T), respectively. Let $\beta_1 = \beta" - \beta'$. We show that $\beta = \beta'\pi\beta_1$ is equivalent to visible(α,T) and serial.

By Lemma 31, $\beta'\beta_1$ is a serial schedule which is equivalent to β".

Let P be a primitive other than T"'. Then $\beta|P = \beta'\beta_1|P = \beta"|P$ = visible(α',T)|P, = visible(α,T)|P by Lemma 62. Also, since T"' is visible to T in α, visible(α,T)|T"' = visible(α,T"')|T"', = visible(α',T"')π|T"' by Lemma 62, = $\beta'\pi$|T"' = β|T"'. Thus β is equivalent to visible(α,T).

It remains to show that β is serial. This follows from Lemma 33, provided we can show that:
(4.a) $\beta'\pi$ is a serial schedule,
(4.b) T sees everything in $\beta'\beta_1$, and
(4.c) β' = visible(β',T") = visible($\beta'\beta_1$,T").

(4.a) Consider any execution of the serial system having β' as its operation sequence, and let s' be a state of the serial scheduler after β'. We show that π is enabled in state s'. That is, we show that T" $\in$ create−requested(s'), that T" $\notin$ created(s'), and that siblings(T") $\cap$ created(s') $\subseteq$ returned(s').

Consider any execution of the weak concurrent system having α as its operation sequence, and let s be the state of the weak concurrent scheduler after α', with components s_c (the weak controller state), and s_X for every object X (the lock managers).

Since the weak concurrent scheduler is able to perform ABORT(T") in state s, it must be that T" is in create−requested(s_c), and so it must be that T"' issues a REQUEST−CREATE(T") in α'. Since T"' is visible to itself, and β' is equivalent to visible(α',T"), it follows that this REQUEST−CREATE(T") operation also occurs in β'. Therefore, T" is in create−requested(s').

Since α cannot contain two ABORT(T") operations, there cannot be an ABORT(T") operation in α', and so there cannot be one in β'. Assume that there is a CREATE(T") in β'. Then T" is visible to T"' in α', so COMMIT(T") occurs in α'. But then a COMMIT(T") and and ABORT(T") both occur in α, which cannot occur. Therefore, there is neither an ABORT(T") nor a CREATE(T") in β', and so T" is not in created(s').

Fix U $\in$ siblings(T") $\cap$ created(s'). Then there is a CREATE(U) in β'. But then U is visible to T" in α', so that a COMMIT for U occurs in α', and hence (because parent(U) is visible to T"' in

α') a COMMIT for U occurs in β'. Therefore, U $\in$ returned(s').

(4.b) Immediate from Lemma 10.

(4.c) The first equality follows from Lemma 10. Clearly, β' = visible(β',T''') is a prefix of visible($\beta'\beta_1$,T'''). Equality follows because any operation in β_1 visible to T''' in $\beta'\beta_1$ would also be visible to T''' in α', and so would be in β' and not β_1. ∎

Corollary 69: Every weak concurrent schedule is serially correct for every non-orphan non-access transaction.

Proof: Let α be a weak concurrent schedule. Let T be a non-access transaction that is not an orphan in α. We must show that $\alpha|T$ is a serial schedule. Note that T is not an orphan in any prefix of α.

There are three cases:

(1) $\alpha|T$ is empty.
Then the result is trivial.

(2) T is live in α.
Then Theorem 68 yields a serial schedule β that is equivalent to visible(α,T). Thus, $\alpha|T$ = visible(α,T)$|T$ = $\beta|T$, which suffices.

(3) T is a transaction which is live in some proper prefix of α.
Since α is well-formed, α has a prefix $\alpha'\pi$, where π is a COMMIT operation for T, $\alpha'|T = \alpha|T$ and T is live in α'. Then Theorem 68 yields a serial schedule β that is equivalent to visible(α',T)$|T$. Thus, $\alpha|T = \alpha'|T$ = visible(α',T)$|T$ = $\beta|T$, which suffices. ∎

Now, since T_0 cannot become an orphan (having no ancestors to abort), our first major correctness result is immediate.

Corollary 70: Every weak concurrent schedule is serially correct for T_0.

Having proved correctness properties for weak concurrent schedules, we are now ready to prove the correctness of concurrent schedules.

Lemma 71: Every concurrent execution is a weak concurrent execution.

Proof: The proof is by induction on execution length, with a trivial basis. Let $\alpha = \alpha', s', \pi, s$ be a concurrent execution with (s', π, s) a single step of the concurrent system, and assume the lemma holds for α'. Let s'_c and s_c denote the states of the concurrent controller in system states s' and s. If π is any operation other than an ABORT, the result is immediate, since the pre- and postconditions for all other operations are identical in the concurrent and weak concurrent systems. Assume that π is an ABORT(T). We must show that T $\in$ create-requested(s'_c) - returned(s'_c).

Since π is enabled in state s'_c in the concurrent controller, T $\in$ (create-requested(s'_c) - created(s'_c)) $\cup$ (commit-requested(s'_c) - returned(s'_c)). If T is in create-requested(s'_c) - created(s'_c), Lemma 45 implies that α' contains no CREATE(T) or ABORT(T) operation. By

well-formedness, α' also contains no COMMIT operation for T, and the result follows from Lemma 45. On the other hand, if T is in commit-requested(s'_c) - returned(s'_c), Lemma 45 implies that a REQUEST−COMMIT operation for T occurs in α'. By well-formedness, this is preceded by a CREATE(T) operation, and by the concurrent controller precondition, this is preceded by a REQUEST−CREATE for T. Finally, again by Lemma 45, the result follows. ∎

Now we can prove the second major result of the paper.

Corollary 72: Every concurrent schedule is serially correct.

Proof: Let α be a concurrent schedule. Then α is also a weak concurrent schedule, by Lemma 71, and is well-formed, by Lemma 61. We must show that α is serially correct for every transaction T. There are three cases:

(1) $\alpha|T$ is empty.
Then the result is trivial.

(2) T is live in α.
By Lemma 50, all of T's ancestors are live in α, so that T is not an orphan in α. Then Corollary 69 yields the result.

(3) T is a transaction which is live in some proper prefix of α.
By Lemma 51, α has a prefix $\alpha'\pi$, where π is a return operation for T, $\alpha'|T = \alpha|T$ and T is live in α'. By Lemma 50, all of T's ancestors are live in α', so T is is not an orphan in α'. Then Corollary 69 implies that α' is serially correct for T, so that α is serially correct for T. ∎

For completeness, we include an analog of Theorem 68 for concurrent schedules.

Theorem 73: Let α be a concurrent schedule, and T any transaction which is live in α. Then there is a serial schedule β which is equivalent to visible(α,T).

Proof: Lemma 71 implies that α is a weak concurrent schedule. Since T is live in α, Lemma 50 implies that T is not an orphan in α. Then Theorem 68 yields the result. ∎

8. Discussion

In this paper, we have presented a formal model for describing nested transaction systems and their properties. The model has many features that we believe make it a major contribution to the understanding of transaction systems, and we highlight some of these below.

First, the entire model is based on a very general and very simple underlying model for concurrent computation, the I/O automaton model. The general definitions and properties of this underlying model provide the necessary underpinnings for our entire transaction modelling effort. This modelling is very easy to learn and use, and its usefulness extends much beyond transaction systems. Thus, it seems to us to be a very satisfactory foundation for our work.

Our transaction system model permits simple, yet completely rigorous description of concurrency control

algorithms in ways which correspond very closely to the usual informal ways of understanding the algorithms. The important components of transaction systems, the transactions, data and schedulers, are described explicitly, which greatly facilitates reasoning about them.

There is a substantial amount of work in this area which does not represent all of these components explicitly, but only implicitly, by properties of their behavior [Ly,BBG,Go, for example]. There are problems with this approach. A key ingredient that is usually absent from such implicit models is a clear notion of "causality", describing how particular actions (operations) are triggered by other actions or states. In contrast, our explicit representation of all system components as I/O automata makes it easy to understand exactly what causes all operations to occur. When causality is important in reasoning about algorithms, as in [Go], implicit models can be extraordinarily difficult to use. Even in cases where implicit models can be used, we see the present work as providing a formal and intuitive basis for that work.

Our model represents transactions as essentially arbitrary automata, subject only to simple syntactic constraints. This approach is much more general than representing them as programs in some particular, overly-constrained language.

The I/O automata model permits description of algorithms in an abstract form which is not tied to a particular programming language or system, and which allows maximum nondeterminism. An implementation of an algorithm for a particular system will generally restrict the nondeterminism allowed in our presentation, because of the need to tailor the implementation to the requirements of a particular environment. However, since the implementation is just a restriction of the abstract algorithm, correctness properties of the algorithm within our model will hold *a fortiori* for the implementation.

Formulating nested transaction systems as I/O automata permits precise formulation of the correctness conditions to be satisfied by concurrency control algorithms; those correctness conditions can be stated at the transaction interface, an interface which does not contain explicit information about object representation. Because of this choice of interface, the correctness conditions can be stated in a robust way: the same conditions can be useful for describing the properties of many different kinds of algorithms, some of which manipulate the data in very different ways. Also, the correctness conditions can be described in a way that is meaningful to a user of the system.

The particular correctness conditions that we describe in this paper involve serial correctness at transaction interfaces. We believe that these particular correctness definitions are very interesting, and will be useful for describing the correctness of most of the usual algorithms studied in the concurrency control area. That is, the same conditions appear to be the right ones to use to describe correctness of many different kinds of

algorithms, including those that use locking, timestamps, multiple versions, and replicated data.

The model permits rigorous correctness proofs to be carried out for concurrency control algorithms in ways that follow intuitive understanding of the algorithms. For example, in this paper, we have used the model to describe and show the correctness of a very important nested transaction concurrency control algorithm. Our correctness proofs are constructive and provide considerable intuition about the workings of the algorithm. In contrast to most correctness proofs for concurrent algorithms, our proofs are not voluminous low-level case-analyses; rather, they consist of a large number of clear and natural lemmas about the behavior of the algorithm. These lemmas can be understood individually, and build upon each other in the manner of good mathematics. Many of the lemmas should be reusable in extensions of this work as well.

A successful model of nested transactions should contain the classical theory as a special case, in a way which is natural and sheds some light on that case. We believe that our model has contributed much to the classical theory. For example, the I/O automaton model provides a general underlying model, a missing component of the classical theory. Also, our explicit and general modelling of the transactions unifies the earlier collection of somewhat arbitrary approaches. Our use of the transaction interface for stating correctness conditions is also an improvement.

Another contribution to the classical theory is in motivating serializability as a correctness condition. Serializability consists of two criteria: individually, each transaction must see a consistent state, and together, they must appear to run in a serial order. (A schedule in which each transaction reads and writes the initial state of the database provides a consistent state to each transaction, but is not serializable.) Why is this second ordering property a part of the generally accepted correctness condition of the classical theory? Clearly, because of implicit nesting in the context of the transaction system. In practice, transactions perform tasks on behalf of some external entity or entities, which may expect one transaction to see the results of the next. In the natural formulation of classical systems within the present model, the classical transactions are children of T_0, with data accesses as their only children. The root is an explicit representation of the external environment in which the system runs. Thus, the ordering property of serializability is a natural consequence of the requirement that all transactions see serial schedules, including T_0. It does not have to be introduced as an independent requirement in need of separate justification.

By now, there has been a large amount of systems design and algorithms work that uses or implements nested transactions. It seems likely that these ideas will form the basis of future programming languages for distributed computing. However, there is currently a problem with the presentation of this work. Some of these algorithms are presented in the context of specific systems and programming languages. Very useful and general ideas are too intimately connected with details of the systems to be fully appreciated, particularly

for readers with only a passing understanding of those systems. This level of detail also makes careful reasoning about the algorithms very difficult.

We believe that our model has provided the necessary framework and some of the necessary vocabulary, for describing this work in a clear and unambiguous way. We are currently studying much of this work on systems design and algorithms using our model, and our preliminary results indicate that it works very well.

Throughout the paper, we have described connections with other people's work as appropriate. Here, we mention some of the particular modelling work that relates most closely to ours, and describe the connections in more detail.

First, the pioneering work of Bernstein and Goodman [BG, etc.] has had a strong influence on this work. Quite early, they recognized the need for a model for single-level transaction systems, that would have many of the characteristics which we have sought for nested transaction systems. They have carried out extensive research on precise understanding of single-level transaction concurrency control algorithms. They have presented formal statements of correctness conditions, in terms of serializability of the accesses to data objects by different transactions. They have described some concurrency control algorithms with precision, and have proved correctness of some algorithms, using a lemma which characterizes serializability by absence of cycles in a certain dependency relation. Their work has gone a long way toward providing precise understanding of the work in this area.

However, the particular models used by Bernstein and Goodman have some problems which limit their applicability. For instance, the basic correctness condition is stated in terms of the interface between the data objects and the algorithm. There are many algorithms which handle objects in very different ways, e.g. using multiple versions, or making multiple copies in order to permit "backing out" of operations. Since these algorithms do not preserve the specified object interface, they would not be considered correct under the same correctness condition. Thus, the correctness condition must be modified. Another limitation is that the proof technique, which involves proving absence of cycles, is a proof by contradiction; it does not give much insight into the operation of the algorithms. For many reasons, it is not at all clear how to extend these frameworks to handle nesting of transactions.

Earlier attempts in [Ly,Go,BBGLS] to model nested transactions have made significant contributions. For example, [Ly] contains a language-independent model, which is used to give precise correctness conditions and a proof for a locking algorithm. Many of the ideas in that work have been useful in providing a vocabulary for talking about nested transactions. However, attempts to extend the model of [Ly] to handle correctness of orphans [Go] demonstrate that it is not sufficiently expressive. Certain aspects of the model

lead to technical difficulties; for example, it fails to model the transactions explicitly, using instead a specification of their behavior. Our new model builds on the strengths of the earlier work, while managing (we believe) to avoid its weaknesses.

Finally, the very recent work in [BBG] proposes another general model for nested transactions. While on the surface the models appear quite different, they are actually "compatible", in that the concepts described in [BBG] seem to be easily definable within our model. The style of the model in [BBG] is different from ours: their work models transactions and the scheduler implicitly, for instance. However, we believe that their important axiomatic statements of properties can be described as assumptions and lemmas about behaviors of components in our model. Also, the partial orders which they use to model executions can actually be defined simply and directly in terms of our linearly-ordered executions. There are many points of agreement: the use of the transaction interface for stating correctness conditions, and the use of the virtual root transaction T_0, to mention two.

On the other hand, the emphasis in [BBG] is on a different example than the one studied in this paper. They consider multiple levels of abstraction for the data, and regard transactions at any level of the transaction tree as accesses to data at a corresponding level of abstraction. This view meshes quite well with our model, where, for example, we use the same CREATE notation for creation of a transaction and invocation of an operation on data. Their paper clarifies the concurrency control requirements for data at different levels, when the correctness condition is serial correctness at T_0. We hope and expect that it will be easy to restate their results as claims about our model.

We note that the work in [BBG] only treats concurrency control, but does not address the very critical and difficult issues of resiliency.

9. Further Work

This paper is an embarkation on a major project to formulate a unified presentation of the most important algorithms for concurrency control and resiliency, especially those for nested transactions. So far, we have defined a general framework meeting the requirements outlined above. We have demonstrated the power of this framework by using it to specify two correctness conditions for nested transactions, to present two locking algorithms for implementing nested transactions, and to prove that the algorithms satisfy their respective requirements.

Future extensions to this work will include treatment of many other algorithms in the same framework. Among the algorithms we will consider are timestamp and multiversion algorithms, algorithms which take advantage of special properties of the transactions and objects (semantic information), algorithms for orphan

management, and algorithms which use replicated data objects. Although our focus so far has been on nested transactions, we believe that our viewpoint contributes new insight to the special case of single-level transactions as well; thus, we will examine algorithms for non-nested transactions as well as nested transactions.

We are particularly interested in studying algorithms which give rise to live orphans, i.e. live transactions whose ancestors have aborted [Go,Li,Wa,HM]. Our serial correctness condition provides a formal definition of *orphan correctness* – that all transactions (including orphans) "see consistent data" [Go]. In fact, in work currently in progress [HLMW], we are describing and proving correctness of several of the recently-developed algorithms for orphan management. This work now seems to be quite easy, given the foundation provided by the present paper. In fact, some of the key results of this paper are used as lemmas in that work.

Another direction of interest is the explicit representation of distribution within the model. It is fairly natural to model each transaction and object as located at different sites, each with a local automaton representing the resident portion of the (distributed) scheduler. These automata would communicate with each other in order to implement the (centralized) scheduler studied here. The natural next step would be to model failure resilience, as various components lose information or fail altogether.

The reader might have noted that our correctness conditions do not guarantee anything about the system making progress, but only about "safety" properties. Further work is needed to incorporate guarantees of progress. This work is likely to be difficult, however. Only recently, in [LT], have we achieved what we consider to be a satisfactory understanding of the eventuality and fairness issues for the basic I/O automaton model, so that we can even formulate the conditions we want to satisfy. But even with the ability to state such conditions, the algorithmic issues still seem difficult.

10. Acknowledgments

We thank Bill Weihl for many, many comments and questions, and much encouragement, during the course of this project. We also thank all the other members of the ARGUS design and implementation group at MIT, for providing a concrete model for us to try to abstract and generalize. Also, we thank Yehuda Afek for his comments on an early draft, and Sharon Perl for her comments on later drafts.

11. References

[AM] Allchin, J. E., and McKendry, M. S., "Synchronization and Recovery of Actions," *Proc. 1983 Second Annual ACM Symposium on Principles of Distributed Computing*, Montreal, Quebec, Canada, August 17-19, 1982, pp. 31-44.

[BBG] Beeri, C., Bernstein, P. A., and Goodman, N., "A Model for Concurrency in Nested

Transaction Systems," Manuscript.

[BBGLS] Beeri, C., Bernstein, P. A., Goodman, N., Lai, M. Y., and Shasha, D. E., "A Concurrency Control Theory for Nested Transactions," *Proc. 1983 Second Annual ACM Symposium on Principles of Distributed Computing*, Montreal, Quebec, Canada, August 17-19, 1983, pp. 45-62.

[BG] Bernstein, P. A., and Goodman, N., "Concurrency Control in Distributed Database Systems," *ACM Computing Surveys* 13,2 (June 1981), pp. 185-221.

[EGLT] Eswaren, K. P., Gray, J. N., Lorie, R. A., and Traiger, I. L., "The Notions of Consistency and Predicate Locks in a Database Systems," *Communications of the ACM*, Vol. 19, No. 11, November 1976, pp. 624-633.

[Go] Goree, Jr., John A., "Internal Consistency of a Distributed Transaction System With Orphan Detection," MS Thesis, Technical Report MIT/LCS/TR-286, MIT Laboratory for Computer Science, Cambridge, MA., January 1983.

[Gr] Gray, J., "Notes on Database Operating Systems," in Bayer, R., Graham, R. and Seegmuller, G. (eds), Operating Systems: an Advanced Course, Lecture Notes in Computer Science, Vol. 60, Springer-Verlag, 1978.

[HLMW] Herlihy, M., Lynch, N., Merritt, M., and Weihl, W., "On the Correctness of Orphan Elimination Algorithms," In progress.

[HM] Herlihy, M., and McKendry, M., "Time-Driven Orphan Elimination", in *Proc. of the 5th Symposium on Reliability in Distributed Software and Database Systems*, Los Angeles, CA., January 1986, pp. 42-48.

[Ho] Hoare, C.A.R., "Communicating Sequential Processes", Prentice Hall International Englewood Cliffs, NJ, 1985.

[KS] Kedem, Z., and Silberschatz, A., "A Characterization of Database Graphs Admitting a Simple Locking Protocol", *Acta Informatica* 16 (1981) pp. 1-13.

[LaS] Lampson, B. W., and Sturgis, H. E., "Crash Recovery in a Distributed Data Storage System," Tech. Rep., Computer Science Lab., Xerox Palo Alto Research Center, Palo Alto, Calif., 1979.

[LHJLSW] Liskov, B., Herlihy, M., Johnson, P., Leavens, G., Scheifler, R., and Weihl, W., "Preliminary Argus Reference Manual," Programming Methodology Group Memo 39, October 1983.

[Li] Liskov, B., and Ladin, R., "Highly-Available Distributed Services and Fault-Tolerant Distributed Garbage Collection," to appear in the *Proc. of the 5th ACM Symposium on Principles of Distributed Computing*, Calgary, Alberta, Canada, August 1986.

[LiS] Liskov, B., and Scheifler, R., "Guardians and Actions: Linguistic Support for Robust,

Distributed Programs", *ACM Transactions on Programming Languages and Systems* 5, 3, (July 1983), pp. 381-404.

[LT] Lynch, N., and Tuttle, M., "Correctness Proofs for Distributed Algorithms", in progress.

[Ly] Lynch, N. A., "Concurrency Control For Resilient Nested Transactions," *Advances in Computing Research* 3, 1986, pp. 335-373.

[Mi] Milner, R., "A Calculus of Communicating Systems", *Lecture Notes in Computer Science*, #92, Springer-Verlag, Berlin, 1980.

[Mo] Moss, J. E. B., "Nested Transactions: An Approach To Reliable Distributed Computing," Ph.D. Thesis, Technical Report MIT/LCS/TR-260, MIT Laboratory for Computer Science, Cambridge, MA., April 1981. Also, published by MIT Press, March 1985.

[R] Reed, D. P., "Naming and Synchronization in a Decentralized Computer System," Ph.D Thesis, Technical Report MIT/LCS/TR-205, MIT Laboratory for Computer Science, Cambridge, MA 1978.

[RLS] Rosenkrantz, D. J., Lewis, P. M., and Stearns, R. E., "System Level Concurrency Control for Distributed Database Systems," *ACM Transactions on Database Systems*, Vol. 3, No. 2, June 1978, pp. 178-198.

[Wa] Walker, E. F., "Orphan Detection in the Argus System," M.S. Thesis, Technical Report/MIT/LCS/TR-326, MIT Laboratory for Computer Science, Cambridge, MA., June 1984.

[Wel] Weihl, W. E., "Specification and Implementation of Atomic Data Types," Ph.D Thesis, Technical Report/MIT/LCS/TR-314, MIT Laboratory for Computer Science, Cambridge, MA., March 1984.

[We2] Weihl, W. E., Personal communication.

CS-TR Scanning Project
Document Control Form

Date : 4/11/95

Report # LCS-TR-367

Each of the following should be identified by a checkmark:
Originating Department:

☐ Artificial Intellegence Laboratory (AI)
☒ Laboratory for Computer Science (LCS)

Document Type:

☒ Technical Report (TR) ☐ Technical Memo (TM)
☐ Other:_____

Document Information

Number of pages: 62(69-IMAGES)
Not to include DOD forms, printer intstructions, etc... original pages only.

Originals are:
☒ Single-sided or
☐ Double-sided

Intended to be printed as :
☐ Single-sided or
☒ Double-sided

Print type:
☐ Typewriter ☐ Offset Press ☐ Laser Print
☐ InkJet Printer ☒ Unknown ☐ Other:_____

Check each if included with document:

☒ DOD Form ☐ Funding Agent Form ☐ Cover Page
☒ Spine ☒ Printers Notes ☐ Photo negatives
☐ Other: _____

Page Data:

Blank Pages(by page number):_____

Photographs/Tonal Material (by page number):_____

Other (note description/page number):

Description :	Page Number:
IMAGE MAP (1-2) UN#'ED TITLE & INTRO PAGES	
(3-62) PAGES #'ED 1-60	
(63-66) SCANCONTROL, SPINE, PRINTERS NOTES, POD	
(67-69) TRGTS (3)	

Scanning Agent Signoff:

Date Received: 4/11/95 Date Scanned: 4/18/95 Date Returned: 4/20/95

Scanning Agent Signature: Michael W. Cobb

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER MIT/LCS/TR-367	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Introduction to the theory of nested transactions		5. TYPE OF REPORT & PERIOD COVERED 6. PERFORMING ORG. REPORT NUMBER MIT/LCS/TR-367
7. AUTHOR(s) Nancy Lynch and Michael Merritt		8. CONTRACT OR GRANT NUMBER(s) DARPA/DOD N00014-85-K-0168
9. PERFORMING ORGANIZATION NAME AND ADDRESS MIT Laboratory for Computer Science 545 Technology Square Cambridge, MA 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS DARPA/DOD 1400 Wilson Boulevard Arlington, VA 22209		12. REPORT DATE July 1986 13. NUMBER OF PAGES 62
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) ONR/Department of the Navy Information Systems Program Arlington, VA 22217		15. SECURITY CLASS. (of this report) Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release, distribution is unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

unlimited

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

nested transactions, transacations, concurrency control, resiliency, databases, serializability, orphans

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A new formal model is presented, for studying concurrency and resiliency properties for nested transactions. The model is used to state and prove correctness of a well-known locking algorithm.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Scanning Agent Identification Target

Scanning of this document was supported in part by the **Corporation for National Research Initiatives,** using funds from the **Advanced Research Projects Agency** of the **United states Government** under Grant: **MDA972-92-J1029.**

The scanning agent for this project was the **Document Services** department of the **M.I.T Libraries.** Technical support for this project was also provided by the **M.I.T. Laboratory for Computer Sciences.**

Scanned

Date: 4/18/1995

M.I.T. Libraries
Document Services