Gaston, Maple and Mike

Michael Monagan

Center for Experimental and Constructive Mathematics
Simon Fraser University
British Columbia

GNOME 2014, Zurich,
July 4th, 2014

Me, Gaston and Maple

May 1982 – Dec 1982 Waterloo, Masters student
Jan 1983 – Aug 1989 Waterloo, PhD student
Aug 1989 – Oct 1995 Zurich, Assistent

Me, Gaston and Maple

May 1982 – Dec 1982 Waterloo, Masters student
Jan 1983 – Aug 1989 Waterloo, PhD student
Aug 1989 – Oct 1995 Zurich, Assistent

Gaston gave me this paper for my Masters essay

> Shafi Goldswasser and Silvio Micali.
> *Probabilistic encryption & how to play mental poker keeping secret all partial information.* STOC '82, June 1982

which we implemented in Maple.

Gaston's number theory package, the first Maple package.

```
    |\^/|     Maple V Release 4 (WMI Campus Wide License)
._|\|   |/|_. Copyright (c) 1981-1996 by Waterloo Maple Inc. All rights
 \  MAPLE  /  reserved. Maple and Maple V are registered trademarks of
 <____ ____>  Waterloo Maple Inc.
      |       Type ? for help.

> with(numtheory);
Warning, new definition for order

[ F, M, cyclotomic, divisors, factorset, fermat, ifactor, imagunit,

    isprime, issqrfree, ithprime, jacobi, lambda, legendre, mcombine,

    mersenne, mlog, mroot, msqrt, nextprime, order, phi, prevprime,

    pprimroot, primroot, quadres, rootsunit, safeprime, sigma, tau]
```

I chose not to pursue cryptography for a PhD.

Life as a graduate student with Gaston ...

First Maple retreat, Sparrow lake, summer, 1983

What was Gaston's main contribution to Maple?

What was Gaston's main contribution to Maple?

Maple's Sum-of-Products representation and hashing of all subexpressions.

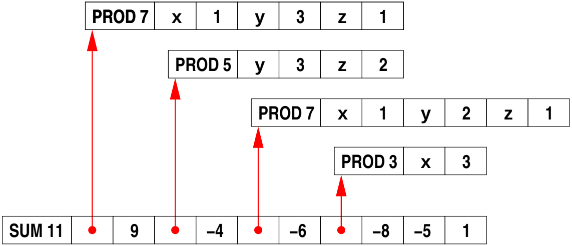

$$9\,xy^3z - 4\,y^3z^2 - 6\,xy^2z - 8\,x^3 - 5$$

What was Gaston's main contribution to Maple?

Maple's Sum-of-Products representation and hashing of all subexpressions.

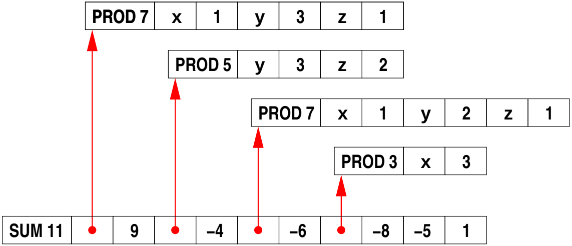

$$9\,xy^3z - 4\,y^3z^2 - 6\,xy^2z - 8\,x^3 - 5$$

What is the most important operation to make efficient?

What was Gaston's main contribution to Maple?

Maple's Sum-of-Products representation and hashing of all subexpressions.

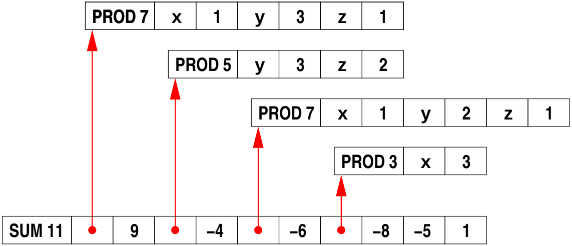

$$9\,xy^3z - 4\,y^3z^2 - 6\,xy^2z - 8\,x^3 - 5$$

What is the most important operation to make efficient?
Polynomial multiplication (and division).
But monomial multiplication cost $>$ 200 cycles.

Singular's representation

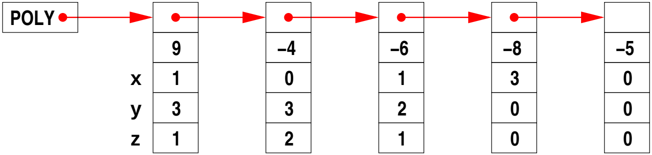

$$9\,xy^3z - 4\,y^3z^2 - 6\,xy^2z - 8\,x^3 - 5$$

Our new POLY representation (default in Maple 17)

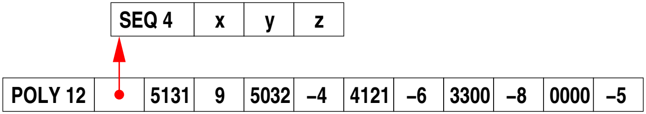

$$9\,xy^3z - 4\,y^3z^2 - 6\,xy^2z - 8\,x^3 - 5.$$

6 advantages

Our new POLY representation (default in Maple 17)

SEQ 4	x	y	z

POLY 12	•	5131	9	5032	–4	4121	–6	3300	–8	0000	–5

$$9\,xy^3z - 4\,y^3z^2 - 6\,xy^2z - 8\,x^3 - 5.$$

6 advantages

1. It's about $4\times$ more compact.
2. Memory access is sequential.
3. Kernel operations become $O(\#terms)$, some $O(1)$.
4. Monomial multiplication is one 64 bit integer $+$
   Monomial comparison is one 64 bit integer $>$
5. The simpl table is not filled with PRODs.
6. Division cannot cause exponent overflow in graded lex order.

What will fast multiplication using POLY do for the Maple library?

Intel Core i7 920 2.66 GHz (**4 cores**) Times in seconds

	Maple	Maple 16		Magma	Singular	Mathem
multiply	13	1 core	4 cores	2.16-8	3.1.0	atica 7
$p_4 := f_4(f_4+1)$	95.97	2.14	0.643	13.25	30.64	273.01
divide						
$q_4 := p_4/f_4$	192.87	2.25	0.767	18.54	14.96	228.83
factor	**Hensel lifting is mostly polynomial multiplication!**					
p_4 135751 terms	2953.54	59.29	46.41	332.86	404.86	655.49

$$f_4 = (1+x+y+z+t)^{20}+1 \qquad 10626 \text{ terms}$$

Parallel speedup for $f_4 \times (f_4+1)$ is 2.14 / .643 = **3.33**$\times$. **Why?**

What will fast multiplication using POLY do for the Maple library?

Intel Core i7 920 2.66 GHz (**4 cores**) Times in seconds

multiply	Maple 13	Maple 16 1 core	Maple 16 4 cores	Magma 2.16-8	Singular 3.1.0	Mathematica 7
$p_4 := f_4(f_4 + 1)$	95.97	2.14	0.643	13.25	30.64	273.01
divide						
$q_4 := p_4 / f_4$	192.87	2.25	0.767	18.54	14.96	228.83
factor	**Hensel lifting is mostly polynomial multiplication!**					
p_4 135751 terms	2953.54	59.29	46.41	332.86	404.86	655.49

$$f_4 = (1 + x + y + z + t)^{20} + 1 \qquad \text{10626 terms}$$

Parallel speedup for $f_4 \times (f_4 + 1)$ is 2.14 / .643 = **3.33**$\times$. **Why?**
Conversion overhead between POLY and SUM of PRODs!

After brainstorming with Roman, I asked Laurent if we could make POLY the default in Maple. Maple 17 uses POLY if all monomials in a polynomial with integer coefficients fit in 64 bits - otherwise we use SUM-of-PRODs. Conversions between POLY and SUM-of-PRODs are automatic and invisible to the Maple user.

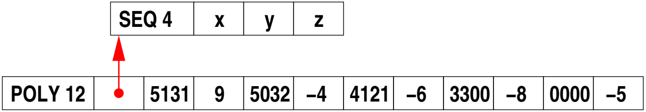

So we coded POLY for each kernel routine.

Faster at everything except op, map, etc.

command	Maple 16	Maple 17	speedup	notes
$\mathtt{coeff}(f, x, 20)$	2.140 s	0.005 s	420x	terms easy to locate
$\mathtt{coeffs}(f, x)$	0.979 s	0.119 s	8x	reorder exponents and radix
$\mathtt{degree}(f, x)$	0.073 s	0.003 s	24x	stop early using monomial de
$\mathtt{diff}(f, x)$	0.956 s	0.031 s	30x	terms remain sorted
$\mathtt{eval}(f, x = 6)$	3.760 s	0.175 s	21x	use Horner form recursively
$\mathtt{expand}(2 * x * f)$	1.190 s	0.066 s	18x	terms remain sorted
$\mathtt{indets}(f)$	0.060 s	0.000 s	$\rightarrow O(1)$	first word in dag
$\mathtt{op}(f)$	0.634 s	2.420 s	0.26x	has to construct old structur
`for t in f do`	0.646 s	2.460 s	0.26x	has to construct old structur
$\mathtt{taylor}(f, x, 50)$	0.668 s	0.055 s	12x	get coefficients in one pass
$\mathtt{type}(f, polynom)$	0.029 s	0.000 s	$\rightarrow O(n)$	type check variables only
$f;$	0.162 s	0.000 s	$\rightarrow O(n)$	evaluate the variables

For f with $n = 3$ variables and $t = 10^6$ terms created by

```
f := expand(mul(randpoly(v,degree=100,dense),v=[x,y,z])):
```

multiply	Maple 16 1 core	Maple 16 4 cores	Maple 17 1 core	Maple 17 4 cores
$p_4 := f_4(f_4+1)$	2.140	0.643	1.770	0.416
factor p_4 135751 terms	59.27	46.41	24.35	12.65

Intel Core i5 750 2.66 GHz **4 cores**. Real times in seconds.

$$f_4 = (1+x+y+z+t)^{20}+1 \qquad \text{10626 terms}$$

Parallel speedup for $f_4 \times (f_4+1)$ is $1.77/0.416 = \mathbf{4.2}\times$. **How ?**

Joris van der Hoven: Do you use the extra bits for the total degree?
My answer: No, because ...

I changed my mind. Roman Pearce recoded everything for Maple 18.

	per variable		total degree		$V_n = \det n \times n$ Vandermonde			
n	#bits	maxdeg	**#bits**	maxdeg	deg	Maple 16	17	18
7	8	255	8	255	21	0.012s	0.005	0.004
8	7	127	8	255	28	0.093s	0.027	0.026
9	6	63	10	1023	36	1.35 s	0.218	0.150
10	5	31	14	16383	45	15.95s	25.44	1.57
11	5	31	9	511	55	–	–	18.87
12	4	15	16	65535	66			236.4
13	4	15	12	4095	78			–
14	4	15	8	255	91			
15	4	15	4	15	105			
16	3	7	16	65535	120			

Maple retreat, Sparrow lake, circa 1992

Thank you Gaston for Waterloo, Zurich and Maple. Mike.

Notes on integration of POLY for Maple 17

Given a polynomial $f(x_1, x_2, ..., x_n)$, we store f using POLY if

(1) f is expanded and has integer coefficients,

(2) $d > 1$ and $t > 1$ where $d = \deg f$ and $t = \#\text{terms}$,

(3) we can pack all monomials of f into **one 64 bit word**, i.e. if $d < 2^b$ where $b = \lfloor \frac{64}{n+1} \rfloor$

Otherwise we use the sum-of-products representation.

- The representation is invisible to the Maple user. Conversions are automatic.
- POLY polynomials will be displayed in sorted order.
- Packing is fixed by $n = \#\text{variables}$.
- Maple 18 uses remaining bits for total degree.

Parallel multiplication using a binary heap.

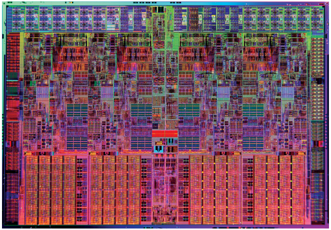

Target architecture

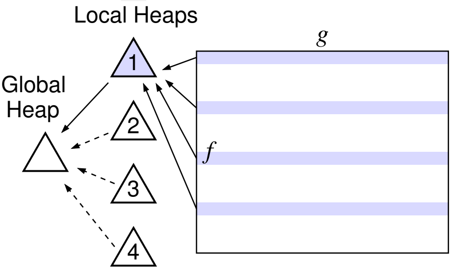

One thread per core.

Threads write to a finite circular buffer.

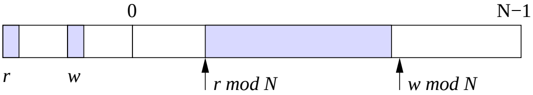

Threads try to acquire global heap as buffer fills up to balance load.