THE UNIVERSITY OF MICHIGAN

COMPUTING RESEARCH LABORATORY[1]

Connections Between Two Theories of Concurrency: Metric Spaces and Synchronization Trees

W. G. Golson and W. C. Rounds

CRL-TR-3-83

JANUARY 1983

Computing Research Laboratory
Room 1079, East Engineering Building
Ann Arbor, Michigan 48109
USA
Tel: (313) 763-8000

[1]All correspondence should be sent to Professor W.C. Rounds. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies.

Connections Between Two Theories of Concurrency:

Metric Spaces and Synchronization Trees

October 1982

W.G. Golson

W.C. Rounds

Department of Computer and Communication Sciences
University of Michigan
Ann Arbor, Michigan 48109

Abstract

We establish a connection between the semantic theories of concurrency and communication in the works of de Bakker and Zucker, who develop a denotational semantics of concurrency using metric spaces instead of complete partial orders, and Milner, who develops an algebraic semantics of communication based upon observational equivalence between processes. We endow his _rigid synchronization trees_ (RSTs) with a simple pseudometric distance induced by Milner's _weak equivalence_ relation and show the quotient space to be complete. We establish an isometry between our space and the solution to a domain equation of de Bakker and Zucker, presenting the solution in a conceptually simpler framework. Under an additional assumption, we establish the equivalence of the weak equivalence relation over RSTs and the elementary equivalence relation induced by the sentences of a modal logic due to Hennessy and Milner.

0. Introduction

In this paper we establish a fundamental connection between the semantic theories of concurrency and communication in the works of de Bakker and Zucker [BaZ] and Milner [Mil]. In [BaZ] de Bakker and Zucker develop a denotational semantics of concurrency using metric spaces (see for example [Niv] or [ArN]) instead of complete partial orders as the underlying mathematical structures. They solve several reflexive domain equations, and the solutions of two equations in particular, involving nondeterministic processes, entail the abstract completion of a metric space recursively constructed from metric spaces which utilize a Hausdorff distance between closed sets. Milner develops an algebraic semantics of communication based upon behavioral or observational equivalence between processes. We take his rigid synchronization trees (RSTs), with countable branching and arc labels from an arbitrary alphabet, and endow them with a simple pseudometric distance induced by Milner's weak observational equivalence relation to construct a concrete representation of the solution to the first domain equation above. We prove that our quotient space is complete under the corresponding metric distance, and show that it is isometric to the de Bakker-Zucker completion by identifying an appropriate dense subset. As a result, one does not necessarily have to use the complicated notions of Hausdorff distance and the attendant machinery of metric space completions; one can work directly with trees as graphs and use a simple metric defined directly on the graph structure.

The structure of our metric space has additional properties of interest. For example, unlike Milner we cannot restrict ourselves to finitely branching trees, since infinitely branching trees are necessary for the completion of the metric space. This need for unbounded branching arises quite naturally, a development which we are pleased to see. In another vein, while the construction in this paper allows the alphabet Σ to be infinite, we can prove that our metric space is compact if and only if Σ is finite. In this case it turns out that the weak observational equivalence relation is exactly the elementary equivalence relation induced by the sentences of a simple modal logic due to Hennessy and Milner [HeM]. The statement that our space is compact is exactly the assertion of the Compactness Theorem for the Hennessy-Milner logic (HML). Since the HML compactness theorem follows from a direct translation into first order logic, this gives us an elegant but nonconstructive proof of completeness for the case when Σ is finite. On the other hand, since our proof of the metric space compactness is constructive, the HML compactness theorem is true without the axiom of choice.

The rest of the paper is organized as follows. Section 1 is preliminary, defining the domain of trees and establishing some necessary properties. Section 2 presents the rigid synchronization trees of Milner and defines weak equivalence. The third section constructs the metric space and proves its completeness. The fourth section recalls the necessary definitions and results from [BaZ] and establishes the isometry between the metric spaces of this paper and [BaZ]. Finally section 5 establishes the connections between HML and our metric space.

1. Preliminaries

We regard a tree as a directed, unordered graph on a countable set of nodes with arcs labeled from an alphabet Σ. The graph must have the obvious tree shape and two arcs leaving the same node may have the same label. More formally we define the set of trees, T, as follows:

Definition: S is a tree ($S \varepsilon T$) iff S is a 4-tuple $S=(V,E,\ell,v_0)$

where V is a set of vertices or nodes;

$v_0 \varepsilon V$ is the root;

$E \subseteq V \times V$ is the edge relation, antisymetric and irreflexive;

$\ell : E \to \Sigma$ assigns a label to each edge.

In addition the following properties are satisfied:

(1) all nodes are reachable from the root:

$\forall v \varepsilon V - \{v_0\}$ $\langle v_0, v\rangle \varepsilon E^+$ where E^+ is the transitive closure of E;

(2) each node has only one ancestor:

$\forall u,v,w \varepsilon V$, $\langle u,w\rangle \varepsilon E$ and $\langle v,w\rangle \varepsilon E$ implies $u=v$.

We say two trees are isomorphic if both can be transformed into the other preserving structure and labeling:

Definition: $S=(V,E,\ell,v_0)$ and $S'=(V',E',\ell',v_0')$ are isomorphic iff there is a bijection $f:V \to V'$ such that

(1) $f(v_0)=v_0'$ (identification of roots);

(2) $\langle v,w\rangle \varepsilon E \iff \langle f(v),f(w)\rangle \varepsilon E'$ (identification of edges);

(3) $\forall \langle v,w\rangle \varepsilon E$, $\ell(\langle v,w\rangle)=\ell'(\langle f(v),f(w)\rangle)$ (identified edges have same label).

For example, 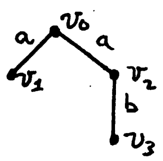 and 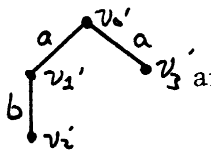 are isomorphic under the correspondence $v_0 \to v_0'$, $v_1 \to v_3'$, $v_2 \to v_1'$, $v_3 \to v_2'$. When S and S' are isomorphic, we shall write $S=S'$.

The notions of path, path length, and finite and infinite paths are the usual ones. We say a tree is <u>bounded</u> if there is a finite bound on all path lengths. A node is <u>finitely branching</u> if it has a finite number of direct descendants. A tree is finitely branching if all its nodes are. We allow <u>countable branching</u> at any node.

The k^{th} <u>cross section</u> $S^{(k)}$ of a tree S is just S restricted so that no path has a length exceeding k:

<u>Definition</u>: For $S \in \mathcal{T}$, let the k^{th} <u>cross</u> section of $S=(V,E,\ell,v_0)$ be:

$S^{(0)}=(\{v_0\},\emptyset,\emptyset,v_0)$, $k=0$;

$S^{(k)}=(V_k,E_k,\ell_k,v_0)$, $k \geq 1$;

where $E_k=\{\langle v,w\rangle \in E \mid \text{the path } \langle v_0,w\rangle \text{ has length at most } k\}$;

$V_k=V|E_k$;

$\ell_k=\ell|E_k$.

<u>Examples</u>:

1) $S^{(0)}$ is just the root, which we call <u>nil</u>.

2) If S= 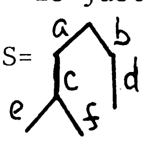

then $S^{(0)}$=nil, $S^{(1)}$= 

, $S^{(2)}$= 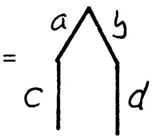

, $S^{(k)}=S$ for $k \geq 3$.

We have the following relationship between a tree and its cross sections:

<u>Lemma 1.1</u>: For any $S \in \mathcal{T}$, let $\{S^{(k)}\}$ be the set of all its cross sections, $k \geq 0$.

Then (a) $\forall k \geq 0$ $E_k \subseteq E_{k+1}$ and $E=\bigcup E_k$

(b) $\forall k \geq 0$ $V_k \subseteq V_{k+1}$ and $V=\bigcup V_k$

(c) $\forall k \geq 0$ $\ell_k \subseteq \ell_{k+1}$ and $\ell=\bigcup \ell_k$ (viewing ℓ_k as a set of ordered pairs $\langle e_k,a\rangle$ from E_k and Σ)

<u>Proof</u>: We prove (a). $E_k \subseteq E_{k+1}$ directly from the definition. Now clearly $E_k \subsetneq E$ for all k so $\bigcup E_k \subsetneq E$.

Let $<v,w>\varepsilon E$. Then there is a path $<v_0,w>$ and therefore $<v,w>\varepsilon E_k$ for any k not less than the path length of $<v_0,w>$. Therefore $<v,w>\varepsilon \cup E_k$, whereby $E \subseteq \cup E_k$.

This lemma suggests that any tree can be represented as a union of its cross sections, leading to the following definitions:

<u>Definition</u>: Let $\{S_k\} \subseteq T$. $\{S_k\}$ is a <u>cross sectional sequence</u> (written $<S_k>$ a XSS) iff (1) each S_k is bounded, say with maximum path length of b(k);

(2) $\forall m \geq k \quad S_m{}^{b(k)} = S_k{}^{b(k)}$ (writing $S^{b(k)}$ for $S^{(b(k))}$)

The last condition insures that the b(k)-th cross sections of $S_k, S_{k+1} \ldots$ are all equal, i.e. only the leaves of S_k with path length b(k) can be extended to form S_{k+1}. For convenience, in any sequence $<S_k>$, we shall take S_0 to be the <u>nil</u> tree and b(0)=0.

<u>Definition</u>: Let $<S_k>$ be a XSS. The <u>union tree</u> of $<S_k>$ is

$$\cup S_k = (\cup V_k, \cup E_k, \cup \ell_k, v_0).$$

We collect some facts about XSS which will be useful later:

<u>Lemma 1.2</u>: Let $<S_k>$ be a XSS.

(a) $k \leq n$ implies $b(k) \leq b(n)$

(b) $S_k = S_k{}^{b(k)}$

(c) $\forall m \geq k \ \forall j \leq b(k) \ S_m{}^{(j)} = S_k{}^{(j)}$

(d) $\cup S_k$ is a tree and $(\cup S_k)^{b(k)} = S_k$

<u>Proof</u>: omitted.

We wish to define two additional operators on trees, prefixing and joining, enabling us to create complex trees from simpler ones.

<u>Notation</u>: S[v/w] means the tree S with the node w replaced by v.

efinition: For $S=(V,E,\ell,v_0)$ and $a\varepsilon\Sigma$

let $aS=\{V\cup\{v_a\},E\cup\{<v_a,v_0>\},\ell\cup\{<<v_a,v_0>,a>\},v_a\}$

where $v_a \notin V$.

ictorially we have  . We call aS a prefixed (sub)tree.

efinition: We say $\{S_k\}$ are disjoint if $\{V_k\}$ are pairwise disjoint.

efinition: Let $\{S_k\}\subseteq\mathcal{T}$, $S_k=(V_k,E_k,\ell_k,v_{0,k})$, $\{S_k\}$ disjoint. The join of $\{S_k\}$ is

$$\oplus S_k = \cup S_k[v_0/v_{0,k}]$$

o $S\oplus T$ becomes the tree  . We view the expression $S\oplus S$ to be well efined, representing the joining of two disjoint isomorphic copies of S. We epresent by S^n the joining of n copies of S for $1\leq n\leq\omega$. In a similar spirit, $\oplus T$ will always be taken to be well defined through an inessential relabeling of odes if necessary.

emma 1.3: aS and $\oplus S_k$ are trees.

roof: clear.

Finally we establish another representation of an arbitrary tree:

emma 1.5: For $S\varepsilon\mathcal{T}$, there is a set $\{a_iS_i\}\subseteq\mathcal{T}$ such that $S=\oplus a_iS_i$

roof: Clearly we can represent S as the join of its prefixed subtrees.

2. Rigid Synchronization Trees and Weak Equivalence

In the spirit of [Mil] we regard a rigid synchronization tree (RST) as the "unfolding" of a state transition graph of a nondeterministic machine. For example, given the graph 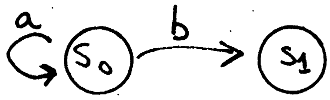 we associate the RST: 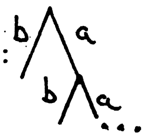

Note that that state names are no longer important; the tree nodes are nameless. The arc labels are chosen from an event alphabet Σ, reflecting the communication requirements of the process from its environment. We depart from [Mil] and allow the nodes to have countable branching.

Nondeterministic choice exists in the tree 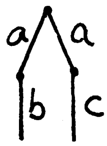 . Given an "a", the machine must choose between two paths, arriving at either a state where only a "b" is acceptable or one in which only a "c" is. Now consider the tree 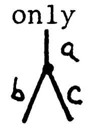

If viewed as acceptors, both of these trees are equivalent, accepting the language {ab,ac}. But are they equivalent behaviorally? After one step the second tree is in a state where either a "b" or "c" is acceptable, and so it never deadlocks on input from {ab,ac}. However, the first tree can deadlock on either "ab" or "ac": after "a" has been consumed, it will be in a state waiting for a specific event and will fail if the environment offers an incompatible input. Note that nondeterministic trees do not necessarily "choose correctly"; they react only to the current event, not to future ones. Since the trees behave differently on inputs from {ab,ac}, it is reasonable to maintain that they are not equivalent behaviorally.

Several different equivalence relations have been proposed to describe behavioral or observational equivalence [Mil]. The relation appropriate for this paper is the weak equivalence relation and is defined as follows.

Notation: When we write $S \xrightarrow{a} T$ we mean there is some a-transition from the root of S leading to T, or that aT is a prefixed subtree of S.

<u>Definition</u>: For $S,T\varepsilon\mathcal{T}$, S is weakly equivalent to T, $S \equiv_w T$, iff $\forall_k\ S \equiv_k T$, where the equivalences $\equiv_k$ are defined as:

$S \equiv_0 T$ for all S,T;

$S \equiv_{k+1} T \iff \forall a\varepsilon\Sigma\forall S'\varepsilon\mathcal{T},\ S \xrightarrow{a} S' \implies \exists T'\varepsilon\mathcal{T}\ \ T \xrightarrow{a} T'$ and $S' \equiv_k T'$ and

$\forall a\varepsilon\Sigma\forall T'\varepsilon\mathcal{T},\ T \xrightarrow{a} T' \implies \exists S'\varepsilon\mathcal{T}\ \ S \xrightarrow{a} S'$ and $T' \equiv_k S'$.

We write $S \equiv T$ for $S \equiv_w T$.

An alternate way of presenting k+1-equivalence which we shall find convenient is the following:

$S\equiv_{k+1}T \iff$ for every prefixed subtree aS' of S, there is a prefixed subtree aT' of T such that $S'\equiv_k T'$ (and vice versa).

<u>Examples</u>:

(1) 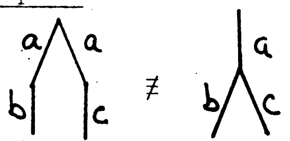 since they are $\not\equiv_2$. To see this note that nodes are $\equiv_1$ if the set of events which can occur next are the same. The tree  is $\not\equiv_1$ to either |b or |c.

(2) Let A_k be the tree 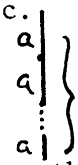 k times. Let $A_* = \oplus A_k$, $k\geq 1$. So A_* has arbitrarily long finite paths and an infinitely branching root: 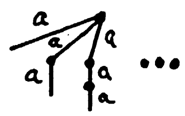

Let A_ω be the infinite tree  and let $A_\infty = A_* \oplus A_\omega$. Note that for all k, $A_\infty^{(k)} = A_*^{(k)}$ as each k-th cross section contains one path each of lengths 1,...,k-1 and a countable number of paths of length k. We claim that $A_\infty \equiv_k A_*$ for all k and thus $A_\infty \equiv A_*$, as can be seen from the following lemma:

<u>Lemma 2.1</u>: If $S^{(k)}=T^{(k)}$ then $S\equiv_k T$.

<u>Proof</u>: Induction on k.

For k=0 the result is immediate.

Assume the lemma holds for k.

Suppose now $S^{(k+1)}=T^{(k+1)}$. As the prefixed subtrees of S and T are in 1-1 correspondence, we can write $S^{(k+1)}=\bigoplus a_i S_i^{(k)} = \bigoplus a_i T_i^{(k)} = T^{(k+1)}$ where $S_i^{(k)}=T_i^{(k)}$. Therefore by the induction hypothesis we have $S_i \equiv_k T_i$. Clearly now we have $S \equiv_{k+1} T$.

We remark that the converse is false:  $\equiv$ 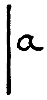 but not equal.

Finally we collect some easy and useful facts:

<u>Lemma 2.2</u>: (1) $S \equiv_k T$ implies $\forall j \leq k \quad S \equiv_j T$

(2) $S \not\equiv_k T$ implies $\forall j \geq k \quad S \not\equiv_j T$

(3) $S \equiv_k S^{(k)} \equiv_k S^{(n)} \quad n \geq k$

<u>Proof</u>: omitted.

3. The Metric Space of RSTs

In this section the completeness of the metric space on $\mathcal{T}$ induced by the weak equivalence relation is demonstrated. For topological definitions and related items, the reader is refered to [Dug].

We define the following metric on $\mathcal{T}$:

Definition: For $S,T \varepsilon \mathcal{T}$, let $d_w(S,T) = 2^{-k}$ where $k = \max_j S \equiv_j T$. If the maximum does not exist, we take k to be infinite.

As k-equivalence examines no nodes which are along paths of length greater than k from the root, we see that the larger the value of k above, the more alike the two trees are, the smaller the value of d_w.

Examples: $d_w($

$, | a) = 1$ since $S \not\equiv_1 T$

$d_w($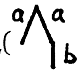

$, | a) = \frac{1}{2}$ since $S \equiv_1 T$ but $S \not\equiv_2 T$

$d_w($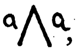

$, | a) = d_w(A_*, A_\infty) = 0$

Lemma 3.1: $\langle \mathcal{T}, d_w \rangle$ is a ultra pseudo metric space.

Proof: (1) $d_w(S,T)=0 \iff \forall k\ S \equiv_k T \iff S \equiv T$ (pseudo)

(2) $d_w(S,T) = d_w(T,S)$

(3) $d_w(S,T) \leq \max(d_w(S,U), d_w(U,T))$ (ultra)

Let $d_w(S,T) = 2^{-k}$ and suppose (wlog) $d_w(S,U) < 2^{-k}$. Then $S \equiv_{k+1} U$. Since both $S \equiv_k U$ and $S \equiv_k T$, we have $U \equiv_k T$. However, $U \not\equiv_{k+1} T$ as $S \not\equiv_{k+1} T$. Therefore $d_w(U,T) = 2^{-k}$.

We define the notions of Cauchy sequence and limit:

Definition: $\langle S_n \rangle$ is a Cauchy sequence (CS) iff

$$\forall k \geq 0\ \exists N \forall m,n \geq N,\ S_m \equiv_k S_n .$$

<u>Definition</u>: S is a <u>limit</u> of a CS $\langle S_n \rangle$ (written $S \varepsilon \llcorner \lim S_n \lrcorner$)

iff $\forall k \geq 0 \, \exists N \forall n \geq N, \; S \equiv_k S_n$.

<u>Remarks</u>: (1) The above definitions are equivalent to the more usual presentations, e.g., $\forall \varepsilon > 0 \, \exists N \forall m,n \geq N, \; d_w(S_m, S_n) < \varepsilon$.

(2) We must deal with equivalence classes of CS limits. Recall that $\langle T, d_w \rangle$ is a pseudometric space, and for example, if $S_n = \bigoplus_{j=1}^{n} A_j$,

e.g., $S_3 =$ 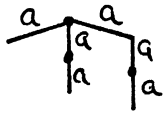 , we have that $\langle S_n \rangle$ is a CS and so for all n

$A_* \equiv_n S_n \equiv_n A_\infty$, and therefore $\{A_*, A_\infty\} \subseteq \llcorner \lim S_n \lrcorner$.

Proceeding to the completeness proof, we will establish that any XSS $\langle S_n \rangle$ in $\langle T, d_w \rangle$ is a CS with a well defined constructible limit, the union tree, $US_n \varepsilon \llcorner \lim S_n \lrcorner$. An operator on trees, C, yielding a fully expanded countably branching tree in a sense made precise below, will be defined and shown to possess the following special properties:

(1) weak equivalence is the same is isomorphism, i.e.

$C(S) = C(T) \iff C(S) \equiv C(T)$, for bounded S,T;

(2) for any bounded S, $S \equiv C(S)$.

By standard argument, given a CS $\langle S_n \rangle$, we can select a subsequence $\langle S'_n \rangle$ such that $\langle {S'_n}^{(n)} \rangle$ has the same limit as $\langle S_n \rangle$, if indeed such a limit exists. Now since $\langle C({S'_n}^{(n)}) \rangle$ is a XSS (implied by (1)) and therefore has a limit which by (2) is the same as $\langle S_n \rangle$, the completeness of $\langle T, d_w \rangle$ will follow directly.

<u>Lemma 3.2</u>: If $\langle S_n \rangle$ is a XSS, then it is also a CS in $\langle T, d_w \rangle$.

<u>Proof</u>: Recall $\langle S_n \rangle$ a CS $\iff \forall k \geq 0 \, \exists N \forall m,n \geq N, \; S_m \equiv_k S_n$

We have two cases:

(a) $\langle S_n \rangle$ is bounded (i.e. $\{b(n)\}$ is bounded). Then after some N_0, $\forall m,n \geq N_0, \; S_m = S_n$. Then for any $k, S_m^{(k)} = S_n^{(k)}$ and so $S_m \equiv_k S_n$ (Lemma 2.1).

(b) $\langle S_n \rangle$ is not bounded. Choose N such that $b(N) \geq k$. Then as $S_N = S_N^{b(N)}$ (Lemma 1.2.b) we have $\forall m,n \geq N \; S_m^{(k)} = S_n^{(k)}$ (Lemma 1.2.c) and therefore $S_m \equiv_k S_n$ (Lemma 2.1).

Theorem 3.3: Let $\langle S_n \rangle$ be a XSS. Then $\lfloor \lim S_n \rfloor$ exists and $US_n \in \lfloor \lim S_n \rfloor$.

Proof: US_n exists by Lemma 1.2.d. The reader may now proceed in a fashion similar to the proof of Lemma 3.2.

Our C operator is defined as

Definition: For any bounded tree S, let $C(S)$ be:

$$C(\text{nil}) = \text{nil}$$

$$C(\oplus a_i S_i) = [\oplus a_i C(S_i)]^{\omega}.$$

To aid the intuition, $C(S)$ can be constructed for any bounded tree S as follows:

(1) mark all leaf nodes as ready;

(2) repeat until the root is marked ready

if all of a node's descendants are ready

then replace each prefixed subtree of the node

by ω copies of the subtree and

mark the node ready;

For example, if S = 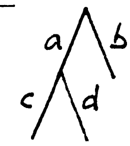

, then $C(S)$ = 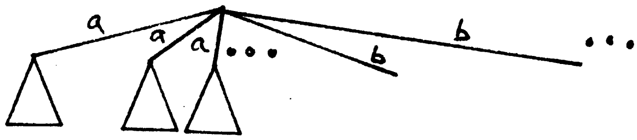

where  = 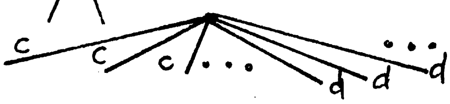

Lemma 3.4: For S bounded, $C(S)$ is a tree.

Proof: omitted.

The utility of C-trees becomes evident in the theorem and corollary below, in which weak equivalence is seen to be the same as isomorphism.

Theorem 3.5: Let $C=\mathcal{C}(S)$ and $D=\mathcal{C}(T)$ for some bounded S,T.

Then $C\equiv_k D \iff C^{(k)}=D^{(k)}$

Proof: (<==) Lemma 2.1.

(==>) Induction on k:

case k=0: immediate.

Assume for k.

case k+1:

Suppose $C\equiv_{k+1}D$. Partition the prefixed subtrees of both C and D into k+1-equivalence classes. As $C\equiv_{k+1}D$, these equivalence classes of C and D are in 1-1 correspondence. By the induction hypothesis, the representatives of corresponding classes have isomorphic k+1 cross sections, so the two trees obtained by the joining of the representatives are k+1-isomorphic.

As C and D are $\mathcal{C}$-trees, each equivalence class represents at most ω prefixed subtrees. Now as every prefixed subtree of C or D contributes ω copies of itself to C or D, the number of subtrees represented by any class is ω. Therefore we have $C^{(k+1)}=D^{(k+1)}$.

Corollary: $C\equiv D \iff C=D$ for bounded C,D.

Proof: As C,D bounded, the isomorphisms constructed above will stabilize.

The last result we need prior to priving completeness is the following:

Lemma 3.6: For S bounded, $S\equiv\mathcal{C}(S)$.

Proof: We show $\forall k\ S\equiv_k \mathcal{C}(S)$ by induction on k.

case k=0: immediate.

Assume for k.

case k+1:

Let $S=\bigoplus a_iS_i$, $\mathcal{C}(S)=[\bigoplus a_i\mathcal{C}(S_i)]^{\omega}$

Now $S\equiv_{k+1}\mathcal{C}(S) \iff \forall a \forall S'\ S\xrightarrow{a}S'$ implies $\exists C'\ \mathcal{C}(S)\xrightarrow{a}C'$ and $S'\equiv_k C'$ and vice versa.

If a_iS_i is a prefixed subtree of S, then $a_iC(S_i)$ is a prefixed subtree of $C(S)$. We have $S_i \equiv_k C(S_i)$ by the induction hypothesis and so the required C' exists. A similar argument for the reverse direction establishes the lemma.

We are now ready to prove

Theorem 3.7: $<T, d_w>$ is complete.

Proof: Let $<S_n>$ be any arbitrary CS in $<T,d_w>$,

$$\text{i.e. } \forall k \geq 0 \ \exists N \ \forall m,n \geq N \ S_m \equiv_k S_n.$$

By passing to a subsequence if necessary, we can assume $\forall n \geq k \ \ S_k \equiv_k S_n$. Consider now the sequence $<S_k^{(k)}>$. Clearly $<S_k^{(k)}>$ is a CS as $S_k^{(k)} \equiv_k S_n^{(n)} \forall n \geq k$. Since $S_k^{(k)}$ is bounded, $S_k^{(k)} \equiv C(S_k^{(k)})$ by Lemma 3.6. Therefore $<S_k^{(k)}>$ has a limit iff $<C(S_k^{(k)})>$ does. But $<C(S_k^{(k)})>$ is a XSS (by Theorem 3.5) and has a limit (Theorem 3.3).

Finally we observe that by construction $<S_k>$ has the same limit as $<S_k^{(k)}>$, completing the proof of the theorem.

At this point we would like to remark that our construction not only incorporates countably branching trees, but requires them for our space to be complete. That arbitrary finite branching is not enough can be seen from the following. Recall that A_j is the tree $\left.\begin{array}{c} a| \\ a| \\ \vdots \\ a| \end{array}\right\}$ j times. Now suppose that $S \equiv_k A_j$ for $j<k$. Then both the minimum and maximum path lengths in S have size j, so that all the paths in S have length j.

Now suppose $S \equiv_{k+1} A_*$, where we now write $A_* = \bigoplus aA_j$ for j a natural number. Then for all $j \leq k$ there is a prefixed subtree aS_j of S such that $A_j \equiv_k S_j$. Therefore, for each $j<k$, S_j has path lengths of exactly j and S_k has path lengths of at least k. So we have established that

Lemma 3.8: If $S \equiv_{k+1} A_*$, then S has at least a k-way branching root.

Theorem 3.9: $<T,d_w>$ is incomplete if trees cannot have countable branching.

Proof: $<A_1, A_1 \oplus A_2, \ldots>$ is a finitely branching CS with limit A_*, which by the lemma is not equivalent to any finitely branching tree.

4. An Isometry with a Metric Space of de Bakker and Zucker

En route to their denotational semantics of concurrency, de Bakker and Zucker [BaZ] wish to find a metric space $\langle P,d_B\rangle$ which solves

$$P=\{p_0\}\cup\mathcal{P}_c(\Sigma \times P) \qquad (4.1)$$

where $\mathcal{P}_c$ refers to the set of all subsets closed with respect to d_B. Their solution turns out to be isometric to a quotient space of $\langle T,d_w\rangle$. In this section we shall describe their solution $\langle P,d_B\rangle$ and establish the isometry.

<u>Definition</u>: Let $\langle P_n,d_n\rangle$ be a series of metric spaces defined by

$P_0 = \{p_0\}$, $\quad$ p_0 is the <u>nil</u> process,

$P_{n+1} = \{p_0\}\cup\mathcal{P}(\Sigma \times P_n)$ $\quad$ $\mathcal{P}$ is the power set operator,

and $d_0(p,q) = 0$ $\quad$ for all $p,q\in P_0$,

$$d_{n+1}(p,q) = \begin{cases} 0 & \text{for } p=q=p_0 \\ 1 & \text{for } p=p_0 \text{ or } q=p_0\text{, but not both} \\ \max(\sup\limits_{p'\in p}\inf\limits_{q'\in q} d'_{n+1}(p',q'),\ \sup\limits_{q'\in q}\inf\limits_{p'\in p} d'_{n+1}(p',q')) & \\ & \text{for both } p,q\subseteq\Sigma \times P_n \end{cases}$$

where

$$d'_{n+1}(p',q') = \begin{cases} 0 & \text{for } p'=q'=p_0 \\ 1 & \text{for } p'=p_0 \text{ or } q'=p_0\text{, but not both} \\ d_n(p'',q'')/2 & \text{for } p'=\langle a,p''\rangle,\ q'=\langle b,q''\rangle \text{ and } a=b \\ 1 & \text{above, except } a\neq b. \end{cases}$$

Note that d_{n+1} is the Hausdorff metric distance between the subsets of P_{n+1} induced by the metric d'_{n+1} on the points of P_{n+1}.

<u>Definition</u>: Let $\langle P,d_B\rangle$ be the completion of $\langle\cup P_n,\cup d_n\rangle$.

<u>Theorem 4.1</u> [BaZ]: $\langle P,d_B\rangle$ satisfies (4.1).

The quotient space through which the isometry will be established is the space of <u>reduced</u> trees. We need a preliminary definition:

<u>Definition</u>: For S bounded let $\Box S$ be:

$\Box(\text{nil})=\text{nil}$

$\Box(\oplus a_iS_i)=\oplus a_i'S_i'$ where $\{a_i'S_i'\}$ is the maximal collection of pairwise nonisomorphic prefixed subtrees of S.

We shall write for convenience $\Box a_iS_i$ for $\Box(\oplus a_iS_i)$.

<u>Examples</u>: (1) S =  , $\Box S$ = |a

(2) S = 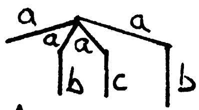 , $\Box S$ = 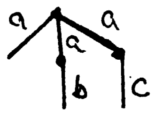

(3) S =  , $\Box S$ = S

<u>Lemma 4.2</u>: $\Box S$ is a tree.

<u>Proof</u>: We must verify that $\Box S$ is well defined. If $\{a_i'S_i'\}$ and $\{a_i''S_i''\}$ are two maximal collections of nonisomorphic prefixed trees of S, then the sets must necessarily be in 1-1 correspondence and so $\oplus a_i'S_i' = \oplus a_i''S_i''$.

<u>Definition</u>: (reduction operator) For S bounded let $R(S)$ be:

$$R(\text{nil}) = \text{nil}$$

$$R(\oplus a_iS_i) = \Box a_iR(S_i).$$

<u>Example</u>: For S = 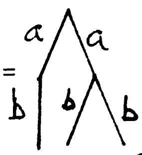 , R(S) =  (see example 3 above).

For convenience, let $R_n=\{R(S) \mid S \text{ bounded by } n\}$.

<u>Lemma 4.3</u>: (1) $R(S)$ is a tree;

(2) R_n is the set of all reduced trees of height S_n.

Proof: omitted.

We shall now establish an isometric bijection between $\langle UR_n, d_w\rangle$ and $\langle UP_n, Ud_n\rangle$.

Definition: Let $\phi: UR_n \to UP_n$ by $\phi(\mathcal{R}(nil)) = \phi(nil) = p_0$

$$\phi(\mathcal{R}(S)) = \phi(\Box a_i \mathcal{R}(S_i)) = \{\langle a_i, \phi(\mathcal{R}(S_i))\rangle\}$$

Theorem 4.3: $\phi|_{R_n}$ is a bijection between R_n and P_n.

Proof: Induction on n:

case n=0: immediate.

Assume for n.

case n+1:

ϕ is 1-1: Let $\mathcal{R}(S)$, $\mathcal{R}(T) \in R_{n+1}$ and suppose $\phi(\mathcal{R}(S)) = \phi(\mathcal{R}(T))$.

Now if $\phi(\mathcal{R}(S)) = \phi(\mathcal{R}(T)) = p_0$, then by the induction hypothesis, $\mathcal{R}(T) = \mathcal{R}(S) = nil$.

Suppose $\mathcal{R}(S) = \Box a_i \mathcal{R}(S_i)$ and $\mathcal{R}(T) = \Box b_j \mathcal{R}(T_j)$ where $\mathcal{R}(S_i)$, $\mathcal{R}(T_j) \in R_n$

$\therefore \{\langle a_i, \phi(\mathcal{R}(S_i))\rangle\} = \{\langle b_j, \phi(\mathcal{R}(T_j))\rangle\}$

$\therefore \forall_i \exists j \; \langle a_i, \phi(\mathcal{R}(S_i))\rangle = \langle b_j, \phi(\mathcal{R}(T_j))\rangle$ and vice versa

$\therefore a_i = b_j$ and $\phi(\mathcal{R}(S_i)) = \phi(\mathcal{R}(T_j))$

$\therefore \mathcal{R}(S_i) = \mathcal{R}(T_j)$ by the induction hypothesis

$\therefore a_i \mathcal{R}(S_i) = b_j \mathcal{R}(T_j)$

$\therefore \Box a_i \mathcal{R}(S_i) = \Box b_j \mathcal{R}(T_j)$ $\qquad \therefore \phi$ is 1-1

ϕ is onto: Let $p \in P_{n+1}$. If $p = p_0$, choose $\phi^{-1}(p) = nil$. Else $p = \{\langle a_i, p_i\rangle\}$, $p_i \in P_n$. By the induction hypothesis ϕ is onto P_n. Denote by $\phi^{-1}(p_i)$ the unique (ϕ is 1-1) element of R_n such that $\phi(\phi^{-1}(p_i)) = p_i$. Let $\phi^{-1}(p) = \Box a_i \phi^{-1}(p_i)$. Because $\phi^{-1}(p_i) \in R_n$ we have $a_i \phi^{-1}(p_i) \in R_{n+1}$ and therefore $\phi^{-1}(p) \in R_{n+1}$. Furthermore $\phi(\phi^{-1}(p)) = \{\langle a_i, p_i\rangle\} = p$

$\therefore \phi$ is onto.

Corollary: ϕ is a bijection between UR_n and UP_n.

The following will be useful in establishing the isometry:

Lemma 4.4: For $S,T \in UR_n$, $S \equiv T \iff S=T$, i.e. $\langle UR_n, d_w \rangle$ is a metric (not pseudometric) space.

Proof: (<=) immediate.

(=>) standard induction argument on $\equiv_k$.

Theorem 4.5: ϕ is an isometry between UR_n and UP_n.

Proof: We shall establish that $\forall S,T \in R_n$, $d_w(S,T)=d_n(\phi(S),\phi(T))$ from which the conclusion follows.

Induction on n:

case n=0: $d_w(S,T)=0$ since S=T=nil

and $d_0(\phi(S),\phi(T))=d_0(p_0,p_0)=0$.

Assume for n.

case n+1: we shall establish

$\forall S,T \in R_{n+1}$, $d_w(S,T)=2^{-k} \iff d_{n+1}(\phi(S),\phi(T))=2^{-k}$

Induction on k:

case k=0: $d_w(S,T)=0 \iff S=T$ (Lemma 4.4)

$\iff \phi(S)=\phi(T) \iff d_{n+1}(\phi(S),\phi(T))=0$

Assume for k.

case k+1: We know $d_w(S,T)=2^{-(k+1)} \iff S\equiv_{k+1}T$ and $S\not\equiv_{k+2}T$. We claim that $S\equiv_{k+1}T \iff d_{n+1}(\phi(S),\phi(T))\leq 2^{-(k+1)}$. If the claim is established, then the induction and theorem follow as

$$d_w(S,T)=2^{-(k+1)} \iff 2^{-(k+2)}<d_{n+1}(\phi(S),\phi(T))\leq 2^{-(k+1)}$$
$$\iff d_{n+1}(\phi(S),\phi(T))=2^{-(k+1)}$$

The first inequality above arises from the claim and the fact that $S\not\equiv_{k+2}T$. It remains to establish the claim.

Claim: $S \equiv_{k+1} T \iff d_{n+1}(\phi(S),\phi(T)) \leq 2^{-(k+1)}$

Proof:

(=>) $S \equiv_{k+1} T \iff \forall a \forall S'\ S \xrightarrow{a} S' \Rightarrow \exists T'\ T \xrightarrow{a} T'$ and $S' \equiv_k T'$ and vice versa

$\therefore\ d_n(\phi(S'),\phi(T')) \leq 2^{-k}$ by the induction hypothesis for n and k

$\therefore\ d'_{n+1}(\phi(aS'),\phi(aT')) \leq 2^{-(k+1)}$

$\therefore\ \inf_j d'_{n+1}(\phi(aS'),\phi(a_jT_j)) \leq 2^{-(k+1)}$

Since $S \equiv_{k+1} T$, $\forall i\ \inf_j d'_{n+1}(\phi(a_iS_i),\ \phi(a_jT_j)) \leq 2^{-(k+1)}$

$\therefore\ \sup_i \inf_j d'_{n+1}(\phi(a_iS_i),\ \phi(a_jT_j)) \leq 2^{-(k+1)}$

A similar argument establishes $\sup_j \inf_i d'_{n+1}(\cdot,\cdot) \leq 2^{-(k+1)}$

$\therefore\ d_{n+1}(\phi(S),\phi(T)) \leq 2^{-(k+1)}$.

(<=) Suppose now $d_{n+1}\phi(S),\phi(T)) \leq 2^{-(k+1)}$

Then $\sup_i \inf_j d'_{n+1}(\phi(a_iS_i),\ \phi(b_jT_j)) \leq 2^{-(k+1)}$

$\therefore\ \forall i \exists j\ d'_{n+1}(\phi(a_iS_i),\phi(b_jT_j)) \leq 2^{-(k+1)}$

$\therefore\ \forall i \exists j\ \ a_i = b_j$ and $d_n(\phi(S_i),\phi(T_j)) \leq 2^{-k}$

By applying the induction hypothesis for the claim for each prefixed subtree aS_i there is a corresponding aT_j such that $S_i \equiv_k T_j$

$\therefore$ one half of the definition of k+1-equivalence is satisfied. We obtain the other half from $\sup_j \inf_i d'_{n+1}(\cdot,\cdot) \leq 2^{-(k+1)}$.

This completes the proof of the claim and the theorem.

Since $\cup R_n \subseteq T$, the completion $\langle P,d_B\rangle$ of $\langle \cup P_n, \cup d_n\rangle$ is isometric to a complete subspace of $\langle T,d_W\rangle$ (modulo $\equiv$), say $\langle R,d_W\rangle$. We need to demonstrate that $\cup R_n$ is dense in T, i.e. that $[R/\equiv]=[T/\equiv]$, so that $\langle P,d_B\rangle$ will be isometric to $\langle T,d_W\rangle/\equiv$. We need a preliminary lemma.

Lemma 4.6: For S bounded, $S \equiv R(S)$.

Proof: Induction on k:

case *k=0*: immediate.

Assume *for* k.

case *k+1*: Let $S = \bigoplus a_i S_i$, $R(S) = \square\, a_i R(S_i)$

Suppose $a_i S_i$ is a prefixed subtree of S. Then $a_i R(S_i)$ is a prefixed subtree of R(S) (or there is some $a_i R'$ branch of R(S) such that $R' = R(S_i)$). Then by the induction hypothesis, $R(S_i) \equiv_k S_i$ and we are done in one direction.

The reverse direction is similar.

Theorem 4.7: For any $S \varepsilon \mathcal{T}$, there is some $T \varepsilon R$ such that $S \equiv T$.

Proof: Recall that $S = \cup S^{(n)}$.

Since $S^{(n)}$ bounded, $S^{(n)} \equiv R(S^{(n)})$ by Lemma 4.6.

$\therefore$ $\langle R(S^{(n)}) \rangle$ is a CS in $\langle R, d_W \rangle$ and therefore has a limit $T \varepsilon R$. Clearly $T \equiv S$.

5. A Connection with Programming Logic.

In this section we treat the case when our RST's are labeled from a finite set Σ. We introduce the small modal logic HML (Hennessy-Milner logic). It turns out that for any trees S,T, that $S \equiv T$ iff for every $\phi \varepsilon$HML, $S \models \phi \iff T \models \phi$. We exploit this fact to show that completeness of the space $\langle \mathcal{T}, d_w \rangle$ is a consequence of the Compactness Theorem for HML. This theorem in turn follows from the Compactness Theorem for first-order logic, so we have an alternative proof of completeness in this case. Finally, we observe that if our metric space is compact, then the HML Compactness Theorem follows as a consequence.

These results are in a sense already known in model theory. The relation $\equiv$ can be defined on arbitrary first order structures, and the equivalence $\mathcal{A} \equiv \mathcal{B}$ iff for all sentences ϕ, $\mathcal{A} \models \phi \iff \mathcal{B} \models \phi$ is part of the Ehrenfeucht-Fraissé characterization of elementary equivalence [Mon, p. 408]. HML can be considered as a fragment of first-order logic and the general theory applied. However, the proofs in the HML case are simple and revealing, so we think it worth while to present them here.

Definition: The set of formulas HML is given by the following inductive clauses:

tt, ff ε HML (two Boolean constants)

$\phi, \psi \; \varepsilon$ HML imply

$\phi \wedge \psi \; \varepsilon$ HML and $\neg\phi \; \varepsilon$ HML (Boolean operations)

$\phi \; \varepsilon$ HML and $a \varepsilon \Sigma$ imply

$a\langle\phi\rangle \varepsilon$HML ("possible" modality)

The formula $a\langle\phi\rangle$ is to be read: "From the initial state (root) it is possible to execute the atomic action a and arrive in a state satisfying ϕ". Note: Σ is henceforth finite.

Definition (semantics of HML): Let S be an RST over Σ, and let $\phi \varepsilon \text{HML}$. We say S satisfies ϕ ($S \models \phi$) in case we can apply the following inductive clauses:

$S \models tt$ always;

$S \models ff$ never;

$S \models \phi \wedge \psi$ iff $S \models \phi$ and $S \models \psi$;

$S \models \neg\phi$ iff not ($S \models \phi$);

$S \models a\langle\phi\rangle$ iff $(\exists S')(S \xrightarrow{a} S'$ and $S' \models \phi)$.

We proceed to develop some facts about HML and the relation $\equiv$.

Definition: The depth $|\phi|$ of an HML formula is given by:

$|tt| = |ff| = 0$;

$|\phi \wedge \psi| = \max(|\phi|, |\psi|)$;

$|\neg\phi| = |\phi|$;

$|a\langle\phi\rangle| = 1 + |\phi|$.

Let $\text{HML}_n = \{\phi \mid |\phi| \leq n\}$.

Lemma 5.1: For all T,U, and n, if $T \equiv_n U$ then for all $\phi \varepsilon \text{HML}_n (T \models \phi \Leftrightarrow U \models \phi)$.

Proof: easy induction on n.

The converse of 5.1 requires a little work, and is false unless Σ is finite.

Definition: Two HML formulas ϕ, ψ are logically equivalent iff for all T, $T \models \phi$ iff $T \models \psi$.

Lemma 5.2: For each n, the relation of logical equivalence restricted to HML_n has only finitely many equivalence classes.

Proof: Use induction on n; the proof amounts to finding a DNF for the formulas in HML_n. Here the finiteness of Σ must be used.

Lemma 5.3: For any n, and any T,U, if for all $\phi\varepsilon HML_n$, $T \models \phi \Leftrightarrow U \models \phi$, then $T \equiv_n U$.

Proof: Again, by induction on n. The result is clear when n=0. Assume it for k, and all T',U', and $\phi\varepsilon HML_k$. Suppose $T \xrightarrow{a} T'$. Let

$$F_k = \{\Theta_1, \ldots, \Theta_p\}$$

be a set of representatives of the equivalence classes of logical equivalence restricted to HML_k, and suppose $\Theta_1, \ldots \Theta_i$ are the formulas in F_k satisfied by T'. Then $T \models a\langle\Theta_1 \wedge \ldots \wedge \Theta_i \wedge \neg\Theta_{i+1} \wedge \ldots \wedge \neg\Theta_p\rangle$. This is a formula in HML_{k+1}, so by hypothesis, U satisfies it too. This gives a tree U' with $U \xrightarrow{a} U'$ and T' and U' satisfying exactly the same formulas in F_k. Since F_k is a complete set of representatives for logical equivalence, T' and U' satisfy exactly the same HML_k formulas. By inductive hypothesis, $T' \equiv_k U'$.

The case $U \xrightarrow{a} U'$ is of course exactly similar, so the proof of 5.3 is complete.

Corollary 5.4: $S \equiv T$ iff $\forall\phi\varepsilon HML$, $S \models \phi \Leftrightarrow T \models \phi$.

Corollary 5.5 ("Master formula" theorem for HML): For each $n \geq 0$ and each T, there is a formula $\phi(n,T)$ such that

(i) $T \models \phi(n,T)$;

(ii) For any U, if $U \models \phi(n,T)$ then $U \equiv_n T$.

Proof: As in 5.3 let F_n be a representative system for logical equivalence in HML_n. Given T, let

$$\phi(n,T) = \bigwedge\{\phi\varepsilon F_n \mid T \models \phi\}.$$
$$\wedge \bigwedge\{\neg\phi \mid \phi\varepsilon F_n \text{ and not } T \models \phi\}.$$

Clearly $T \models \phi(n,T)$. Further if $U \models \phi(n,T)$ then U and T agree on all formulas in F_n and thus on HML_n. The result follows from Lemma 5.3.

Theorem 5.6 (Compactness theorem for HML): Let $\Gamma \subseteq$ HML. If for any finite $\Delta \subseteq \Gamma$ there is a tree T such that $T \models \phi$ for all $\phi \varepsilon \Delta$, then there is a tree U such that for all $\phi \varepsilon \Gamma$, $U \models \phi$.

Proof: We translate (the semantics of) HML into first-order logic. For each $a \varepsilon \Sigma$ let $\underset{\sim}{a}$ be a binary relation symbol, and let $\underset{\sim}{k}$ be a constant symbol. Let L be the first-order language determined by these symbols.

For each $\phi \varepsilon$HML, we define a formula $\phi^* \varepsilon L$ with at most one free variable. Let tt* be some fixed tautological sentence in L, and let $ff^* = \neg(tt^*)$. Further define

$$(\phi \wedge \psi)^* = \phi^* \wedge \psi^*;$$

$$(\neg \phi)^* = \neg(\phi^*);$$

$(a\langle\phi\rangle)^* = \exists y(\underset{\sim}{a}(x,y) \wedge \phi^*(y))$, where y is the free variable in ϕ^* (if one exists) and x is a new free variable.

For any set Γ of formulas in HML, let

$$\Gamma^* = \{\phi^*(\underset{\sim}{k}) \mid \phi \varepsilon \Gamma\}.$$

The Γ^* is a set of sentences in L, and it is easy to show that Γ^* has a model if and only if Γ has a tree model. Now 5.6 follows immediately from the Compactness Theorem for first-order logic.

We can now prove that $\langle T, d_w \rangle$ is a complete metric space. Let $\langle T_k \rangle$ be a Cauchy sequence of trees. By passing to a subsequence if necessary, we may assume that for all k, $T_k \equiv_k T_{k+1}$. Now define

$$\Gamma = \{\phi(k, T_k) \mid k \geq 1\}$$

where the $\phi(k,T_k)$ are given by 5.5. We claim that for any U, if $U \models \phi(k,T_k)$ then for any $j \leq k$, $U \models \phi(j,T_j)$. The proof is by induction on k, and k=0 is trivial. Now if $U \models \phi(k+1,T_{k+1})$ then by 5.5, $U \equiv_{k+1} T_{k+1}$. Since $T_{k+1} \equiv_k T_k$, we have $U \equiv_k T_k$. But $|\phi(k,T_k)| \leq k$, so by Lemma 5.1 $U \models \phi(k,T_k)$. The claim follows by induction.

From the claim, if Δ is a finite subset of Γ, then Δ has a tree model. By 5.6, Γ has a tree model T; i.e. $T \models \phi(k,T_k)$ for all k. By 5.5 again, we have $T \equiv_k T_k$ for all k; i.e., $d_w(T,T_k) \to 0$ as desired.

Finally, we observe that from the compactness of $\langle \mathcal{T}, d_w \rangle$ we can derive the Compactness Theorem for HML. Let Γ be an arbitrary set of formulas such that every finite subset has a tree model. Enumerate $\Gamma = \{\phi_1, \phi_2, \ldots\}$. For each i let Δ_i be the set $\{\phi_1, \ldots, \phi_i\}$. Then each Δ_i has a tree model T_i. Since $\langle \mathcal{T}, d_w \rangle$ is compact, the sequence $\langle T_i \rangle$ has a convergent subsequence, say to some tree T. It is easy to see that T is a tree model for Γ. (The compactness of $\langle \mathcal{T}, d_w \rangle$ can be proved directly. One need only show completeness as in the previous sections, and the use the fact that Σ is finite to show that for any ε, a finite number of ε-spheres cover $\mathcal{T}$.)

References

[ArN] Arnold, A. and Nivat, M., "Metric Interpretations of Infinite Trees and Semantics of Nondeterministic Recursive Programs", Theoretical Computer Science 11, 1980, pp. 181-205.

[BaZ] de Bakker, J.W. and Zucker, J.I., "Denotational Semantics of Concurrency", Proceedings of the 14th STOC, May 1982, pp. 153-8.

[Dug] Dugundji, J., Topology, Allen & Bacon, Boston, 1966.

[HeM] Hennessy, M., and Milner, R., "On Observing Nondeterminism and Concurrency", 7th ICALP, Automata, Languages and Programming, LNCS 85, Springer-Verlag, New York, 1980, 299-309.

[Mil] Milner, R. A Calculus of Communicating Systems, LNCS 92, Springer-Verlag, New York, 1980.

[Mon] Monk, J.D., Mathematical Logic, Springer-Verlag, New York 1976.

[Niv] Nivat, M., "Infinite Words, Infinite Trees, Infinite Computations", Foundations of Computer Science III, 2, (J.W. de Bakker and J. Van Leeuwen, eds.), Mathematical Centre Tract 109, 1979.