

Knowledge-Based Synthesis of Distributed Systems Using Event Structures

Mark Bickford

Robert C. Constable

Joseph Y. Halpern

Sabina Petride

Department of Computer Science

Cornell University

Ithaca, NY 14853

{markb,rc,halpern,petride}@cs.cornell.edu

Abstract

To produce a program guaranteed to satisfy a given specification one can synthesize it from a formal constructive proof that a computation satisfying that specification exists. This process is particularly effective if the specifications are written in a high-level language that makes it easy for designers to specify their goals. We consider a high-level specification language that results from adding *knowledge* to a fragment of Nuprl specifically tailored for specifying distributed protocols, called *event theory*. We then show how high-level *knowledge-based programs* can be synthesized from the knowledge-based specifications using a proof development system such as Nuprl. Methods of Halpern and Zuck [1992] then apply to convert these knowledge-based protocols to ordinary protocols. These methods can be expressed as heuristic transformation tactics in Nuprl.

1 Introduction

Errors in software are extremely costly and disruptive. NIST (the National Institute of Standards and Technology) estimates the cost of software errors to the US economy at \$59.5 billion per year. One approach to minimizing errors is to synthesize programs from specifications. Synthesis methods have produced highly reliable moderate-sized programs in cases where the computing task can be precisely specified. One of the most elegant synthesis methods is the use of so-called *correct-by-construction* program synthesis [Bates and Constable 1985; Constable et al. 1986; Smith and Green 1996; Kreitz 1998; Paulin-Mohring and Werner 1993; Geuvers et al. 2001; Geuvers et al. 2001]. Here programs are constructed from *proofs* that the specifications are satisfiable. That is, a constructive proof that a specification is satisfiable gives a program that satisfies the specification. This method has been successfully used by several research groups and companies to construct large complex *sequential* programs, but it has not yet been used to create substantial realistic distributed programs.

The Cornell Nuprl proof development system was among the first tools used to create correct-by-construction functional and sequential programs [Constable et al. 1986]. Nuprl has also been used extensively to optimize distributed protocols [Liu et al. 1999; Birman et al. 2000], and to specify them in the language of I/O Automata [Bickford et al. 2001; Bickford et al. 2001; Liu et al. 2001]. Recent work by two of the authors has resulted in the definition of a fragment of the higher-order logic used by Nuprl tailored to specifying distributed protocols, called *event theory*, and the extension of Nuprl methods to synthesize distributed protocols from specifications written in event theory [Bickford 2003; Bickford and Constable 2003].

Event logic is a specification language closely related to I/O automata. As has long been recognized [Halpern and Moses 1990], designers typically think of specifications at a high level, which often involves knowledge-based statements. For example, the goal of a program might be to guarantee that a

certain process knows certain information. It has been argued [Fagin et al. 1995; Fagin et al. 1997] that a useful way of capturing these high-level knowledge-based specifications is by using high-level *knowledge-based programs*. Knowledge-based programs are an attempt to capture the intuition that what an agent does depends on what it knows. For example, a knowledge-based program may say that process 1 should stop sending a bit to process 2 once process 1 knows that process 2 knows the bit. Such knowledge-based programs and specifications can be given precise semantics [Fagin et al. 1995; Fagin et al. 1997; Halpern 1999]. They have already met with some degree of success, having been used in papers such as [Dwork and Moses 1990; Hadzilacos 1987; Halpern et al. 2001; Halpern and Zuck 1992; Mazer and Lochovsky 1990; Mazer 1990; Moses and Tuttle 1988; Neiger and Toueg 1993; Stulp and Verbrugge 2002] both to help in the design of new protocols and to clarify the understanding of existing protocols.

In this paper, we add knowledge operators to event theory raising its level of abstraction and show by example that knowledge-based programs can be synthesized from constructive proofs that specifications in event theory with knowledge operators are satisfiable. Our example uses the *sequence-transmission problem*, where a sender must transmit a sequence of bits to a receiver in such a way that the receiver eventually knows arbitrarily long prefixes of the sequence. Halpern and Zuck [1992] provide two knowledge-based programs for the sequence-transmission, prove them correct, and show that many standard programs for the problem in the literature can be viewed as implementations of their high-level knowledge-based program. Here we show that these two knowledge-based programs can be synthesized from the specifications of the problem, expressed in event theory augmented by knowledge. We can then translate the arguments of Halpern and Zuck to Nuprl, to show that the knowledge-based programs can be transformed to the standard programs in the literature.

Engelhardt, van der Meyden, and Moses [1998, 2001] have also provided techniques for synthesizing knowledge-based programs from knowledge-based specifications, by successive refinement. We see their work as complementary to ours. Since our work is based on Nuprl, we are able to take advantage of the huge library of tactics provided by Nuprl to be able to generate proofs. The expressive power of Nuprl also allows us to express all the high-level concepts of interest (both epistemic and temporal) easily. Engelhardt, van der Meyden, and Moses do not have a theorem-proving engine for their language. However, they do provide useful refinement rules that can easily be captured as tactics in Nuprl.

2 Synthesizing Distributed Programs From Constructive Proofs

2.1 Nuprl: a brief overview

Much current work on formal verification using theorem proving, including Nuprl, is based on type theory (see [Constable 2002] for a recent overview). A type can be thought of as a set with structure that facilitates its use as a data type in computation; this structure also supports constructive reasoning. The set of types is closed under constructors such as \times and \rightarrow , so that if A and B are types, so are $A \times B$ and $A \rightarrow B$. where, intuitively, $A \rightarrow B$ represents the computable functions from A into B . *Constructive* type theory, upon which Nuprl is based, was developed to provide a foundation for constructive mathematics. The key feature of constructive mathematics is that “there exists” is interpreted as “we can construct (a proof of)”. A consequence of this approach is that, for example, the law of excluded middle does not hold.

At an abstract level, a program in Nuprl is just an object of some type Pgm . A *program semantics* is a function $S \in \text{Pgm} \rightarrow \text{Sem}$ assigning to each *program* $\text{pr} \in \text{Pgm}$ a *meaning* in type Sem . A *semantic property* is a predicate X on meanings. We say that a *program* pr *satisfies a semantic property* X if X holds of the semantic meaning of pr . Formally, we write $\text{pr} \models X$ as an abbreviation of $X (S \text{ pr})$; thus, this fact can be expressed in Nuprl. A semantic property X is *satisfiable* if there is some program that satisfies it. Satisfiability can also be expressed in Nuprl: we take $\text{Sat}(X)$ to be an abbreviation for $\exists \text{pr} : \text{Pgm}. \text{pr} \models X$. The key point for the purposes of this paper is that from a

constructive proof of $\text{Sat}(X)$, we can *extract* a program that satisfies X .

For an instance of this general constructive framework to be useful *in practice*, the parameters Pgm , Sem , and S must be chosen so that (a) programs are concrete enough to be compiled, and (b) specifications are naturally expressed as predicates over Sem , and (c) there is a small set of *rules* for producing proofs of satisfiability. To use this general framework for synthesis of *distributed, asynchronous* algorithms, we choose the programs in Pgm to be *distributed message automata*.

Message automata are closely related to Lynch's *IO-Automata* [Lynch and Tuttle 1987; Lynch and Tuttle 1989] and are roughly equivalent to *UNITY* programs [Chandy and Misra 1988] (but with message-passing rather than shared-variable communication). We describe distributed message automata in Section 2.3. As we shall see, they satisfy criterion (a) above.

The semantics of a program is the *system*, or set of *runs*, consistent with it. Typical specifications in the literature are predicates on runs. We can view a specification as a predicate on systems by saying that a system satisfies a specification exactly if all the runs in the system satisfy it.

To satisfy criterion (b) above, we choose a formal definition of runs that builds in the fundamental order structure and provides the operators for appropriately abstract specifications. To do this we formalize runs as structures that we call *event structures*, much in the spirit of Lamport's [1978] model of events in distributed systems. Event structures are explained in more detail in the next section.

2.2 Event Structures

Following Lamport, we characterize a run of a program as a set of events. Each event is associated with a unique *agent* or *process* (we use the two words interchangeably in this paper). Formally, events are elements of a type E , and there is a function agent of type $E \rightarrow \text{AG}$, where AG is some set of agents. For each $i \in \text{AG}$, the set of events e such that $\text{agent}(e) = i$ is totally ordered. Intuitively, this set of events is the *history* of events at agent i . If $\text{first}(e)$ holds, then e is the first event in the history associated with $\text{agent}(e)$; if not, then e has a predecessor $\text{pred}(e)$.

Events are further partitioned by their *kinds*. The $\text{kind}(e)$ of event e is either $\text{rcv}(l)$ —a receive event on *link* l , or else $\text{local}(a)$ —a local event of kind a .¹ Every receive event has a *sender event* $\text{sender}(e)$. The sender event of a message is the event when the message was sent.

Following Lamport [1978], we can define a *causal order* on events as the transitive closure of the sender-receiver and predecessor relations. Thus, \rightarrow is the least relation on events such that $e \rightarrow e'$ if

- e' is a receive event and e is the corresponding send event,
- $\text{agent}(e) = \text{agent}(e)'$ and e precedes e' in the total order associated with $\text{agent}(e)$, or
- for some event e'' we have $e \rightarrow e''$ and $e'' \rightarrow e'$.

Intuitively, if $e \rightarrow e'$, then e is guaranteed to happen before e' .

The local state of an agent is represented as the values of a collection of state variables. Formally, state variables are just identifiers that are assigned a value at each event.² There are binary functions when and after that describe the values of state variables before and after an event takes place. We typically write these functions using infix notation. Thus, if $\text{agent}(e) = i$ then $(x \text{ when } e)$ describes the value of the state variable x at agent i just before e , and $(x \text{ after } e)$ describes its value after e . Note that state variables are local variables and two agents may have state variables with the same name.

¹Receive events are further partitioned by a *tag* so that we can restrict the kinds of events sending messages on a given link with a given tag without restricting other uses of that link. To simplify the discussion in this paper we have suppressed all mention of these tags.

²State variables are typed, but to simplify our discussion we suppress all type declarations.

Every event e also has a *value* $\text{val}(e)$. The value of a receive event is the message that is received, and the value of a local event represents a value (satisfying some constraints) chosen (non-deterministically) by the agent when the local event is generated. For example, if whenever a local event of kind a occurs the agent chooses an integer value and sends twice that value on link l , then events with $\text{kind}(e)=\text{rcv}(l)$ and $\text{kind}(\text{sender}(e))=\text{local}(a)$ will have $\text{val}(e)=2*\text{val}(\text{sender}(e))$.

The axioms of the event structure say that an event is the sender of only a finite number of messages; the predecessor function is one to one; causal order is well-founded; the local predecessor of an event has the same agent; the sender of an event has the agent of the source of the link on which the message was received; and the observation of state variables is related to the order structure in the obvious way, namely: $(x \text{ after } \text{pred}(e)) = (x \text{ when } e)$

2.3 Distributed message automata

Programs are built from a small set of basic clauses. With each basic clause c we associate a formula φ_c in the language of event structures, and the event structures consistent with c are the ones satisfying φ_c . If we prove a specification ψ using a set of assumptions $\{\varphi_c | c \in C\}$, then this set C of the clauses used in the proof will be a program realizing ψ . A finite set C of basic clauses is *feasible* if there is an event structure (a run) consistent with all the clauses in C (i.e. satisfying all the φ_c). Every basic clause is feasible and we have defined a syntactically checkable notion of *compatibility* that guarantees that the union of compatible feasible sets is again feasible.

Accordingly, a distributed message automaton is a finite set of basic clauses. Each basic clause has an agent, and as part of that agent it does one of six things:

1. defines the initial value of one state variable,
2. defines the effect of one kind of event on one state variable,
3. defines the messages sent on one link when events of one kind occur,
4. defines the precondition for one kind of local event,
5. lists all the kinds of event that affect one state variable,
6. lists all the kinds of event that send on a given link.

The first four basic clauses correspond to lines of code (and can easily be compiled into code). The last two basic clauses are called *frame conditions*, and correspond to promises *not* to add code. A frame condition and an effect clause or a send clause may be *incompatible*. We can form more complicated programs (i.e., automata) from simpler automata by composition. The composition $A \oplus B$ of two programs A and B is just the union of the clauses from A and B . The rules restrict composition of message automata to automata whose clauses are pairwise compatible. The class of distributed message automata (defined as a recursive type in Nuprl) is the smallest class containing the six basic clauses above and closed under \oplus .

The semantics of a distributed message automaton is the set of event structures that are consistent with it. This is formally defined as a relation $\text{Consistent}(R; es)$ between a program (i.e., message automaton) R and event structure es . Let Sys-R consist of all event structures consistent with R . A specification for us is a predicate on sets of runs. A program R is said to satisfy a specification if Sys-R does. As we show in the full paper, the relation $\text{Consistent}(R; es)$ is expressible in Nuprl. Programs are also expressible in Nuprl. Thus, we can talk about the system consistent with a program R in Nuprl.

Recall that a (standard) specification is a predicate on runs. A program R is said to satisfy a specification if every run in Sys-R does.

2.4 Realizability Rules

If P is a predicate on runs, R is a program, and the set of runs consistent with R is non-empty and every run consistent with R satisfies P , then we write $R \Vdash P$ and say that R *realizes* P . Note that $R \Vdash P$ is just an abbreviation for the formula $\forall es. (\text{Consistent}(R; es) \Rightarrow P(es)) \wedge \exists es. \text{Consistent}(R; es)$.

We have derived from the formal semantics of distributed message automata a set of nine rules for proving the realizability of a specification $P[es]$. There are six base rules, one for each basic clause, an additional rule for the combination of precondition and initialization clauses, a composition rule, and a refinement rule. We now briefly explain these rules.

The refinement rule says that if P refines Q (that is, if a run satisfies P , then it also satisfies Q) and A realizes P , then A also realizes Q :

$$A \Vdash P \Rightarrow (\forall es: ES. (P[es] \Rightarrow Q[es])) \Rightarrow A \Vdash Q .$$

The composition rule requires the notion of *compatibility*. Two programs A and B are *compatible*, denoted $A \parallel B$, if the sends and effect clauses of one obey the constraints imposed by the frame clauses of the other. The composition rule just captures the fact that $A \oplus B$ combines the constraints of A and B :

$$(A \Vdash P \wedge B \Vdash Q \wedge A \parallel B) \Rightarrow A \oplus B \Vdash P \wedge Q .$$

Each basic rule says that a basic clause c realizes its associated realizability rule φ_c . For the purposes of this paper we don't need the details of the syntax of the six basic clauses and the full listing of the six rules with all of their parameters. Instead, we give some examples.

The basic clause “at agent i initialize x to 5” realizes

$$\forall e@i. \neg \text{first}(e) \Rightarrow (x \text{ when } e) = 5,$$

where $\forall e@i. P$ is an abbreviation of $\forall e. \text{agent}(e)=i \Rightarrow P$.

A more important example for this paper is the precondition clause of the form “at agent i the precondition for a local action $a(v)$ is $P(v, x, \dots, z)$ ”. In this clause, the expression $P(v, x, \dots, z)$ is a predicate on the values of the state variables x, \dots, z and a value v . The intended meaning is that “infinitely often” agent i decides whether there is a value v satisfying predicate P in the current state, and if so the agent chooses some such v and performs a local action of kind a and value v .³ Also, an action of kind a may occur only when its value satisfies the precondition. The realizability rule associated with this precondition clause is the conjunction of a liveness property

$$\forall e@i. \exists e' \geq e. \text{kind}(e') = \text{local}(a) \vee \forall v. \neg P(v) \text{ after } e'$$

and a safety property

$$\forall e@i. \text{kind}(e) = \text{local}(a) \Rightarrow P@e,$$

where $P(v) \text{ after } e \equiv P(v, x \text{ after } e, \dots, z \text{ after } e)$ and $P@e \equiv P(\text{val}(e), x \text{ when } e, \dots, z \text{ when } e)$

2.5 Example

As an example of a parameterized specification that we will need later, consider the following predicate $\text{Fair}(P, f, l)$ on event structures, where P is a precondition, f is a function, and l is a link. $\text{Fair}(P, f, l)$ is a conjunction of a safety condition and a liveness condition. The safety condition asserts that every

³Performing an action of kind k means updating all the state variables in accordance with any effect clauses for k and sending all messages in accordance with any sends clauses for k .

receive event on link l has a value that is f of the state of the sender and that state satisfies the precondition P . The liveness condition says that “infinitely often” either a receive event on l occurs or else the precondition P fails.

$$\begin{aligned} \text{Fair}(P, f, l) \equiv & \\ & (\forall e'. \text{kind}(e') = \text{rcv}(l) \Rightarrow P @ \text{sender}(e') \wedge \text{val}(e') = f @ \text{sender}(e')) \\ & \wedge \forall e @ \text{source}(l). \exists e' \geq e. \text{kind}(e') = \text{rcv}(l) \vee \forall v : A. (\neg(P(v) \text{ after } e')) \end{aligned}$$

This specification is realized by the following program $\text{Fair-Pg}(P, f, l)$ consisting of the three basic clauses at agent $\text{source}(l)$:

```
precondition for a(v) is P
local(a)(v) sends [f v] on l
only events in [local(a)] send on l
```

We can prove that the program realizes the specification using the realizability rules. The intuitive reason is that the precondition is checked infinitely often. When it is true, the local event a occurs and sends the message. If this happens infinitely often, eventually a receive event will occur. Since only local action a sends this kind of message, the sender’s precondition must be the given one.

3 Adding knowledge to Nuprl

3.1 Consistent cut semantics for knowledge

To reason about knowledge in event structures we use a standard first-order modal logic of knowledge and time. Assume that there are n processes. Consider a propositional logic of knowledge, where formulas are formed by starting with a set Φ of function symbols, predicate symbols, and constant symbols of various arities. We form atomic predicates and terms as usual in first-order logic, and close under conjunction, negation, universal quantification, the temporal operator \Box , and the modal operators $K_i, i = 1, \dots, n$, one for each process i .⁴

Typically semantics for knowledge are given with respect to a pair (r, m) consisting of a run r and a time m , assumed to be the time on some external global clock (that none of the processes necessarily knows about) [Fagin et al. 1995]. In event structures, there is no external notion of time. Fortunately, Panangaden and Taylor [1992] give a variant of the standard definition with respect to what they call *asynchronous runs*, which are essentially identical to event structures. Thus, we just apply their definition in our framework.

The truth of formulas is defined relative to a triple $(\text{Sys}, \mathbb{E}, c)$, consisting of a system Sys (i.e., a set of event structures), and event structure \mathbb{E} in Sys , and a *consistent cut* c of \mathbb{E} , where a *consistent cut* c in \mathbb{E} is a set of events in \mathbb{E} closed under the causality relation. That is, if e' is an event in c and e is an event in \mathbb{E} that precedes e' (i.e., $e \rightarrow e'$), then e must also be in c .

Define the equivalence relations $\sim_i, i = 1, \dots, n$, on consistent cuts by taking $c \sim_i c'$ if i ’s history is the same in c and c' . Intuitively, $c \sim_i c'$ if process i cannot tell c and c' apart, given its information. Given two consistent cuts c and c' , we say that $c \preceq c'$ if, for each process i , process i ’s history in c is a prefix of process i ’s history in c' .

Given a nonempty set of objects D and a system Sys , an interpretation function π associates to each cut c and symbol s in Φ its interpretation, denoted $\pi(c, s)$, which is a predicate or function on D of the right arity. To extend this interpretation to terms, we start with a valuation V , which associates with each variable an element of D . For each variable x , we define $\pi(c, x) = V(x)$. We then define $\pi(c, f(t_1, \dots, t_k))$ by induction on the structure of terms, taking $\pi(c, f(t_1, \dots, t_k)) =$

⁴We can also define other standard modal operators, such as common knowledge, but we do not need them for the discussion in this paper. We can also define temporal operators such as *until*.

$\pi(c, f)(\pi(c, t_1), \dots, \pi(c, t_k))$. Using V and π , we define what it means for a formula φ to be true at the consistent cut \mathbf{c} in event structure \mathbf{E} in system \mathbf{Sys} , denoted $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models \varphi$, by induction on the structure of φ , in the usual way:

- if P is a predicate symbol in Φ of some arity k , and t_1, \dots, t_k are terms, then

$$(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models P(t_1, \dots, t_k) \text{ iff } \pi(c, P)(\pi(c, t_1), \dots, \pi(c, t_k))$$

- $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models \neg\varphi$ iff $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \not\models \varphi$
- $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models \varphi_1 \wedge \varphi_2$ iff $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models \varphi_1$ and $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models \varphi_2$
- $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models \forall x.\varphi$ iff, for all $d \in D$, $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V[x/d]) \models \varphi$, where $V[x/d]$ is the valuation that agrees with V on all variables except possible x , and $V[x/d](x) = d$
- $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models K_i\varphi$ iff for all $\mathbf{E}' \in \mathbf{Sys}$ and cuts \mathbf{c}' of \mathbf{E}' such that $\mathbf{c}' \sim_1 \mathbf{c}$, $(\mathbf{Sys}, \mathbf{E}', \mathbf{c}', \pi, V) \models \varphi$
- $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, \pi, V) \models \Box\varphi$ iff for all cuts \mathbf{c}' of \mathbf{E} such that $\mathbf{c} \preceq \mathbf{c}'$, $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}', \pi, V) \models \varphi$.

As usual, we take $\Diamond\varphi$ to be an abbreviation of $\neg\Box\neg\varphi$, so that $\Diamond\varphi$ is true at (\mathbf{E}, \mathbf{c}) if there is some cut \mathbf{c}' extending \mathbf{c} where φ is true.

Just as programs and the property of a program satisfying a specification can be expressed as a formula in Nuprl, the satisfaction relation \models_V can be expressed as a formula in Nuprl. More precisely, we would like to define a translation T such that for all tuples $(\mathbf{Sys}, \mathbf{E}, \mathbf{c})$, domains D , interpretations π , valuations V , and formulas φ , $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \varphi)$ is provable iff $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, V, \pi) \models \varphi$.

To present the translation, some notation is necessary. The type of systems (sets of event structures) is $\mathbf{Set}(\mathbf{ES})$ and the type of all consistent cuts in event structure \mathbf{E} is written as $\mathbf{CC}(\mathbf{E})$. The translation T is defined inductively on the structure of formula φ as follows:

- if P is a predicate symbol in Φ of some arity k , and $\mathbf{t}_1, \dots, \mathbf{t}_k$ are terms, then

$$T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, P(\mathbf{t}_1, \dots, \mathbf{t}_k)) = \pi(\mathbf{c}, P)(\pi(\mathbf{c}, \mathbf{t}_1), \dots, \pi(\mathbf{c}, \mathbf{t}_k))$$

- $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \neg\varphi) = \neg(T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \varphi))$
- $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, (\varphi_1 \wedge \varphi_2)) = (T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \varphi_1)) \wedge (T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \varphi_2))$
- $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, (\forall \mathbf{x}.\varphi)) = \forall \mathbf{x} : D. (T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, (\varphi(\mathbf{V} \mathbf{x})))$
- $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, K_i\varphi) = \forall \mathbf{E}' : \mathbf{Set}(\mathbf{ES}). \mathbf{E}' \in \mathbf{Sys} \Rightarrow \forall \mathbf{c}' : \mathbf{CC}(\mathbf{E}'). \mathbf{c}' \sim_1 \mathbf{c} \Rightarrow (T(\mathbf{Sys}, \mathbf{E}', \mathbf{c}', D, \pi, V, \varphi))$
- $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \Box\varphi) = \forall \mathbf{c}' : \mathbf{CC}(\mathbf{E}). \mathbf{c} \preceq \mathbf{c}' \Rightarrow (T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}', D, \pi, V, \varphi))$

As we said, we would now like to show that $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \varphi)$ is provable iff $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, V, \pi) \models \varphi$. However, since first-order epistemic logic relies on the principle of excluded middle, and Nuprl is a constructive type theory that does not, in general, assume the law of the excluded middle, we can prove that T has the desired effect only if we assume the law of the excluded middle. We remark that this assumption is necessary only for the purpose of this translation and not for the proofs we present in Section 4. With this assumption, we get the desired result.

Proposition 3.1: *Assuming the principle of excluded middle, for all tuples $(\mathbf{Sys}, \mathbf{E}, \mathbf{c})$, domains D , interpretations π , valuations V , and formulas φ , we have that $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \varphi)$ is provable iff $(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, V, \pi) \models \varphi$.*

To simplify the notation, we abbreviate $T(\mathbf{Sys}, \mathbf{E}, \mathbf{c}, D, \pi, V, \varphi)$ as $\varphi@^{\mathbf{Sys}, \mathbf{E}, D, \pi, V}$; we often abbreviate this further as $\varphi@c$, when the other components are clear from context.

3.2 Knowledge-based programs and specifications

In this section we show how we can extend the notions of program and specification presented in Section 2 to knowledge-based programs and specifications. This will allow us to employ the large body of tactics and libraries already developed in Nuprl to synthesize knowledge-based programs from knowledge-based specifications.

We have identified programs with distributed message automata, where a distributed message automaton is characterized by a set of clauses. We take a *knowledge-based message automaton* to be a function that associates to each system (i.e., set of event structures) a message automaton; intuitively, a knowledge-based message automaton allows preconditions on actions to depend on the knowledge of processes about the whole system. For the purposes of this paper, we take knowledge-based programs (hereafter abbreviated *kb programs*) to be knowledge-based message automata. Note that each standard program Pg corresponds to the kb program that associates to each system the program Pg .

What should the semantics of a kb program be? As discussed in Section 2, in the case of standard programs, a program semantics is a function of type $S \in Pgm \rightarrow Sem$; S associates with every program $Pg \in Pgm$ the system $S(Pg)$ consisting of all the runs consistent with Pg . (Recall that Sem is the type consisting of all systems.) As we have seen, the truth of a knowledge test in a kb program depends on the whole system. Once we have a system, we can determine the truth of the knowledge tests. A kb program then reduces to a standard program. Thus, a kb program has type $KbPgm = Sem \rightarrow Pgm$. Note that composing the semantic function S with a knowledge-based program yields a function from systems to systems. A system Sys is said to *represent* a kb program $kbPg$ if it is fixed point of this function. That is, Sys represents the kb program $kbPg$ if $S(kbPg)(Sys) = Sys$. Following Fagin et al. [1995, 1997], we take the semantics of a kb program $kbPg$ to be the set of systems that represent $kbPg$. That is, a kb program semantics S^{kb} is a function of type $KbPgm \rightarrow set-of-systems$, where *set-of-systems* is the type whose elements are sets of systems. As observed by Fagin et al. [1995, 1997], it is possible to construct kb programs that are represented by no systems, exactly one system, or more than one system. It is also possible to construct sufficient conditions (which are often satisfied in practice) that guarantee that a kb program is represented by exactly one system. Note that, in particular, standard programs when viewed as knowledge-based programs are represented by a unique system; indeed, $S^{kb}(Pg) = \{S(Pg)\}$. Thus, we can view S^{kb} as extending S .

We next consider knowledge-based specifications (hereafter abbreviated *kb specifications*). Recall that a (standard) specification is a predicate on runs. Following [Halpern 1999], we take a kb specification to be a predicate on systems. That is, a kb specification has type $Sem \Rightarrow P$. If kbX is a kb specification, then $kbX(Sys)$ holds if Sys satisfies the specification kbX .

As in [Fagin et al. 1995; Halpern 1999], we say that a kb program $kbPg$ satisfies the kb specification kbX if all the systems representing $kbPg$ satisfy X . Thus, we take $kbPg \models kbX$ as an abbreviation of $\forall Sys \in (S^{kb} kbPg). kbX(Sys)$.

Example 3.2: Recall that in Section 2 a specification $Fair(P, f, l)$ was considered that requires that, infinitely often, either a precondition P fails at the state of the source of some link or a message is received on the link; the message is constructed by applying the function f at the source of the link. The specification is satisfied by a standard program $Fair-Pg(P, f, l)$.

$Fair(P, f, l)$ can be generalized to a kb specification $Fair^{kb}(P^{kb}, f^{kb}, l)$, where, instead of using a precondition P and function f , we use a *knowledge-based* predicate P^{kb} and a *knowledge-based* function f^{kb} , both of which take a system as an extra argument (in addition to the other arguments of P and f). $Fair^{kb}(P^{kb}, f^{kb}, l)$ asserts that, in every run of the system, infinitely often either the kb precondition fails or a receive event with the value given by f^{kb} occurs on line l . $Fair^{kb}(P^{kb}, f^{kb}, l)$ is satisfied by a kb program $Fair-Pg^{kb}(P^{kb}, f^{kb}, l)$, which associates to each system Sys the program $Fair-Pg(P^{kb}(Sys), f^{kb}(Sys), l)$; in system Sys , a process following $Fair-Pg^{kb}(P^{kb}, f^{kb}, l)$ sends a message with value determined by $f^{kb}(Sys)$ exactly when predicate $P^{kb}(Sys)$ holds.

4 The sequence transmission problem

The main motivation for formalizing epistemic logic in event theory is that it allows us to write kb specifications in Nuprl, and then to synthesize kb programs. In this section, we consider one example of how this can be done. We show how a kb program that solves the *sequence transmission problem* (stp from now on) discussed by Halpern and Zuck [1992] can be synthesized from a kb specification of the problem. In fact, Halpern and Zuck consider two different kb programs that solve the problem; we show how both can be synthesized.

The stp involves a sender S who has an input tape with a (possibly infinite) sequence $X = [x_0, x_1, \dots]$ of bits, and wants to transmit X to a receiver R ; R must write this sequence on an output tape Y . A solution to the problem must satisfy two conditions:

1. (safety): at all times, the sequence Y of bits written by R is a prefix of X , and
2. (liveness): every bit x_k is eventually written by R on the output tape.

If messages cannot be lost, duplicated, reordered or corrupted, then S could simply send the bits in X to R in order; however, just as Halpern and Zuck, we are interested in solutions to the stp in contexts where communication is not reliable. Following Halpern and Zuck, we assume (a) that all corruptions are detectable and (b) a weak fairness condition: all messages sent infinitely often are eventually received.

The safety and liveness conditions above are run-based specifications. As argued in [Fagin et al. 1995], it is often better to think in terms of knowledge-based specifications for this problem. The real goal of the stp is to get the receiver to know the bits. That is, writing $K_R(x_i)$ as an abbreviation for $K_R(x_i = 0) \vee K_R(x_i = 1)$, we really want a knowledge-based liveness condition of the form $\forall i \diamond K_R(x_i)$. If we further require $\forall i \forall k \square (Y[i] = k \Rightarrow K_R(x_i = k))$ (that is, the receiver never sets the i th bit of the output to k unless it knows that the k bit of the sequence X is actually k) and $\forall i, j, k \square ((i \geq j \Rightarrow K_R(x_j)) \wedge K_R(x_i = k) \Rightarrow \diamond Y[i] = k)$ (that is, if the receiver knows the first k bits, then it eventually writes them out), then clearly the safety and liveness specifications will hold. The formula $\forall j. i \geq j \Rightarrow K_R(x_j)$ is abbreviated as $K_RX[0 \dots i]$.

Recall that a kb specification is a predicate on systems; the condition $\forall i, k. \square (Y[i] = k \Rightarrow K_R(x_i = k))$ is written in Nuprl as the kb-specification STP_1 that associates to each system \mathbf{Sys} a predicate that ensures that for any i and k , at any run E in the system and consistent cut c , the condition $(Y[i] = k) \Rightarrow K_R(\mathbf{x}_i = k)$ holds. Similarly, the specification $\forall i k \square (K_RX[0 \dots i] \wedge \wedge K_R(x_i = k) \Rightarrow \diamond Y[i])$ is translated into Nuprl as the following kb specification STP_2 , which requires, for each system \mathbf{Sys} , natural number i , bit k , run E of \mathbf{Sys} , and consistent cut c in E ,

$$(K_RX[0 \dots i] \Rightarrow K_R(\mathbf{x}_i = k)) @ c \Rightarrow \exists c' \succeq c. (Y[i] = k) @ c'$$

Our goal is to prove using Nuprl that $STP_1 \wedge STP_2$ is satisfiable. As we said before, the constructive proof will actually yield a kb program satisfying this specification. The proof that $STP_1 \wedge STP_2$ is satisfiable is carried out in Nuprl by refining it to subgoals that are, intuitively, easier to prove. The system can often suggest appropriate subgoals using its library of tactics. In addition, the system has a library of standard proofs that can be applied at appropriate times. We are extending the library so that it includes standard facts about knowledge. For example, the well-known *knowledge axiom* that whatever is known must be true ($K_i \varphi \Rightarrow \varphi$) becomes a short lemma in Nuprl that can be invoked in any subsequent proof and used for refining a goal φ to $K_i \varphi$. Similarly, the library includes the fact that, in event structures, there is *perfect recall* [Fagin et al. 1995]: processes do not forget *stable facts* (that is, facts which, once true, remain true), so the formula $K_i \varphi \Rightarrow \square K_i \varphi$ is valid if φ is stable. Including these results in the library facilitates reuse of code. These lemmas can be invoked repeatedly in the course of a proof.

Reusability of code can also be exploited at the level of kb specifications; having formally defined distributed programs in Nuprl, once we prove that a certain kb specification \mathbf{kbX} is satisfiable we also

construct a program Pg that satisfies kbX that can be manipulated and reused during the proof. Ingenuity is required to identify the specifications likely to appear in many proofs. But once good choices are made, they can significantly simplify proofs. As we shall see, the specification in Section 3.2 is invoked a number of times in the course of synthesizing a program that satisfies $\text{STP}_1 \wedge \text{STP}_2$.

Returning to our problem, we remark that performing a few standard Nuprl refinement steps allows us to reduce the proof of $\text{Sat}(\text{STP}_1 \wedge \text{STP}_2)$ to a few subgoals, the most interesting of which is that

$$K_{RX}[0 \dots k - 1] @c \Rightarrow \exists c' \succeq c. K_{RX}[0 \dots k] @c' \quad (1)$$

is satisfiable. (To simplify notation, we have omitted the universal quantification over runs E and cuts c in E .) This subgoal says that we want to find a program that ensures that R makes progress: if at some cut R knows the first k bits, then later on R knows all these bits, and the next one in the sequence.

The fact that R continues to know the bits it already knows is an instance of the rule we mentioned before regarding perfect recall in event structures. Thus, what we really have to do is find a program that ensures that, if R knows the first k bits, it will eventually know the $(k + 1)$ st bit.

Up to this point in the proof, we have used only standard facts about logic, standard refinement rules for first-order logic, and natural induction. All of this is built into Nuprl. At this stage in the refinement, ingenuity is required to find suitable rules specific to the particular kb specification of interest. It turns out that only one new tactic is needed; we are guided to it by the weak fairness condition we have imposed on communication. Observe that the fairness condition guarantees that if at S a message is sent infinitely often, then R will eventually get it.

Intuitively, R learns x_k from S ; we would like to be the case that infinitely often either there is no value for which S doesn't know that R knows it, or if such a value exists, then R receives it. This is just an instance of the kb specification $\text{Fair}^{\text{kb}}(\text{p}^{\text{kb}}, \text{f}^{\text{kb}}, 1)$ presented in example 3.2 and satisfied by the program denoted by $\text{Fair-Pg}(\text{p}^{\text{kb}}, \text{f}^{\text{kb}}, 1)$. We just need to choose the p^{kb} , f^{kb} , and 1 appropriate for our problem. Let P_S^{kb} be the kb predicate that tests whether $\neg K_S K_{RX}$ holds; let f_S^{kb} be the kb function that, given S 's state, if P_S^{kb} holds, returns the maximum index i such that $K_S K_{RX}[0 \dots i - 1] \wedge \neg K_S K_{RX}[i]$ holds; let 1 be the link between R and S . $\text{Fair}^{\text{kb}}(\text{P}_S^{\text{kb}}, \text{f}_S^{\text{kb}}, 1)$ is the kb specification that requires that infinitely often either $K_S K_{RX}$, or R receives $\langle i, x_i \rangle$.

Recall that $\text{Fair-Pg}^{\text{kb}}(\text{P}_S^{\text{kb}}, \text{f}_S^{\text{kb}}, 1) = \text{Fair}^{\text{kb}}(\text{P}_S^{\text{kb}}, \text{f}_S^{\text{kb}}, 1)$. If S uses the kb program $\text{Fair-Pg}^{\text{kb}}(\text{P}_S^{\text{kb}}, \text{f}_S^{\text{kb}}, 1)$, and S knows that R knows all the bits, then S sends no message; otherwise, S sends to R the maximum value i such that S knows that R knows the first i bits, but does not know that R knows x_i . Thus, S is essentially using the following kb program, also used in [Halpern and Zuck 1992] (which we write using their notation):

if $K_S K_{RX}[0 \dots i - 1] \wedge \neg K_S K_{RX}[i]$ **then** $\text{send}(\langle i, x_i \rangle)$ **else skip.**

It is clear that if S uses this kb program and does not know that R knows the k th bit, then S will send the k th bit repeatedly as long as S does not know that R knows the k th bit. Since a message is eventually received if sent infinitely often, R will eventually know the k th bit. Thus, we have reduced the problem to finding a way to guarantee that, if R knows the first i bits and does not know x_i , then S eventually knows that R knows the first i bits (and so will send x_i). This is again just an instance of the kb specification $\text{Fair}^{\text{kb}}(\text{p}^{\text{kb}}, \text{f}^{\text{kb}}, 1)$. Let P_R^{kb} be the kb predicate that tests whether there is a bit that R does not know; let f_R^{kb} be the kb function that, given R 's state that, if P_R^{kb} holds, returns the maximum index j such that R knows the first j bits in the sequence; again, 1 is the link between R and S . In other words, f_R^{kb} returns the (unique) value j such that $K_{RX}[0 \dots j - 1] \wedge \neg K_{RX}[j]$ holds in R 's state when interpreted with respect to system Sys . $\text{Fair}^{\text{kb}}(\text{P}_R^{\text{kb}}, \text{f}_R^{\text{kb}}, 1)$ is the kb specification that requires that infinitely often either P_R^{kb} fails, so the receiver knows all the bits, or S receives from R a message and S knows that f_R^{kb} holds at all points where R 's last action was to send this message. Now, when R follows $\text{Fair-Pg}^{\text{kb}}(\text{P}_R^{\text{kb}}, \text{f}_R^{\text{kb}}, 1)$, if R knows all the bits, then R sends no message; otherwise R sends to S the maximum value j such that R knows the first j bits but does not know x_j . Thus, R is essentially

using the following kb program used in [Halpern and Zuck 1992]:

if $K_{RX} [0 \dots j - 1] \wedge \neg K_{RX} [j]$ **then** send(j) **else** skip.

Recall goal (1); from $\text{Fair}^{\text{kb}}(\mathbf{p}_R^{\text{kb}}, \mathbf{f}_R^{\text{kb}}, 1)$ we infer that either $K_{RX} @ c$, which implies $K_{RX} [0 \dots k] @ c$, or there is some j such that R knows all the bits up to but excluding j . Since R knows all bits up to k , it must be that $j \geq k$; if $j > k$ then, since R knows all bits up to j , R also knows \mathbf{x}_k ; if $j = k$ then $\text{Fair}^{\text{kb}}(\mathbf{p}_R^{\text{kb}}, \mathbf{f}_R^{\text{kb}}, 1)$ assures us that there is some cut \tilde{c} extending c such that S receives k at \tilde{c} . At this moment it must be that S knows that R knows at least the bits up to k . From $\text{Fair}^{\text{kb}}(\mathbf{p}_S^{\text{kb}}, \mathbf{f}_S^{\text{kb}}, 1)$ we infer that either $K_S K_{RX} @ \tilde{c}$ holds, from which the knowledge axiom can be applied to deduce $K_{RX} @ \tilde{c}$, or there is some i such that S knows that R knows the bits up to i , but $\neg K_S K_{RX} [i] @ \tilde{c}$ holds. Thus, $i \geq k$; if $i > k$ then S knows that R knows \mathbf{x}_k , and so R knows \mathbf{x}_k . If $i = k$, then $\text{Fair}^{\text{kb}}(\mathbf{p}_S^{\text{kb}}, \mathbf{f}_S^{\text{kb}}, 1)$ assures us that there must be a later cut c' at which R receives $\langle k, \mathbf{x}_k \rangle$, and so R learns \mathbf{x}_k . From the perfect recall rule, we deduce that R knows all the bits up to and including \mathbf{x}_k , which completes our proof.

It is now straightforward to prove in Nuprl that if R uses the kb program $\text{Fair-Pg}^{\text{kb}}(\mathbf{p}_R^{\text{kb}}, \mathbf{f}_R^{\text{kb}}, 1)$, then S gets the requisite knowledge regarding which bit to send. The composition of the two programs

$$(\text{Fair-Pg}^{\text{kb}}(\mathbf{p}_S^{\text{kb}}, \mathbf{f}_S^{\text{kb}}, 1) \text{ A}) \oplus (\text{Fair-Pg}^{\text{kb}}(\mathbf{p}_R^{\text{kb}}, \mathbf{f}_R^{\text{kb}}, 1) \text{ A})$$

is thus essentially the kb program from [Halpern and Zuck 1992] that solves the stp. What we have shown here is that this program can be synthesized in Nuprl, using just standard facts in the Nuprl library, and the kb program Fair-Pg , which we will build into the Nuprl library, since it seems that it will be generally useful in all systems where communication is fair.

Halpern and Zuck show that a number of standard programs used in the literature are implementations of this kb programs. We can prove this in Nuprl, essentially by emulating the Halpern-Zuck proof. We omit the details in this abstract. Halpern and Zuck also present another kb program which is *receiver-driven*: rather than S sending the i th bit when it S does not know that R knows the bit, S sends the bit only when S knows that R does *not* know it. That is, the precondition of the sender's program has the test $K_S \neg K_{RX} [i]$ rather than $\neg K_S K_{RX} [i]$. In the full paper, we show how this kb program can be synthesized in Nuprl, using a relatively minor variant of the argument given above.

5 Conclusion

We have shown how to embed epistemic logic notions in constructive type theory. This allows us to express knowledge-based specifications and knowledge-based programs in Nuprl. We then showed by example how, from the constructive proof that a knowledge-based specification is satisfiable, we can extract a knowledge-based program that is guaranteed to satisfy that specification. What is notable about our proof is that it uses only standard logical arguments in Nuprl involving straightforward first-order logic and induction (as is the case for the synthesis of many sequential programs that have been carried out in Nuprl), standard arguments in epistemic logic (the knowledge axiom and the fact that in event structures agents have perfect recall), and one new argument involving the knowledge-based specification Fair^{kb} and the knowledge-based program $\text{Fair-Pg}^{\text{kb}}$ that implements it.

The structure of the argument gives us reason to believe that we will be able to incorporate into the Nuprl library standard tactics that will apply to a wide variety of knowledge-based specifications, and that we will be able to reuse these tactics in many proofs. We are currently in the process of synthesizing other knowledge-based programs to verify that, and to deepen our understanding of the techniques needed to knowledge-based verification. We believe that, once we have shown the power of this approach, the ideas will quickly spread to other verification approaches based on higher-order logic. We are optimistic about the prospects.

References

- Bates, J. L. and R. L. Constable (1985). Proofs as programs. *ACM Transactions on Programming Languages and Systems* 7(1), 53–71.
- Bickford, M. (2003). Experiments with theory modification in the FDL. *In Progress*.
- Bickford, M. and R. L. Constable (2003). A logic of events. Technical Report TR2003-1893, Cornell University.
- Bickford, M., C. Kreitz, R. van Renesse, and R. Constable (2001). An experiment in formal design using meta-properties. In J. Lala, D. Mughan, C. McCollum, and B. Witten (Eds.), *DARPA Information Survivability Conference and Exposition II (DISCEX-II)*, Volume II of *IEEE Computer Society Press*, Anaheim, CA, pp. 100–107.
- Bickford, M., C. Kreitz, R. van Renesse, and X. Liu (2001, September). Proving hybrid protocols correct. In R. Boulton and P. Jackson (Eds.), *14th International Conference on Theorem Proving in Higher Order Logics*, Volume 2152 of *Lecture Notes in Computer Science*, Edinburgh, Scotland, pp. 105–120. Springer-Verlag.
- Birman, K., R. Constable, M. Hayden, J. Hickey, C. Kreitz, R. van Renesse, O. Rodeh, and W. Vogels (2000). The Horus and Ensemble projects: Accomplishments and limitations. In *DARPA Information Survivability Conference and Exposition (DISCEX 2000)*, Hilton Head, SC, pp. 149–161. IEEE Computer Society Press.
- Chandy, K. M. and J. Misra (1988). *Parallel Program Design: A Foundation*. Reading, Mass.: Addison-Wesley.
- Constable, R. L. (2002). Naïve computational type theory. In H. Schwichtenberg and R. Steinbrüggen (Eds.), *Proof and System-Reliability, Proceedings of International Summer School Marktoberdorf, July 24 to August 5, 2001*, Volume 62 of *NATO Science Series III*, Amsterdam, pp. 213–260. Kluwer Academic Publishers.
- Constable, R. L. et al. (1986). *Implementing Mathematics with the Nuprl Proof Development System*. NJ: Prentice-Hall.
- Dwork, C. and Y. Moses (1990). Knowledge and common knowledge in a Byzantine environment: crash failures. *Information and Computation* 88(2), 156–186.
- Engelhardt, K., R. v. d. Meyden, and Y. Moses (1998). A program refinement framework supporting reasoning about knowledge and time. In J. Tiuryn (Ed.), *Proc. Foundations of Software Science and Computation Structures (FOSSACS 2000)*, pp. 114–129. Berlin/New York: Springer-Verlag.
- Engelhardt, K., R. v. d. Meyden, and Y. Moses (2001). A refinement theory that supports reasoning about knowledge and time for synchronous agents. In *Proc. Int. Conf. on Logic for Programming, Artificial Intelligence, and Reasoning*, pp. 125–141. Berlin/New York: Springer-Verlag.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1995). *Reasoning about Knowledge*. Cambridge, Mass.: MIT Press.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1997). Knowledge-based programs. *Distributed Computing* 10(4), 199–225.
- Geuvers, H., F. Wiedijk, and J. Zwanenburg (2001). A constructive proof of the fundamental theorem of algebra without using the rationals. In P. Callaghan, Z. Luo, J. McKinna, and R. Pollack (Eds.), *Types for Proofs and Programs, Proceedings of the International Workshop TYPES 2000*, Volume 2277 of *Lecture Notes in Computer Science*, pp. 96–111. Springer.
- Geuvers, J. H. et al. (2001). The ‘fundamental theorem of algebra’ project. FTA web page. See <http://www.cs.kun.nl/gi/projects/fta/>.

- Hadzilacos, V. (1987). A knowledge-theoretic analysis of atomic commitment protocols. In *Proc. 6th ACM Symp. on Principles of Database Systems*, pp. 129–134.
- Halpern, J. Y. (1999). Hypothetical knowledge and counterfactual reasoning. *International Journal of Game Theory* 28(3), 315–330.
- Halpern, J. Y. and Y. Moses (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM* 37(3), 549–587.
- Halpern, J. Y., Y. Moses, and O. Waarts (2001). A characterization of eventual Byzantine agreement. *SIAM Journal on Computing* 31(3), 838–865.
- Halpern, J. Y. and L. D. Zuck (1992). A little knowledge goes a long way: knowledge-based derivations and correctness proofs for a family of protocols. *Journal of the ACM* 39(3), 449–478.
- Kreitz, C. (1998). Program synthesis. In W. Bibel and P. Schmitt (Eds.), *Automated Deduction – A Basis for Applications*, Volume III, Chapter III.2.5, pp. 105–134. Kluwer.
- Lamport, L. (1978). Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM* 21(7), 558–565.
- Liu, X., C. Kreitz, R. van Renesse, J. J. Hickey, M. Hayden, K. Birman, and R. Constable (1999, December). Building reliable, high-performance communication systems from components. In *17th ACM Symposium on Operating Systems Principles (SOSP'99)*, Volume 33(5) of *Operating Systems Review*, pp. 80–92. ACM Press.
- Liu, X., R. van Renesse, M. Bickford, C. Kreitz, and R. Constable (2001). Protocol switching: Exploiting meta-properties. In L. Rodrigues and M. Raynal (Eds.), *International Workshop on Applied Reliable Group Communication (WARGC 2001)*, pp. 37–42. IEEE.
- Lynch, N. A. and M. R. Tuttle (1987). Hierarchical correctness proofs for distributed algorithms. Technical Report MIT/LCS/TR-387, MIT.
- Lynch, N. A. and M. R. Tuttle (1989). An introduction to input/output automata. *CWI Quarterly* 2(3), 219–246. Also available as MIT Technical Memo MIT/LCS/TM-373.
- Mazer, M. S. (1990). A link between knowledge and communication in faulty distributed systems. In *Theoretical Aspects of Reasoning about Knowledge: Proc. Third Conference*, pp. 289–304.
- Mazer, M. S. and F. H. Lochovsky (1990). Analyzing distributed commitment by reasoning about knowledge. Technical Report CRL 90/10, DEC-CRL.
- Moses, Y. and M. R. Tuttle (1988). Programming simultaneous actions using common knowledge. *Algorithmica* 3, 121–169.
- Neiger, G. and S. Toueg (1993). Simulating real-time clocks and common knowledge in distributed systems. *Journal of the ACM* 40(2), 334–367.
- Panangaden, P. and S. Taylor (1992). Concurrent common knowledge: defining agreement for asynchronous systems. *Distributed Computing* 6(2), 73–93.
- Paulin-Mohring, C. and B. Werner (1993). Synthesis of ML programs in the system Coq. *Journal of Symbolic Computations* 15, 607–640.
- Smith, D. R. and C. Green (1996). Toward practical applications of software synthesis. In *FMSP'96, The First Workshop on Formal Methods in Software Practice*, pp. 31–39.
- Stulp, F. and R. Verbrugge (2002). A knowledge-based algorithm for the Internet protocol (TCP). *Bulletin of Economic Research* 54(1), 69–94.