Introduction

Motivation

A power domain construction maps every domain $\mathbf{X}$ of some distinguished class of domains into a so-called power domain over $\mathbf{X}$ whose points represent sets of points of the ground domain. Power domain constructions were originally proposed to model the semantics of non-deterministic programming languages [Plo76, Smy78, HP79, Mai85]. Other motivations are the semantic representation of a set data type [Hec90c], or of relational data bases [BDW88, Gun89b].

In 1976, Plotkin [Plo76] proposed the first power domain construction. Because his construction goes beyond the category of bounded complete algebraic domains, Plotkin proposed the larger category of *SFP-domains* that is closed under his construction. A short time later, Smyth [Smy78] introduced a simpler construction, the upper or Smyth power construction, that respects bounded completeness. In [Smy83], a third power domain construction occurs, the lower power domain, that completes the trio of classical power domain constructions.

Starting from problems in data base theory, Buneman et al. [BDW88] proposed to combine lower and upper power domain to a so-called sandwich power domain. Gunter investigated the logic of the classical power domains [Gun89a]. By extending the logic of Plotkin's domain in a natural way, he developed a so-called mixed power domain [Gun89b, Gun90]. Plotkin's power domain is a subset of the mixed one, and this in turn is a subset of the sandwich power domain.

We independently found the sandwich and mixed power domains in an isomorphic form as big and small set domains when developing domain constructions that would give semantics to an abstract data type of sets in a functional programming language (see [Hec90c]).

Given at least five different power domain constructions, the question arises what is the essence of these constructions, i.e. what are their common features which allow the application of the notion 'power domain'. Thus, we look for a theory of power domain constructions that covers the existing ones and provides answers to the following questions:

(1) What are power domain constructions?

(2) How are different power domain constructions related to each other?

(3) Are there more than the five constructions enumerated above?

(4) If so, how are these five constructions distinguished among all the others?

In addition, a general theory of power constructions provides — if it is to be useful — general theorems that are applicable to all specific power domain constructions.

The authors of the early papers about power domain constructions all mentioned some algebraic operations that are possible in their power domains. They however did not yet indicate the algebraic laws of these operations and the relations among different operations. In [HP79] then, the three classical power domains were characterized as free algebras w.r.t. some algebraic theories. Basically, these theories describe commutative idempotent semigroups, but employ various additional axioms that are mathematically not well motivated. In [Gun89b], Gunter characterized his mixed power domains in roughly the same line. Instead of employing additional axioms for the semigroup, he uses an additional unary operation. In [Hec90c], the sandwich power domains were also characterized as free algebras w.r.t. a theory of semigroups with a partial binary operation whose computational meaning is unclear (see chapter 24).

We felt that these results are not the general theory of power domain constructions. Their answer to question (1) above would be: power domains are free algebras according to algebraic theories of CI semigroups with more or less strange additional operations and axioms. Moreover, this theory has the drawback to deal more with single power domains than the whole power domain construction, i.e. the mapping from ground domains into power domains.

The algebraic theory of power domain constructions

Gunter presents in [Gun90] the semantics of a non-deterministic language in terms of a generic power domain construction using the three basic operators of singleton, binary union, and extending set-valued functions from points to sets. These generic semantics may then be instanciated by choosing a concrete construction instead of the generic one. The concrete construction only has to provide the necessary basic operations.

Thus, we define a power domain construction by axioms concerning the existence of some basic operations. In addition, we specify some axioms that should be satisfied by the basic operations. One might worry about the actual choice of these axioms, but we think that our choice is quite natural. This opinion is strengthened by the fact that our definition leads to a rich theory presented in Part II of the thesis, covers the known power constructions, and allows to characterize them algebraically.

Specification of power constructions

In the sequel, domains are just directed complete partially ordered sets, i.e. we neither require a least element nor algebraicity. As explained in chapter 9, a *power (domain) construction* $\mathcal{P}$ maps *ground domains* $\mathbf{X}$ into *power domains* over $\mathbf{X}$. The power domains have to satisfy the following axioms:

Empty set: There is a distinguished element θ in every power domain $\mathcal{P}\mathbf{X}$.

Binary union: There is a continuous operation $\uplus : [\mathcal{P}\mathbf{X} \times \mathcal{P}\mathbf{X} \to \mathcal{P}\mathbf{X}]$ in every power domain. '$\uplus$' is commutative and associative, and θ is its neutral element.

Singleton sets: There is a continuous mapping $\iota : [\mathbf{X} \to \mathcal{P}\mathbf{X}]$, $x \mapsto \{\!|x|\!\}$ for every ground domain $\mathbf{X}$.

Extension of functions: For every two domains $\mathbf{X}$ and $\mathbf{Y}$, there is a higher order function $ext : [[\mathbf{X} \to \mathcal{P}\mathbf{Y}] \to [\mathcal{P}\mathbf{X} \to \mathcal{P}\mathbf{Y}]]$ mapping set-valued functions on ground domain elements into set-valued functions on sets. It has to satisfy the following axioms:

(P1) $ext\, f\, \theta = \theta$

(P2) $ext\, f\, (A \uplus B) = ext\, f\, A \uplus ext\, f\, B$

(P3) $ext\, f\, \{\!|x|\!\} = f x$ or: $ext\, f \circ \iota = f$
(see the figure to the right)

(S1) $ext\, (\lambda x.\, \theta)\, A = \theta$

$$\begin{tikzcd}
\mathcal{P}\mathbf{X} \arrow[r, "ext\, f"] & \mathcal{P}\mathbf{Y} \\
\mathbf{X} \arrow[u, "\iota"] \arrow[ur, "f"'] &
\end{tikzcd}$$

(S2) $ext\, (\lambda x.\, f x \uplus g x)\, A = ext\, f\, A \uplus ext\, g\, A$.
Raising '$\uplus$' to functions, one may shortly write $ext\, (f \uplus g) = ext\, f \uplus ext\, g$.

(S3) $ext\, (\lambda x.\, \{\!|x|\!\})\, A = A$ or: $ext\, \iota = id$.

(S4) For every two morphisms $f : [\mathbf{X} \to \mathcal{P}\mathbf{Y}]$ and $g : [\mathbf{Y} \to \mathcal{P}\mathbf{Z}]$,

$$ext\, g\, (ext\, f\, A) = ext\, (\lambda a.\, ext\, g\, (f a))\, A$$

holds for all A in $\mathcal{P}\mathbf{X}$, or: $ext\, g \circ ext\, f = ext\, (ext\, g \circ f)$
(see the figure below)

$$\begin{tikzcd}
\mathcal{P}\mathbf{X} \arrow[r, "ext\, f"] & \mathcal{P}\mathbf{Y} \arrow[r, "ext\, g"] & \mathcal{P}\mathbf{Z} \\
\mathbf{X} \arrow[u, "\iota"] \arrow[ur, "f"'] & \mathbf{Y} \arrow[u, "\iota"] \arrow[ur, "g"'] &
\end{tikzcd}$$

Note that we do not require $ext\, f$ to be the only morphism satisfying (P1) through (P3) for given f. If we had assumed this uniqueness, (S1) through (S4) would have been provable. That is why we call them secondary axioms in contrast to the primary axioms (Pi).

Semirings and modules

The operations as specified above allow to derive many other operations with useful algebraic properties (see chapter 10). We here include the most important ones only.

Extension depends on two ground domains, $\mathbf{X}$ and $\mathbf{Y}$. Particularly interesting instances of extension are obtained if one of $\mathbf{X}$ and $\mathbf{Y}$ is the one-point domain $\mathbf{1} = \{\diamond\}$. In case $\mathbf{X} = \mathbf{1}$, extension has functionality $ext : [[\mathbf{1} \to \mathcal{P}\mathbf{Y}] \to [\mathcal{P}\mathbf{1} \to \mathcal{P}\mathbf{Y}]]$. Dropping the obsolete argument in $\mathbf{1}$, uncurrying and twisting arguments leads to the 'external product' $\cdot : [\mathcal{P}\mathbf{1} \times \mathcal{P}\mathbf{Y} \to \mathcal{P}\mathbf{Y}]$. The definition is $b \cdot S = ext\, (\lambda\diamond.\, S)\, b$. If we additionally choose $\mathbf{Y} = \mathbf{1}$, then multiplication becomes an inner operation of $\mathcal{P}\mathbf{1}$.

If we rename the empty set θ of $\mathcal{P}\mathbf{1}$ into 0, the union '$\uplus$' into '+', and the singleton $\{\!|\diamond|\!\}$ into 1, then the power axioms imply the following algebraic laws for $\mathcal{P}\mathbf{1}$:

$$\begin{array}{lll} a+(b+c)=(a+b)+c & a+b=b+a & a+0=0+a=a \\ a\cdot(b\cdot c)=(a\cdot b)\cdot c & & a\cdot 1=1\cdot a=a \\ a\cdot(b_1+b_2)=(a\cdot b_1)+(a\cdot b_2) & (a_1+a_2)\cdot b=(a_1\cdot b)+(a_2\cdot b) & a\cdot 0=0\cdot a=0 \end{array}$$

Such algebraic structures are called *semirings*. Note that all algebras occurring in this thesis have domains as carriers and continuous operations.

Semirings are common generalizations of rings as well as of distributive lattices. Thus, they may also be interpreted logically: 0 as 'false', 1 as 'true', addition as disjunction, and multiplication as conjunction. The semiring $\mathcal{P}\mathbf{1}$ then reveals the inherent logic of the power domain construction $\mathcal{P}$. This view makes most sense if the semiring addition is idempotent — a fact that holds for the known power constructions.

If we also rename the empty set θ of $\mathcal{P}\mathbf{X}$ into 0 and its union '$\uplus$' into '+', then we obtain the following algebraic laws where members of $\mathcal{P}\mathbf{X}$ are capitalized in contrast to members of $\mathcal{P}\mathbf{1}$:

$$\begin{array}{lll} A+(B+C)=(A+B)+C & A+B=B+A & A+0=0+A=A \\ r\cdot(s\cdot A)=(r\cdot s)\cdot A & & 1\cdot A=A \\ r\cdot(A_1+A_2)=(r\cdot A_1)+(r\cdot A_2) & (r_1+r_2)\cdot A=(r_1\cdot A)+(r_2\cdot A) & r\cdot 0=0\cdot A=0 \end{array}$$

Such algebraic structures are called $\mathcal{P}\mathbf{1}$-*modules*. A function $f : [M \to M']$ between two $\mathcal{P}\mathbf{1}$-modules is called *linear* iff $f(A+B) = fA + fB$ and $f(r\cdot A) = r\cdot fA$ hold.

The results of chapter 10 then allow to state:

Theorem 11.1.3: Let $\mathcal{P}$ be a power construction and let

$$\begin{array}{ll} + \;=\; \uplus : & [\mathcal{P}\mathbf{X}\times\mathcal{P}\mathbf{X}\to\mathcal{P}\mathbf{X}] \\ 0 \;=\; \theta : & \mathcal{P}\mathbf{X} \\ \cdot \;=\; \lambda(a,S).\,\mathit{ext}\,(\lambda\diamond.\,S)\,a : & [\mathcal{P}\mathbf{1}\times\mathcal{P}\mathbf{X}\to\mathcal{P}\mathbf{X}] \\ 1 \;=\; \{\!|\diamond|\!\} : & \mathcal{P}\mathbf{1} \end{array}$$

Then $\mathcal{P}\mathbf{1}$ with these operations is a semiring, and $\mathcal{P}\mathbf{X}$ is a $\mathcal{P}\mathbf{1}$-module for all domains $\mathbf{X}$. For $f : [\mathbf{X}\to\mathcal{P}\mathbf{Y}]$, the extension $\mathit{ext}\,f : [\mathcal{P}\mathbf{X}\to\mathcal{P}\mathbf{Y}]$ is linear, and $\mathit{ext}\,f\circ\iota = f$ holds.

This result connects our work with that of Main [Mai85] where power domains are introduced as free semiring modules. There are however some differences: our constructions may create non-free modules, and our singleton function ι need not be strict.

The semiring $\mathcal{P}\mathbf{1}$ is called the *characteristic semiring* of the power construction $\mathcal{P}$. Generalizing a bit, a power construction $\mathcal{P}$ is said to have characteristic semiring R iff $\mathcal{P}\mathbf{1}$ and R are isomorphic semirings where the isomorphism is assumed to be fixed.

Different power constructions may have the same characteristic semiring. Conversely, Th. 14.5.1 and Th. 15.1.1 provide two distinguished power constructions for any given semiring.

Power homomorphisms

Homomorphisms between algebraic structures are mappings preserving all operations of these structures. Power constructions may be considered algebraic structures on a higher level. Thus, it is also possible and useful to define corresponding homomorphisms (see chapter 12).

A power homomorphism $H : \mathcal{P} \dot{\to} \mathcal{Q}$ between two power constructions $\mathcal{P}$ and $\mathcal{Q}$ is a 'family' of morphisms $H = (H_{\mathbf{X}})_{\mathbf{X}} : [\mathcal{P}\mathbf{X} \to \mathcal{Q}\mathbf{X}]$ commuting over all power operations, i.e.

- The empty set in $\mathcal{P}\mathbf{X}$ is mapped to the empty set in $\mathcal{Q}\mathbf{X}$: $H\,\theta = \theta$.
- The image of a union is the union of the images: $H(A \uplus B) = (HA) \uplus (HB)$.
- Singletons in $\mathcal{P}\mathbf{X}$ are mapped to singletons in $\mathcal{Q}\mathbf{X}$: $H\{\!|x|\!\}_{\mathcal{P}} = \{\!|x|\!\}_{\mathcal{Q}}$.
- Let $f : [\mathbf{X} \to \mathcal{P}\mathbf{Y}]$. Then $H \circ f : [\mathbf{X} \to \mathcal{Q}\mathbf{Y}]$, and $H(ext_{\mathcal{P}}\, f\, A) = ext_{\mathcal{Q}}\,(H \circ f)(HA)$ has to hold for all A in $\mathcal{P}\mathbf{X}$.

If the two constructions $\mathcal{P}$ and $\mathcal{Q}$ share the same characteristic semiring, then one can define: A power homomorphism is linear iff all the functions $H_{\mathbf{X}}$ are linear.

A power isomorphism between two constructions $\mathcal{P}$ and $\mathcal{Q}$ is a family of isomorphisms $H = H_{\mathbf{X}} : [\mathcal{P}\mathbf{X} \to \mathcal{Q}\mathbf{X}]$ such that both $(H_{\mathbf{X}})_{\mathbf{X}}$ and $(H_{\mathbf{X}}^{-1})_{\mathbf{X}}$ are power homomorphisms.

Initial power constructions

A power construction $\mathcal{P}$ is *initial* for semiring R iff for all power constructions $\mathcal{Q}$ with the same characteristic semiring R there is exactly one linear power homomorphism $\mathcal{P} \dot{\to} \mathcal{Q}$. A main result of the theory of power constructions is the existence of such an initial construction for every semiring R.

Initial power domains over $\mathbf{X}$ are obtained as free R-$\mathbf{X}$-modules. $(\mathbf{F}, \eta)$ is a free R-$\mathbf{X}$-module iff $\mathbf{F}$ is an R-module and η is a continuous mapping from $\mathbf{X}$ to $\mathbf{F}$ such that for every R-module M and every continuous $f : \mathbf{X} \to M$, there is exactly one continuous linear $\overline{f} : \mathbf{F} \to M$ with $\overline{f} \circ \eta = f$.

The idea to consider free modules dates back to [HP79]. Hoofman [Hoo87] showed the existence of the free module for semiring $\{0, 1\}$. Main [Mai85] then proposed power constructions defined as free modules for some fancy semirings. In contrast to our work, he requires the singleton mapping to be strict without telling exactly why. The singleton maps of mixed and sandwich power domain are not strict, whence the singleton maps in our general theory of power constructions cannot be strict likewise.

For all semirings R and domains $\mathbf{X}$, free R-$\mathbf{X}$-modules exist and are unique up to isomorphism. Whereas the proof of uniqueness is quite simple, the proof of existence is difficult and usually done by categorical means. In this thesis, we adapted and simplified the proof of [Hoo87] such that categorical theorems are no longer used (section 13.5). Unfortunately, the proof does not provide an explicit representation of the free module.

The main theorem about initial power constructions is formulated in section 14.5.

Theorem 14.5.1:

Let R be a given semiring. Let $\mathcal{P}$ be the construction that maps every ground domain $\mathbf{X}$

into the free R-$\mathbf{X}$-module $(\mathbf{F}, \eta)$. Then $\mathcal{P}$ is the initial power construction for semiring R.

The power operations of union and empty set are given by addition and its neutral element in the free module $\mathbf{F}$, and singleton is η. For functions $f : [\mathbf{X} \to \mathcal{P}\mathbf{Y}]$, *ext* f is given by the unique extension of f. The a priori given external product of the modules $\mathcal{P}\mathbf{X}$ coincides with the external product derived from the power operations.

For any other power construction $\mathcal{Q}$ with semiring R, the unique linear power homomorphism $H : \mathcal{P} \dot{\to} \mathcal{Q}$ is obtained by $H_{\mathbf{X}} = \overline{\lambda x. \{\!|x|\!\}_{\mathcal{Q}}}$.

Final power constructions

A power construction $\mathcal{P}$ is *final* for semiring R iff for every power construction $\mathcal{Q}$ with the same semiring R there is exactly one linear power homomorphism $\mathcal{Q} \dot{\to} \mathcal{P}$. Final constructions are dealt with in chapter 15.

Final power constructions were never proposed in the literature, probably because the notion of a power homomorphism was missing. In contrast to initial constructions, the final power constructions may be explicitly presented in terms of second order predicates. Thus, it turns out that the upper power construction of [Smy83] is final, whereas the original proposal in [Smy78] is initial.

As mentioned above, extension yields particularly interesting operations if one of the two domains it depends on is the one-point domain $\mathbf{1} = \{\diamond\}$. In case $\mathbf{Y} = \mathbf{1}$, one obtains $ext : [[\mathbf{X} \to \mathcal{P}\mathbf{1}] \to [\mathcal{P}\mathbf{X} \to \mathcal{P}\mathbf{1}]]$ which can be re-arranged by uncurrying, twisting, and then currying again. The outcome is a morphism $\mathcal{E} : [\mathcal{P}\mathbf{X} \to [[\mathbf{X} \to \mathcal{P}\mathbf{1}] \to \mathcal{P}\mathbf{1}]]$ mapping formal sets into second order functions. In view of the logical interpretation of $\mathcal{P}\mathbf{1}$, these functions may be coined as second order predicates. The intuitive meaning of $\mathcal{E}$ is then existential quantification. Given a formal set A in $\mathcal{P}\mathbf{X}$ and a predicate $p : [\mathbf{X} \to \mathcal{P}\mathbf{1}]$, the value of $\mathcal{E} A\, p$ intuitively tells whether some member of A satisfies p.

Some of the axioms of extension easily translate into the following properties:

(P1) $\mathcal{E}\ \theta = \lambda p.\, \theta$

(P2) $\mathcal{E}\ (A \uplus B) = \lambda p. (\mathcal{E}\ A\ p) \uplus (\mathcal{E}\ B\ p)$

(P3) $\mathcal{E}\ \{\!|x|\!\} = \lambda p.\, p\, x$

(S4) $\mathcal{E}\ (ext\, f\, A) = \lambda p.\, \mathcal{E} A\, (\lambda a.\, \mathcal{E}\ (f a)\, p)$

These results suggest to define a power construction for given semiring $\mathcal{P}\mathbf{1}$ by $\mathcal{P}\mathbf{X} = [[\mathbf{X} \to \mathcal{P}\mathbf{1}] \to \mathcal{P}\mathbf{1}]$. Indeed, a slight variant of this method to obtain power constructions really works out.

Let R be a given semiring. Then we may define a product $\cdot : [\mathbf{X} \to R] \times R \to [\mathbf{X} \to R]$ by $f \cdot r = \lambda x. f x \cdot r$. Similarly, an addition may be defined as an inner operation on $[\mathbf{X} \to R]$. A function $F : [[\mathbf{X} \to R] \to R]$ is *right linear* iff $F\,(f_1 + f_2) = F f_1 + F f_2$ and $F\,(f \cdot r) = F f \cdot r$ hold. The domain of all right linear function from $[\mathbf{X} \to R]$ to R is denoted by $[[\mathbf{X} \to R] \overset{rlin}{\to} R]$.

Theorem 15.1.1: Let R be a given semiring. The final power construction for R is explicitly given by mapping the ground domain $\mathbf{X}$ to the space of right linear second order predicates over $\mathbf{X}$:

$\mathcal{P}\mathbf{X} = [[\mathbf{X} \to R] \stackrel{rlin}{\to} R]$. Its operations are defined by

- $\theta = \lambda g.\, 0$
- $A \uplus B = \lambda g.\, Ag + Bg$
- $\{\!|x|\!\} = \lambda g.\, gx$ for $x \in \mathbf{X}$.
- $ext\, f\, A = \lambda g.\, A\,(\lambda a.\, fag)$ for $f : [\mathbf{X} \to \mathcal{P}\mathbf{Y}]$ and $A \in \mathcal{P}\mathbf{X}$.

Let $\mathcal{Q}$ be another power construction with semiring R. The unique linear power homomorphism from $\mathcal{Q}$ to $\mathcal{P}$ is given by existential quantification $\mathcal{E}A = \lambda g.\, ext_{\mathcal{Q}}\, g\, A$ for A in $\mathcal{Q}\mathbf{X}$.

The definition of ext as well as those of the remaining operations are motivated by the properties of $\mathcal{E}$ listed above. Note that $\mathcal{P}\mathbf{1}$ is $[[\mathbf{1} \to R] \stackrel{rlin}{\to} R]$ which does not equal R but is isomorphic to R.

Sub-constructions

Let $\mathcal{P}$ be a given power construction. $\mathcal{Q}$ is called a *sub-construction* of $\mathcal{P}$ iff $\mathcal{Q}$ maps ground domains $\mathbf{X}$ into subsets of $\mathcal{P}\mathbf{X}$ such that

- $\mathcal{Q}\mathbf{X}$ is closed w.r.t. lubs of directed sets,
- $\theta \in \mathcal{Q}\mathbf{X}$,
- If A and B are in $\mathcal{Q}\mathbf{X}$, then $A \uplus B$ is in $\mathcal{Q}\mathbf{X}$,
- $\{\!|x|\!\}$ is in $\mathcal{Q}\mathbf{X}$ for all x in $\mathbf{X}$,
- If $f : [\mathbf{X} \to \mathcal{Q}\mathbf{Y}]$ and A in $\mathcal{Q}\mathbf{X}$, then $ext\, f\, A$ is in $\mathcal{Q}\mathbf{Y}$.

In shorter terms, $\mathcal{Q}\mathbf{X}$ is closed w.r.t. all power operations of $\mathcal{P}$. $\mathcal{Q}$ is obviously a power construction since the validity of the power axioms for $\mathcal{Q}$ is inherited from $\mathcal{P}$. Sub-constructions are handled in chapter 14.

One easily verifies that the intersection of a family of sub-constructions of a power construction $\mathcal{P}$ is again a sub-construction of $\mathcal{P}$, if we define $(\bigcap_{i \in I} \mathcal{Q}_i)\mathbf{X} = \bigcap_{i \in I}(\mathcal{Q}_i\mathbf{X})$. Hence, the sub-constructions of $\mathcal{P}$ form a complete lattice.

Let R be a semiring domain. R' is a *sub-semiring* of R iff R' is a subset of R containing 0 and 1, and being closed w.r.t. lubs of directed sets, addition, and multiplication. Because the operations in the characteristic semiring are derived from the power operations, the semiring of a sub-construction $\mathcal{Q}$ of $\mathcal{P}$ is a sub-semiring of the semiring of $\mathcal{P}$. Similarly, the semiring of an intersection of sub-constructions is the intersection of the semirings of the sub-constructions.

We found two general methods to derive sub-constructions from a given construction $\mathcal{P}$ with semiring R. The method of core formation yields the smallest sub-construction of $\mathcal{P}$ that still has semiring R, whereas the method of existential restriction creates the greatest sub-construction of $\mathcal{P}$ with a given sub-semiring R' of R. The core of the existential restriction is then the least sub-construction with the given semiring R'.

The *core* $\mathcal{P}^c\mathbf{X}$ of a power domain $\mathcal{P}\mathbf{X}$ is the least subset of $\mathcal{P}\mathbf{X}$ that contains the empty set θ and all singletons $\{\!|x|\!\}$, and is closed w.r.t. union '$\uplus$', product by factors in $\mathcal{P}\mathbf{1}$, and directed suprema. One can show that the modules $\mathcal{P}^c\mathbf{X}$ are also closed w.r.t. extension of $\mathcal{P}$ such that $\mathcal{P}^c$ is a sub-construction of $\mathcal{P}$. We call it the *core* of $\mathcal{P}$. It is the least sub-construction with the same semiring as $\mathcal{P}$.

A power construction that coincides with its core is called *reduced*. Reduced constructions enjoy particularly nice properties. All initial constructions are reduced.

If $\mathcal{P}$ is a power construction with semiring R, and R' is a sub-semiring of R, then the *existential restriction* of $\mathcal{P}$ to R' defined by

$$\mathcal{Q}\mathbf{X} = \mathcal{P}|_{R'}\mathbf{X} = \{A \in \mathcal{P}\mathbf{X} \mid \forall p : [\mathbf{X} \to R'] : \mathcal{E}A\,p \in R'\}$$

is a sub-construction of $\mathcal{P}$ with semiring R'. It is the greatest such sub-construction.

Because of its definition in terms of existential quantification, one might believe that the existential restriction of a final construction for R is a final construction for R'. However this is not true as we shall see in section 23.4. There are two reasons for this. First, two distinct second order predicates in $[[\mathbf{X} \to R] \stackrel{rlin}{\to} R]$ may produce equal results for predicates in $[\mathbf{X} \to R']$. They are then still different in the restriction of the final construction for R, but equal in $[[\mathbf{X} \to R'] \stackrel{rlin}{\to} R']$. Second, there may be additional members in $[[\mathbf{X} \to R'] \stackrel{rlin}{\to} R']$ that cannot be obtained by restricting predicates in $[[\mathbf{X} \to R] \stackrel{rlin}{\to} R]$.

Despite of this general result, we also meet examples for semirings R and R' where the existential restriction of the final construction for R is final for R' — see Th. 23.3.1.

Products of power constructions

In chapter 16, another method to build new constructions from existing ones is introduced. Given a family $(\mathcal{P}_i)_{i\in I}$ of power constructions, we may build a product construction $\mathcal{P} = \prod_{i\in I}\mathcal{P}_i$:

- $\mathcal{P}\mathbf{X} = \prod_{i\in I}\mathcal{P}_i\mathbf{X}$ for all ground domains $\mathbf{X}$
- $\theta = (\theta_i)_{i\in I}$
- $(A_i)_{i\in I} \uplus (B_i)_{i\in I} = (A_i \uplus B_i)_{i\in I}$
- $\{\!|x|\!\} = (\{\!|x|\!\}_i)_{i\in I}$ for all x in $\mathbf{X}$
- For $f : [\mathbf{X} \to \mathcal{P}\mathbf{Y}]$ let $f_i = \pi_i \circ f$. Then $ext\, f\,(A_i)_{i\in I} = (ext_i\, f_i\, A_i)_{i\in I}$ where ext_i denotes the extension functional of $\mathcal{P}_i$. Here, π_i denotes projection to component i.

The verification of the power axioms for $\mathcal{P}$ is straightforward since the power operations work independently in all dimensions. The characteristic semiring of $\mathcal{P}$ is the product of the characteristic semirings of the $\mathcal{P}_i$.

We derived the following properties of the product:

Finality theorem 16.3.1: If $\mathcal{P}_i$ are final power constructions for the semirings R_i for all $i \in I$, then the product $\prod_{i\in I}\mathcal{P}_i$ is a final construction for the product $\prod_{i\in I} R_i$.

$$\prod_{i\in I}[[\mathbf{X} \to R_i] \stackrel{rlin}{\to} R_i] \;\cong\; [[\mathbf{X} \to \prod_{i\in I} R_i] \stackrel{rlin}{\to} \prod_{i\in I} R_i]$$

Core theorem 16.5.1: The core of the product of two power constructions equals the product of their cores: $(\mathcal{P}_1 \times \mathcal{P}_2)^c = \mathcal{P}_1^c \times \mathcal{P}_2^c$

Initiality theorem 16.5.3:

If $\mathcal{P}_1$ is initial for R_1 and $\mathcal{P}_2$ initial for R_2, then $\mathcal{P}_1 \times \mathcal{P}_2$ is initial for $R_1 \times R_2$.

Factorization theorem 16.6.1:

Let $\mathcal{P}$ be a power domain construction for semiring $R_1 \times R_2$. $\mathcal{P}$ may be factorized, i.e. there are power constructions $\mathcal{P}_1$ for R_1 and $\mathcal{P}_2$ for R_2 such that $\mathcal{P}$ is isomorphic to $\mathcal{P}_1 \times \mathcal{P}_2$ by a linear power isomorphism, if and only if $ext\,(\lambda x.\,(1,\,0) \cdot \{\!|x|\!\})\, B = (1,\,0) \cdot B$ holds for all B in all power domains $\mathcal{P}\mathbf{X}$.

The condition of the factorization theorem holds for many classes of power constructions. For instance, every reduced power construction for $R_1 \times R_2$ can be factorized. The resulting factors are reduced again (cf. Cor. 16.6.4).

Special power constructions

The algebraic theory of power domain constructions may be applied to the known 5 power constructions mentioned in the beginning. In the following summary of our results, we restrict ourselves to algebraic ground domains.

The five known power constructions may be characterized as follows:

- The lower power construction $\mathcal{L}$ is both initial and final for the semiring $\mathbf{L} = \{0,\,1\}$ with $0 < 1$ and $1 + 1 = 1$.
- The upper or Smyth power construction $\mathcal{U}$ is both initial and final for the semiring $\mathbf{U} = \{0,\,1\}$ with $1 < 0$ and $1 + 1 = 1$.
- The convex or Plotkin power construction $\mathcal{C}$ is initial for the semiring $\mathbf{C} = \{0,\,1\}$ with $1 + 1 = 1$ and uncomparable values 0 and 1.
- The mixed power construction $\mathcal{M}$ of Gunter is initial for the semiring $\mathbf{B} = \{\perp,\,0,\,1\}$ with $\perp$ being below the uncomparable values 0 and 1. In $\mathbf{B}$, addition is given by parallel disjunction and multiplication by parallel conjunction.
- The sandwich power construction $\mathcal{S}$ of Buneman et al. is final for the same semiring $\mathbf{B}$.
- The final power construction for the semiring $\mathbf{C}$ is missing in this list. The reason is that it is degenerated and has quite awkward properties.

In chapter 17, we study these semirings and modules. The algebraic structures mentioned in [HP79] to characterize $\mathcal{C}$, $\mathcal{L}$, and $\mathcal{U}$ are nothing else but $\mathbf{C}$-, $\mathbf{L}$-, and $\mathbf{U}$-modules respectively without 0. Similarly, Gunter's mix algebras of [Gun89b, Gun90] are nothing else but $\mathbf{B}$-modules. Gunter's additional unary operation is just multiplication by $\perp$.

The theory of these special power constructions can be made more symmetric by taking the additional semirings $\mathbf{D}$ and $\overline{\mathbf{B}}$ into consideration. The 'double' semiring $\mathbf{D}$ is defined as $\mathbf{L} \times \mathbf{U}$. $\mathbf{D}$ has 4 elements: a least element $\perp = (0,\,1)$, a greatest element $\top = (1,\,0)$, and the two incomparable values $0 = (0,\,0)$ and $1 = (1,\,1)$ in between. $\overline{\mathbf{B}} = \{0,\,1,\,\top\}$ is the order-dual of $\mathbf{B}$. Both $\mathbf{B}$ and $\overline{\mathbf{B}}$ are sub-semirings of $\mathbf{D}$. Their intersection is $\mathbf{C}$.

According to the theory of products of power constructions, the product $\mathcal{D}$ of $\mathcal{L}$ and $\mathcal{U}$ is both initial and final for semiring $\mathbf{D}$. Existential restriction to $\mathbf{B}$ leads to the sandwich construction $\mathcal{S}$ whose core in turn is $\mathcal{M}$. Thus, $\mathcal{S}$ is the greatest sub-construction of $\mathcal{D}$ with semiring $\mathbf{B}$, and $\mathcal{M}$ is the least one. By employing the dual sub-semiring $\overline{\mathbf{B}}$, we obtain the dual sandwich construction $\overline{\mathcal{S}}$ and the dual mix construction $\overline{\mathcal{M}}$. The intersection $\mathcal{S}\overline{\mathcal{S}}$ of $\mathcal{S}$ and $\overline{\mathcal{S}}$ pays some interest since it preserves bounded completeness and has the same semiring $\mathbf{C}$ as Plotkin's construction $\mathcal{C}$. The intersection of $\mathcal{M}$ and $\overline{\mathcal{M}}$ is $\mathcal{C}$ itself.

The sandwich condition that characterizes $\mathcal{S}$ as sub-construction of $\mathcal{D} = \mathcal{L} \times \mathcal{U}$ is given by existential quantification. In chapter 21, we translate this condition into various equivalent forms. We also found an explicit characterization of $\mathcal{M}$ as a sub-construction of $\mathcal{D}$ or $\mathcal{S}$. This mix condition is discussed in chapter 22. In these chapters, the dual conditions for $\overline{\mathcal{S}}$ and $\overline{\mathcal{M}}$ are also investigated.

Part III also contains some more results that are not directly connected to the general algebraic theory of power constructions. The lower and upper power construction are shown to commute generalizing the results of [FM90], i.e. $\mathcal{L}\mathcal{U}\mathbf{X}$ and $\mathcal{U}\mathcal{L}\mathbf{X}$ are isomorphic (sections 20.5 and 20.6). The sandwich power domains $\mathcal{S}\mathbf{X}$ may be characterized as free sandwich algebras over $\mathbf{X}$ (chapter 24).

Domain classes

The results of Part III are proved for as large classes of domains as possible. The coincidence among the initial and the final lower power construction and the one explicitly given in terms of Scott closed sets could be shown for all ground domains. The situation is however much more complex if upper power constructions are considered. The question of coincidence among the initial upper construction $\mathcal{U}_i$, the final one $\mathcal{U}_f$, its core $\mathcal{U}_c$, and the explicitly given construction in terms of compact upper sets $\mathcal{U}_K$ could be answered for certain (though quite large) classes of domains only.

To analyze the situation precisely, we use topological methods that are introduced in Part I of the thesis. With these methods, various classes of domains can be defined and investigated. Some of these classes are novel. For the upper power constructions, $\mathcal{U}_i\mathbf{X} \cong \mathcal{U}_f\mathbf{X}$ could be shown for continuous ground domains (chapter 19). $\mathcal{U}_f\mathbf{X}$ and $\mathcal{U}_c\mathbf{X}$ however definitively coincide for the larger class of multi-continuous domains. The maximal class of coincidence is unknown in both cases. $\mathcal{U}_f\mathbf{X}$ and $\mathcal{U}_K\mathbf{X}$ are isomorphic iff the ground domain $\mathbf{X}$ is sober. $\mathcal{U}_K\mathbf{X}$ does not make sense for all ground domains. It is defined for the superclass K-RD of the class of sober domains.

Novel topological notions such as strong compactness (section 4.7) allow to define even more upper power constructions (chapter 20). The resulting constructions are sub-constructions of $\mathcal{U}_f$ and $\mathcal{U}_K$ respectively. They coincide with these super-constructions for multi-continuous ground domains. The question of full coincidence remains open.

Final remarks

History and publications

The algebraic theory of power constructions evolved from an attempt to specify the semantic properties of a set data type in a functional language. The resulting algebraic specification of 'set domain constructions' was published in [Hec90c]. Two specific set domain constructions that satisfy the specification were found. The big set domains later turned out to be isomorphic to the sandwich power domains, and the small set domains turned out to be isomorphic to the mixed power domains.

The specification of set domain constructions involves an a priori given domain of logical values, namely $\mathbf{B} = \{\perp, 0, 1\}$. In investigating various derived power operations, we found — also inspired by [Mai85] — that a domain of logical values may be derived a posteriori from the set-theoretic power operations empty set, singleton, union, and extension. The set construction specification then turned out to fall into two parts: a specification of more general entities, namely power domain constructions, and the requirement that $\mathcal{P}\mathbf{1}$ be isomorphic to $\mathbf{B}$. Dropping the now artificial restriction $\mathcal{P}\mathbf{1} \cong \mathbf{B}$ led to a first version of the specification of power constructions.

A second starting point was the observation that sandwich power domains $\mathcal{S}\mathbf{X}$ are isomorphic to function spaces $[[\mathbf{X} \to \mathbf{B}] \stackrel{add}{\to} \mathbf{B}]$. A first order-theoretic proof of this isomorphism is contained in the technical report [Hec90a]. In the sequel, we found that much more general function spaces $[[\mathbf{X} \to R] \stackrel{rlin}{\to} R]$ always define power constructions as defined by the algebraic specification. These results together with the characterization of mixed and Plotkin power domains in terms of second order predicates were presented at the MFPS'90 workshop in Kingston, Ontario. The finality of the construction in terms of second order predicates was found in summer 1990. Core, existential restriction, and the product theory are even more recent results.

The fully evolved algebraic theory is contained in the technical report [Hec90b] and accepted for publication as [Hec91]. A paper with emphasis on second order predicates and the sub-constructions of $\mathcal{D}$ is submitted to the TCS journal, and a third paper about the isomorphism between $\mathcal{LU}\mathbf{X}$ and $\mathcal{UL}\mathbf{X}$ is submitted to the MFPS'91 conference.

Acknowledgement

At first, I would like to thank Prof. Reinhard Wilhelm who supported, advised and encouraged me in manifold ways. I am also most grateful to my colleagues Helmut Seidl, Fritz Müller, and Andreas Hense for many fruitful discussions. Fritz also provided many helpful hints to the literature, and Helmut read and controlled hundreds of pages of a draft version of the thesis under great pressure of time. Carl Gunter proposed to attend the MFPS'90 workshop where I found partners for interesting discussions. He and Achim Jung also made valuable proposals to improve notation. Finally, I would like to thank all the other people who are not explicitly mentioned but nevertheless contributed to create this thesis.