CHAPTER 1: INTRODUCTION

1. History of the Problem

The calculation of the stable homotopy groups of spheres is one of the most central and intractable problems in algebraic topology. In the 1950s Serre [57] used his spectral sequence to study this problem. In 1962, Toda [60] used his triple brackets and the EHP sequence to calculate the first 19 stems. These methods were later extended by Mimura, Mori, Oda and Toda [44], [45], [46], [50] to compute the first 30 stems. In the late 1950s the study of the classical Adams spectral sequence was begun [1]. Computations in this spectral sequence are still being pursued using the May spectral sequence and the lambda algebra. The best published results are May's thesis [39] and the computation of the first 45 stable stems by Barratt, Mahowald, Tangora [10], [37] as corrected by Bruner [16]. The use of the BP Adams spectral sequence on this problem was initiated by Novikov [49] and Zahler [62]. Its most spectacular success has been at odd primes [42]. A recent detailed survey of the status of this computation and the methods that have been used has been written by Ravenel [55].

An exotic method for computing stable stems was developed in 1970 by Joel Cohen [19]. Recall [20] that for a generalized homology theory E_* and a spectrum X there is an Atiyah-Hirzebruch spectral sequence:

(1.1.1) $$E^2_{N,p} = H_N(X;E_p) \Longrightarrow E_{N+p}X.$$

Joel Cohen studied this spectral sequence with X an Eilenberg-MacLane spectrum and E equal to stable homotopy or mod p stable homotopy. His idea was to take advantage of the fact that in these cases the spectral sequence is converging to zero in positive degrees. Since the homology of the Eilenberg-MacLane spectra are known, one can inductively deduce the stable stems. This is analogous to the usual inductive computation of the cohomology of Eilenberg-MacLane spaces by the Serre spectral sequence [17]. In that example, however, all the work can be incorporated into the Kudo transgression theorem. Joel Cohen was able to compute a few low stems, but the computation became too complicated to continue. His method was discarded since the Adams spectral sequence computations seemed much more efficient. In 1972, however, Nigel Ray [56] used this spectral sequence with X = MSU and E = MSp. He took advantage of the fact that H_*MSU and MSp_*MSU are known to compute the first 19 homotopy groups of MSp. Again this method was discarded since David Segal had computed the first 31 homotopy groups of MSp by the Adams spectral sequence and his computations were extended to 100 stems in [31].

My interest in Atiyah-Hirzebruch spectral sequences began in 1978. In a joint paper with Snaith [32] we studied the case where X is BSp and E_* is stable homotopy. The methods we developed there, in particular the use of Landweber-Novikov operations to study differentials, were clearly applicable to a wide class of examples. In 1983, I observed that if Joel Cohen's method were applied to the case where X is BP and E_* is stable homotopy then the computations would be greatly simplified over Cohen's case because of the sparseness of H_*BP and because Quillen operations could be used to compute the differentials. So, I began computing at the prime two. I soon discovered that the computations became too complicated to do by hand, but since they were mostly algorithmic they could be done by a computer. Using an IBM PC/AT micro-computer I was able to compute the first 64 stable stems. This work is the account of that computation.

Kaoru Morisugi informed me that in 1972 he attempted to use this method to compute π_*^S at the prime three, but he became bogged down with technical problems.

2. The Brown-Peterson Spectrum and Quillen Operations

In this section we list some of the basic facts about the Brown-Peterson spectrum BP. The notation introduced here will be used throughout the computation.

Let MU denote the unitary Thom spectrum. By the Pontryagin-Thom isomorphism, π_*MU is isomorphic to Ω_*^U, the ring of bordism classes of compact smooth manifolds without boundary which have a complex structure on their stable normal bundles. Using the Adams spectral sequence, Milnor [43] computed π_*MU to be a polynomial algebra over Z with one generator in each even degree. Brown and Peterson [15] discovered that when the spectrum MU is localized at a prime p, it decomposes into a wedge of various suspensions of a spectrum BP. This spectrum defines a generalized homology theory BP_* and a generalized cohomology theory BP^*. We list several basic properties of BP at the prime two. The standard references are the expositions of Adams [7] and Wilson [61].

(1.2.1) There are $M_N \in H_*BP$ of degree $2(2^N-1)$ such that $M_0 = 1$ and

$$H_*BP = Z_{(2)}[M_1, \dots M_N, \dots].$$

(1.2.2) The Hurewicz homomorphism $h: \pi_*BP \longrightarrow H_*BP$ is a monomorphism. Henceforth we consider h as an inclusion.

(1.2.3) Define $V_N \in H_*BP$ of degree $2(2^N-1)$ recursively by $V_0 = 2$ and for $N \geq 1$:

$$V_N = 2M_N - \sum_{k=1}^{N-1} M_k \cdot V_{N-k}^{2^k}.$$

The $V_N/2$, $N \geq 1$, are polynomial generators for H_*BP. Moreover, all the V_N are in the image of h and $\pi_*BP = Z_{(2)}[V_1, \dots, V_N, \dots]$. The V_N are called the Hazewinkel generators [22], [23].

(1.2.4) BP^*BP is the algebra of BP-operations. These operations act on BP_*X for any spectrum X including $BP_*S = \pi_*BP$ and $BP_*KZ = H_*BP$. These operations are natural. In particular, they commute with the Hurewicz homomorphism h.

(1.2.5) $BP^*BP = \pi_*BP[[\, r_\omega \mid \omega$ is a finite sequence of nonnegative integers$]]$.

The r_ω are called the Quillen operations [54]. They have the following properties.

(a) The r_ω are $Z_{(2)}$-module homomorphisms.

(b) If $f: X \longrightarrow Y$ is a map of spectra then $f_* \circ r_\omega = r_\omega \circ f_*$. In particular, $h \circ r_\omega = r_\omega \circ h$.

(c) If X is a ring spectrum and $A, B \in BP_*X$ then we have the Cartan formula

$$r_\omega(A \cdot B) = \sum_{\omega=\omega'+\omega''} r_{\omega'}(A) \cdot r_{\omega''}(B).$$

In [32] we showed how Landweber-Novikov operations act on the Atiyah-Hirzebruch spectral sequences for $\pi_*^S BU$ and $\pi_*^S BSp$. The following theorem shows that the Quillen operations act on Atiyah-Hirzebruch spectral sequences for BP_*X.

THEOREM 1.2.6 Let F be a ring spectrum. Consider the Atiyah-Hirzebruch spectral sequence for F_*BP:

$$E^2_{N,t} = H_N BP \otimes F_t \Longrightarrow F_{N+t}$$

Then each Quillen operation r_ω of degree K induces a map of spectral sequences:

$$r_\omega : E^s_{N,t} \longrightarrow E^s_{N-K,t}.$$

These r_ω have the following properties:

(a) The r_ω are $Z_{(2)}$-module homomorphisms.

(b) The r_ω are natural with respect to maps of spectral sequences induced by maps of spectra.

(c) The r_ω satisfy the Cartan formula

$r_\omega(A \cdot B) = \sum_{\omega=\omega'+\omega''} r_{\omega'}(A) \cdot r_{\omega''}(B)$ for all $A, B \in E^s$.

(d) The action of r_ω on E^2 is given by $r_\omega \otimes 1$ where the latter r_ω is the usual Quillen operation on H_*BP.

(e) $d^s \circ r_\omega = r_\omega \circ d^s$ for all $s \geq 1$.

(f) The action of r_ω on $E^{s+1} = H_*(E^s, d^s)$ is induced by the action of r_ω on E^s.

(g) The action of r_ω on the E^s induce an action of r_ω on $E^\infty = \varinjlim E^s$.

(h) The action of r_ω on E^∞ defined by (g) agrees with the action of r_ω on E^∞ induced by the usual action of the Quillen operations on $F_*BP = BP_*F$.

PROOF. Since $r_\omega \in BP^kBP$, we can represent r_ω by a map of spectra $r_\omega : \Sigma^K BP \longrightarrow BP$. Since the Atiyah-Hirzebruch spectral sequence is natural we have an induced map of spectral sequences. All of the properties are immediate except for the Cartan formula (c). It follows from the observation that the following diagram must commute up to homotopy:

$$\begin{array}{ccccccc}
 & & \Sigma^k BP & & \xrightarrow{\quad r_\omega \quad} & & BP \\
 & \overset{\phi}{\nearrow} & & & & & \uparrow \phi \\
\Sigma^k BP\wedge BP & & & & & & \\
 & \underset{\psi}{\searrow} & & & & & \\
\bigvee_{\omega=\omega'+\omega''} \Sigma^k BP\wedge BP & \xrightarrow{T} & \bigvee_{\omega=\omega'+\omega''} \Sigma^{k'} BP \wedge \Sigma^{k''} BP & & \xrightarrow{r_{\omega'} \wedge r_{\omega''}} & & BP\wedge BP
\end{array}$$

In this diagram ϕ is product map of BP and ψ is the pinching map. In each wedge summand $k = k'+k''$ and T is the switching map. ∎

3. The Inductive Procedure

In this section we will describe in detail the inductive procedure that we will use to compute the stable stems. However, before we apply this procedure in Chapters 5 to 7 we will digress to compute the first eight rows of the spectral sequence in Chapter 3 and to study two of the basic ingredients of our procedure: Toda brackets in Chapter 2 and the image of J in Chapter 4. This section concludes with an exposition of the notation that we will use to denote the elements of π_*^S.

Consider the Atiyah-Hirzebruch spectral sequence:

(1.3.1) $$E^2_{N,t} = H_N BP \otimes \pi_t^S \Longrightarrow \pi_{N+t} BP.$$

Since H_*BP is zero in odd degrees we see that in this spectral sequence:

$$
(1.3.2) \qquad
\begin{aligned}
E^r_{N,*} &= 0 \text{ if } N \text{ is odd,} \\
d^{2r+1} &= 0 \text{ and} \\
E^{2r+1} &= E^{2r+2} \text{ for all } r.
\end{aligned}
$$

The Hurewicz homomorphism is given in terms of this spectral sequence by the following commutative square:

$$
(1.3.3) \qquad
\begin{array}{ccc}
\pi_N BP & \xrightarrow{\quad h \quad} & H_N BP \\
\downarrow & & \uparrow \cong \\
E^\infty_{N,0} & \rightarrowtail & E^2_{N,0}
\end{array}
$$

Since h is one-to-one, it follows that:

$$
(1.3.4) \qquad E^\infty_{N,t} = \begin{cases} 0 & \text{if } t \neq 0 \\ \pi_N BP & \text{if } t = 0 \end{cases} \quad \text{and}
$$

$$
(1.3.5) \qquad E^\infty_{*,0} = Z_{(2)}[v_1,\ldots,v_N,\ldots].
$$

Thus, there must be nonzero differentials originating on the 0 row so that each monomial $K(2^{-e}v_1^{e(1)}\cdots v_M^{e(M)})$ in E^2 survives to E^∞ if and only if K is divisible by 2^e where $e = e(1)+\cdots+e(M)$. We will prove in Chapter 4 that, in our range of computations, all nonzero differentials which originate on the 0 row land in $\mathrm{Im}J \otimes H_*BP$. We will assume that ImJ is known. The first step in our analysis of the spectral sequence (1.3.1) will be to compute all these differentials which originate on the 0 row in degrees 2 through 70. This computation is entirely algorithmic, is done by computer with no human assistance and is carried out in Section 4.4. The purpose of this computation is to record the cokernels of all of these differentials.

The behavior of the following elements in the spectral sequence is the key to the determination of differentials which originate above the 0 row.

DEFINITION 1.3.6 Let $\phi \in \pi_t^S$ have order q and let $V \in H_{2N}BP$. Assume that:

(a) $\phi \cdot V \in E^2_{2N,t}$ survives to an element of $E^{2r}_{2N,t}$ for some $2 \le r \le \infty$;

(b) if $r = \infty$ then $V = 0$;

(c) we know all differentials which originate or land on elements of $E^{2s}_{2k,t}$ which have a representative in $Z_q\phi \otimes H_*BP$ for all s and all $0 \le k < N'$ where $N' = N$ if $r < \infty$ or $N' = \infty$ if $r = \infty$.

We call such an element $\phi \cdot V$ a ϕ-leader.

Note: A ϕ-leader can be zero. In that case our assumption is that we know all differentials which originate or land in $Z_q\phi \otimes H_*BP$.

The following unfortunate phenomenon is the obstruction to using Theorem 1.2.6(e) to computing d^{2r}-differentials on $\phi \cdot V''$, degree V'' > degree V, from the d^{2r}-differential on a ϕ-leader $\phi \cdot V$.

DEFINITION 1.3.7 Let $\phi \cdot V$ be a ϕ-leader, and assume all the notation of Definition 1.3.6. A nonzero differential $d^{2u}(\phi \cdot V')$ is called a hidden differential if:

(a) $\phi \cdot V'$ is also a ϕ-leader;

(b) degree V' > degree V;

(c) $u < r$.

Thus, if there is a hidden differential, the d^{2u}-differentials determined by $d^{2u}(\phi \cdot V')$ must be computed before the d^{2r}-differentials determined by $d^{2r}(\phi \cdot V)$ even though degree $\phi \cdot V'$ > degree $\phi \cdot V$. The inductive computation of π_N^S now proceeds as follows. Assume that the information contained in the following induction hypothesis is known.

(1.3.8) INDUCTION HYPOTHESIS

(1_N) We know π_k^S for $0 \le k < N$.

(2_N) Write each nonzero differential on a ϕ-leader $\phi \cdot V \in E_{2a,b}^{2r}$, with $a+b \le N$, in the form $d^{2r}(\phi \cdot V) = \lambda V' \ne 0$ where $\phi \in \pi_b^S$, $\lambda \in \pi_{b+2r-1}^S$, $V \in H_{2a}BP$ and $V' \in H_{2a-2r}BP$. Assume that we have "computed" $d^{2r}(\phi \cdot V'') = \sum \alpha_I \lambda V_I$ for all $V'' \in H_{2a''}BP$.

(3_N) For each $\phi \in \pi_k^S$, $0 < k < N$, the ϕ-leader of largest known degree is $\phi \cdot V$ where either $V = 0$ or degree $\phi \cdot V \ge N+1$.

The information in (2_N) is called a "tentative differential table" and the information in (3_N) is called a "list of leaders". In condition (2_N), the word computed is in quotation marks because what we assume that we have done is that we have computed $r_\omega \circ d^{2r}(\phi \cdot V'') = d^{2r} \circ r_\omega(\phi \cdot V'')$ for all Quillen operations r_ω of degree $2a''-2a$. This would give an accurate computation of $d^{2r}(\phi \cdot V'')$ if there were no hidden differentials. Unfortunately, there are examples of hidden differentials.

To accomplish the inductive step we must go through the procedure below. We use the terminology "$A \in E_{2N,t}^{2r}$ transgresses" if A survives to E^{2N}. In that case $d^{2N}(A) \in E_{0,2N+t-1}^{2N}$, a subquotient of π_{2N+t-1}^S.

(1.3.9) INDUCTION STEP

(a) Construct the following list of leaders of degrees N+1 and N+2:

<u>Leaders in Degree N+1</u>	<u>Leaders in Degree N+2</u>
α_1	β_1
.	.
.	.
.	.
α_p	β_q

Each $\alpha_i \in E^{2a(i)}_{2a(i),N-2a(i)+1}$ will either be hit by some β_j or it will transgress to determine a nonzero element of π^S_N. In either case α_i transgresses to an element $d^{2a(i)}(\alpha_i) = \hat{\alpha}_i \in \pi^S_N$. In the former case $\hat{\alpha}_i = 0$, and in the latter case $\hat{\alpha}_i \neq 0$.

(b) Search for hidden differentials $d^{2u}(\beta) = \alpha_i$, where $d^{2r}(\beta) = \alpha'$ was one of the differentials in the tentative differential table of 1.3.8(2_N). If a hidden differential is found then α_i must be removed from the list in (a) and replaced with α'. Assume that any necessary adjustments of this sort have been made to the list in (a).

(c) Use Toda bracket methods from Chapter 2 and consequences of differentials which follow from Theorem 1.2.6(e) to make the following deductions:

(i) some of the $\hat{\alpha}_i$ are zero;

(ii) some of the β_j transgress.

This step is complete when

$$\text{card } \{\alpha_i \mid \hat{\alpha}_i = 0\} = \text{card } \{\beta_j \mid \beta_j \text{ is not known to transgress}\}.$$

(d) Construct the following list of all α_i, β_j such that $\hat{\alpha}_i = 0$ and β_j is not known to transgress:

$$\begin{array}{ll} \alpha_{i(1)} & \beta_{j(1)} \\ \cdot & \cdot \\ \cdot & \cdot \\ \cdot & \cdot \\ \alpha_{i(s)} & \beta_{j(s)} \end{array}$$

There is a nonzero differential on each $\beta_{j(k)}$ with image some $\alpha_{i(h)}$. Use Toda bracket methods from Chapter 2, consequences of differentials deduced from Theorem 1.2.6(e) and ad hoc monoid chain arguments to match which $\beta_{j(k)}$s hit which $\alpha_{i(h)}$s.

(e) Use Toda bracket methods from Chapter 2 to solve the additive extension problems to determine π^S_N from its composition series $\{E^{2r}_{0,N} | 1 \leq r \leq [(N+1)/2]\}$. This gives the information required in (1_{N+1}). This step is not absolutely essential and the computation can proceed even if all the additive extension problems can not be solved.

(f) Use the computer program of Section 9.3 to extend the tentative differential table for each of the nonzero differentials determined in (d). This gives the information required in (2_{N+1}).

(g) Update the list of leaders using the new information in the tentative differential table determined in (f). This gives the information required in (3_{N+1}).

In practice this inductive procedure is quite straightforward. There are usually no hidden differentials. Also there are usually very few matchings to be done in (d) and those matchings are obvious. In addition, there are never many possibilities for nontrivial additive extensions and many of these possibilities are quite easy to eliminate. As a final word of encouragement, the reader will see that the above procedure is merely the formalization of the straightforward common sense approach to the analysis of the spectral sequence. The following theorem is widely applicable. (See Appendix 2.)

THEOREM 1.3.10 Assume that $\xi \in \pi_N^S$ is defined as $\xi = d^r(X)$ where ξ is nonzero in $E^r_{N,0}$. If $r > N/2$ then ξ is indecomposable in the ring π_*^S.

PROOF. Assume that ξ is decomposable. Write $\xi = \Sigma\, \alpha_i\beta_i$, where $\alpha_i = d^{s(i)}(A_i)$, $\beta_i = d^{t(i)}(B_i)$ and $s(i) \le t(i)$ for all i. Then $s(i) < r$ for all i and $\xi = \Sigma\, d^s(\beta_i A_i)$ where s is the largest of all the s(i). Since $s < r$, $\xi = 0$ in E^r, a contradiction. Thus, ξ must be indecomposable. ∎

We conclude with the notation that we will use to describe elements of π_*^S. There are competing notations for the elements of the known stable stems. To add to the confusion, most methods of computing stable stems (including ours) only define elements of π_*^S modulo indeterminacy: the indeterminacy of a Toda bracket or of the filtration of a spectral sequence. We will use the usual notation for the elements of Hopf invariant one:

$$\eta \in \pi_1^S,\ \nu \in \pi_3^S \text{ and } \sigma \in \pi_7^S.$$

We will also use the following notation for elements in Im J: $\alpha_N \in \pi_{8N+1}^S$, $\beta_N \in \pi_{8N+3}^S$ and $\gamma_N \in \pi_{8N+7}^S$. If an element $X \in \pi_*^S$ is known to be decomposable then we will usually write it as a product. We will use the following notation for other elements of π_*^S.

DEFINITION 1.3.11 A[N,k] denotes the k^{th} element of π_N^S of order two, B[N,k] denotes the k^{th} element of π_N^S of order four, C[N,k] denotes the k^{th} element of π_N^S of order eight, etc. If there is only one element of π_N^S of a given order then we drop the second entry.

The following examples will help to explain this notation.

1. The element usually denoted $\varepsilon \in \pi_8^S$ of order two will be denoted A[8].
2. The element usually denoted $\bar{\kappa} \in \pi_{20}^S$ of order eight will be denoted C[20].
3. If we write D[45] we are denoting an element of π_{45}^S which has order 16.

We will also use the following notation. Let R be a PID and $\mathcal{A}$ a commutative R-algebra which is a free R-module. If $B, X_1, \ldots, X_t \in \mathcal{A}$ then $RB\{X_1, \ldots, X_t\}$ denotes the free R-submodule of $\mathcal{A}$ with basis $\{BX_1, \ldots, BX_t\}$. For example, let $\xi \in \pi_*^S$ have order N. We may take $R = Z_N$, $\mathcal{A} = Z_N\xi \otimes H_*BP$ and $X_1, \ldots, X_t$ linearly independent elements of H_*BP.

If $\alpha,\ \beta \in \pi_*^S$ and $\alpha\cdot\beta = 0$ then $B_{\alpha\beta}$ denotes a map H from a disc to a sphere such that H restricted to the boundary of its domain is $\alpha' \wedge \beta'$ where α', β' represents α, β respectively.