GROUPS OF ORDER pq AND p^2 DONE BY HAND

Jon A. Sjogren

Air Force Office of Scientific Research

Revised 15 March 2001

für Elisa

Introduction

We give the structure of the groups G of order pq and p^2, where p and q are primes, assuming that some basic results about groups have been covered, including the structure of finite abelian groups. We attempt to give a minimalist argument, nonetheless allowing the student to fully picture both the reasoning and the groups involved. Also, indications of other approaches are given, that would lead to the more standard Sylow Theorem arguments.

Groups of Order pq

Fact 1. *The order of a subgroup $H \subset G$ divides the order of G.*

This is Lagrange's Theorem. It is proved by considering the equivalence relation $g_1 \sim g_2$ when $g_1 = g_2 h$ for some $h \in H$. Any two equivalence classes (cosets) $g_1 H, g_2 H$ are in one-to=one correspondence under $g_1 h \leftrightarrow g_2 h$.

Corollary. *If H has prime order p, then any non-unity element $a \in H$ generates all of H.*

Let $\mathbb{A}$ be some designated finite set and X_p be the index set $\{1, \ldots, p\}$. Denote by α the standard cycle

$$(1, 2, \ldots, p),$$

the permutation that sends $1 \mapsto 2$ and so forth. Then a *word* $w = a_1 a_2 \ldots a_p$ really is the same as a function

$$f : X_p \to \mathbb{A}.$$

Thus $a_i \in \mathbb{A}$ for $1 \leq i \leq p$. Now the cycle α acts on w as follows.

$$\alpha w = a_{\alpha 1} \ldots a_{\alpha p} = a_2 \ldots a_p a_1.$$

Fact 2. *If $\alpha w = w$ as words (or as functions on X_p) then w is a constant word $aa\ldots a$ for some $a \in \mathbb{A}$.*

Proof. From the first index we see that $a_1 = a_2$. From the i-th index we have $a_i = a_{i+1}$, $i < p$ so that by induction we recover $a_{i+1} = a_1$, so now the entries are equal for all indices.

Now α generates a subgroup of Σ_p, the group of permutations of X_p, which is isomorphic to $\mathbb{Z}_p$, the integers modulo p.

Proposition 1. *Let $\sigma \in Z_p$, not the identity, acting as indicated by cycles on X_p. Then given a word w with values in $\mathbb{A}$, σw is well-defined. If $\sigma w = w$, then w is a constant word.*

Proof. By the Corollary to Fact 1 σ is a generator of $\mathbb{Z}_p$. Thus $\alpha = \sigma^k$ for some $1 \le k \le p-1$. Since $\sigma w = w$, we obtain $\alpha w = w$, so by Fact 2, $w = aa\ldots a$ for some $a \in \mathbb{A}$.

Proposition 2. *(After J.H. McKay, [1]) Let the order of a group G be divisible by the prime integer p. Then there exists a subgroup $H \subset G$ of order p. In fact, the number of such subgroups H is of the form $kp+1$ for some integer $k \ge 0$.*

Proof. Let G have order n and let $\mathbb{A}$ be the underlying set of G. Then let S be the set of p-words $w = a_1 \ldots a_p$ such that the product of each word in G equals e, the unity element. Since the first $p-1$ entries of w can be chosen arbitrarily, and the final entry is thereby determined, the cardinality of S is n^{p-1}. Define an equivalence relation on S by the action of the group $\mathbb{Z}_p$ of p-cycles above. That is, two words are considered equivalent if there is a cyclic permuation taking one to the other. If w is a constant word, its class consists of one element. Suppose σ and θ are two cycles, $\sigma \ne \theta$. Then if $\sigma w = \theta w$, the non-unity cycle $\sigma^{-1}\theta$ keeps w, invariant, hence by Proposition 1, w is a constant word. Thus if w is not a constant word, its equivalence class under the cyclic action consists of p elements.

Let r denote the number of solutions to the equation $x^p = e$ in G. Then r counts those elements of S whose equivalence class consists of only one element. We showed that all other elements of S are such that their equivalence class has p elements. If the number of those equivalence classes is s, then we have shown,

$$r + sp = n^{p-1}$$

and hence $p \mid r$ or $r = hp$. The fact that e is a solution to $x^p = e$ shows that h is an integer > 0. Suppose m is the number of distinct subgroups of G of order p. Two such groups H, K are distinct if there is $h \in H$, h not in K (possibly exchanging H and K). If there exists $y \in H \cap K$, $y \ne e$, then by the Corollary to Fact 1, for an integer $k > 0$ we must have $x = y^k \in K$ which is a contradiction. Thus distinct subgroups of G of order p intersect only in the identity. Each of the r solutions of $x^p = e$ must conversely lie in at least one of these m subgroups. (Only the identity element lies in more than one of the subgroups.) The total number of elements in these subgroups is therefore

$$m(p-1) + 1 = r = hp.$$

Hence since $h > 0$ there is an integer $k \ge 0$ such that $m = kp + 1$, so in particular there is at least one subgroup H of G of order p.

The existence of a subgroup of order p is Cauchy's Theorem, and the congruence satisfied by m is part of the Third Theorem of Sylow in the prime case [2].

Fact 3. *Let A and B be subgroups of G such that $A \cap B = \{e\}$. Then the set (subset of the underlying set of G) $S = \{ab \mid a \in A, b \in B\}$ consists of distinct elements for distinct a and b. That is, the cardinality of S is the product of the orders of A and B.*

Proof. For two apparently distinct pairs to be equal would require $ab = a'b'$ which in view of the hypothesis leads to $a = a'$ and $b = b'$.

Proposition 3. *If p and q are prime integers with $p < q$, and G has order pq, then G has a subgroup P of order p and a unique subgroup Q of order q. Therefore also Q is normal in G.*

Proof. By Proposition 2, subgroups P of order p and Q of order q exist. If there were two distinct subgroups Q and Q' of order q, they would intersect only in $\{e\}$. Otherwise they would have a common generator by the Corollary to Fact 1 which means that $Q = Q'$, a contradiction. Now by Fact 3, we immediately have q^2 distinct elements in G, which contradicts $|G| = pq$ and $p < q$. Given $g \in G$, it is easy to see that "conjugation" with g is a one-to-one mapping of a given subset $S \subset G$ to gSg^{-1}. In addition, gQg^{-1} is a subgroup of G, has order q, hence must be equal to Q. Since g was arbitrary, this is the applicable definition for Q to be normal. Instead of this simple counting argument, Sylow theorems are often used to prove the uniqueness and normality of Q!

To continue analysis of a group G of order pq, we now have elements a, b of order p, q respectively. The subgroup P generated by a acts by automorphisms on Q, the unique subgroup of order q, which is generated by b. Specifically, we must have

$$aba^{-1} = b^j$$

for some $1 \leq j \leq q-1$. If $j = 1$ we have $ab = ba$ and in fact since G consists of pq distinct elements $a^i b^k$ (see Fact 3), G is therefore abelian. The theory of finite abelian groups (but see also below), then ensures that up to isomorphism, there is only one such group, the internal direct sum of P and Q, or external direct sum of a group of order p and a group of order q. Thus a case of potential interest arises only when $j \neq 1$. We write the conjugation automorphism as $H(b) = aba^{-1} = b^j$. Now consider the set $U_q = \{1, \ldots, q-1\}$, which are precisely the exponents of b that can arise through the action of H on b. That is, U_q constitutes the possible values of j. Now U_q is an abelian group of order $q-1$ whose group operation is integer multiplication modulo (q) and whose multiplicative identity is 1. The only group property to check that does not follow immediately from the arithmetic of the integers is the existence of an inverse. But since any "integer" in U_q is co-prime to q, given j the congruence

$$fj + tq = 1$$

can always be solved for integers f, t by the Euclidean algorithm. The class modulo (q) represented by f is then the multiplicative inverse of j. Consider the order of j in this abelian group. Writing the automorphism on exponents, $H(1) = j$, and iterated, $H^k(1) = j^k$, referring if necessary to the original definition of H. But $a^p = e$ so $H^p(1) = 1$, thus j has order dividing p, hence p. By Lagrange's Theorem from Fact 1, we must have $p \mid q-1$. Another way of thinking about this is to see that if p did not divide $q-1$, then by the p-fold iteration of H, we would obtain

a relation $b = b^i$ for some $i \neq 1 \mod q$ which would mean that Q had less than q elements.

The fact that $kp + 1 = q$ also follows from a Sylow argument, that the number of subgroups of order p in G which we noted above in Proposition 2 is $kp + 1$ must divide $n = pq$.

The first problem now becomes to show that there is a choice of j for $H(1)$ whose order is p. This follows from Cauchy's Theorem above applied to $p \mid q - 1$. Thus since $j \neq 1$ there esists a non-abelian group of order pq, the *semi-direct sum* of P and Q. The second is whether this choice of j affects the (isomorphism) class of G. Proving this is straightforward due to the rigid cyclic structure of U_q. We examine this structure and in so doing uncover without much effort a number of standard results in Algebra.

Proposition 4. *Assuming that the abelian group U_q has a multiplicative generator, if $p \mid q - 1$ then there is up to isomorphism one non-abelian group of order pq.*

Proof. We already showed a converse, namely that if such a non-abelian group exists, then $kp + 1 = q$ for some integer k. The rest of the proof will show that the given group is up to isomorphism, unique. Choose a generator $u \in U_q$, and let $k \in U_q$ be such that $kp = q - 1$. Then the element $j = u^k$ has order p in U_q, and the automorphism $H(b) = b^j$ gives rise to a consistent multiplication table on the distinct elements $a^r b^s$ of G, $1 \leq r \leq p$, $1 \leq s \leq q$. Thus a non-abelian group of order pq exists under the hypothesis. By the structure of a cyclic group, all the elements of U_q of order p are now of the form $u^{kt}, t = 1, \ldots, p - 1$. Thus we obtain a group G' of order pq by using the automorphism $H'(1) = j^t = u^{kt}$. But we claim that this G' is isomorphic to G (which resulted from the choice $t = 1$). In fact define a mapping $\phi : G' \to G$ which takes the subgroup generators $a' \mapsto a^t$, $b' \mapsto b$, and defines a one-to-one correspondence of the elements $a^r b^s$ consistent with the structure imposed by H, resp. H' and is the required isomorphism.

For completeness, we review enough of the structure theory of finite abelian groups to demonstrate that the multiplicative group of integers modulo a prime q is cyclic. This is not a waste of effort, since a new point of view is introduced. Furthermore, the method is strong enough immediately to conclude also that "the multiplicative group of any finite field is cyclic".

Decomposition of a Finite Abelian Group and the Smith Form

Any finite abelian group G, which we will write additively, can be characterized by a *presentation*, consisting of generators and relations. The set of generators $A = \{\xi_i\}$ forms a subset of the underlying set of G. The relations are words involving the $\{\xi_i\}$, for example

$$w = \sum \mu_i \xi_i \,, \quad \mu_i \in \mathbb{Z}$$

is a word that gives a relation when set to 0. Here we make no distinction between say $a + a$ and $2a$, so we are really treating abelian groups as modules over the integers $\mathbb{Z}$. Let $R = \{w\}$ be a sufficient set of relations for G. Thus G is isomorphic to the free abelian group (or $\mathbb{Z}$-module $\mathbb{Z}A$, modulo the image of the relations $\mathbb{Z}R$.

For example the group multiplication table (written additively) gives rise to a sufficient set of relations. The entry in this table corresponding to row a and

column b is the element c, if and only if $a+b=c$ in G. The corresponding relation $w=a+b-c\ (=0)$. Since the group table is a symmetric matrix of G-values, only relations arising from the main diagonal and above are necessary. We now write all the relations in a different kind of table or matrix. Label n columns by the given symbols $\{\xi_i\}$, $i=1,\ldots,n$, and label m rows by a sufficient set of relations known to hold in G, *e.g.,* those arising in the above manner from the group table. The example relation is now represented by a row consisting entirely of 0's except for a 1 in the columns denoted by a and b and a -1 in the column corresponding to c. We desire to simplify this $m\times n$ matrix M to extract information about G. Entries in M are integers which will be ordered according to absolute value.

We give some *allowable operations* on M, which always result in another relation matrix for an abelian group isomorphic to G. These are the usual row and column operations used in linear algebra to achieve row- and column-echelon form. No special accounting need be made of the rows/relations, but an all-zero row should be made the last relation by interchange, or simply disgarded. Labeling of the columns is however important. A column exchange $v_i \leftrightarrow v_j$ corresponds to an exchange of generators also according to $\xi_i \leftrightarrow \xi_j$, or more precisely, the new matrix D' arises from $\xi_i' = \xi_j$, $\xi_j' = \xi_i$. A column operation (addition) $v_i' = v_i - \beta v_j$ corresponds to the relabeling $D' : \xi_i' = \xi_i$, $\xi_j' = \beta\xi_i + \xi_j$. If now a 1 or -1 occurs alone in a relation (all other entries of the row being 0), the exceptional column must correspond to the identity element 0 and thus both column and row can be discarded. A zero column cannot occur, as this would violate finiteness of the group.

Algorithm I. *Put a given presentation matrix for a finite abelian group into diagonal form.*

Thus we have a matrix M which has no all-zero rows or columns. Row operations of the form $e_j' = e_j - e_i$, corresponding to the subtraction operation of smaller from larger number in the Euclidean algorithm may be performed. These operations produce a single non-zero entry in the first column, namely the greatest common divisor (gcd) in $\mathbb{Z}$ of all elements of the first column. This gcd integer, say α_{k1}, may also turn out to be equal to the gcd of the k-th row as well. If this is so, we may exchange the first and k-th rows and use the new element α_{11} to *zero out* the entire first row and first column, except for the (1,1) entry. Then we continue in the same manner on the matrix L_{11} consisting of the remaining rows and columns of D.

If the element α_{k1} is not the gcd of its row, we use allowed column operations on the current presentation matrix M' to produce a gcd, say α_{ks} of the k-th row. In particular we have $\alpha_{ks} \mid \alpha_{k1}$. If α_{ks} happens to be equal to the gcd of column s, we move row k and column s to row 1 and column 1, and zero out the first row and column, except for α_{11}, and proceed with L_{11} as before. If not, we continue to find the gcd of a succession of rows and columns. Each gcd obtained divides the previous one, so eventually we obtain two successive gcd's which are equal. The first of this pair of equal values is seen to be the gcd of both its row and its column, which can be moved to row 1 and column 1 as before, applying the same procedure to L_{11}. This finally results in D being in a diagonal form, where we may have

discarded zero rows, and columns corresponding to the identity element. Thus,

$$D = \begin{bmatrix} \alpha_{11} & 0 & & 0 \\ 0 & \alpha_{22} & & \\ & & \ddots & \\ 0 & 0 & & \alpha_{kk} \end{bmatrix}.$$

We say that D expresses a *diagonal presentation* of G. A set of software routines is provided that realize such a diagonalization for any group presentation, in addition to the other constructions discussed here.

Lemma O. *Given a diagonal presentation D for G as above, for fixed indices $1 \leq i < j \leq k$ let*

$$g(i,j) = \gcd(\alpha_{ii}, \alpha_{jj})$$
$$l(i,j) = \operatorname{lcm}(\alpha_{ii}, \alpha_{jj}) = \alpha_{ii} \cdot \alpha_{jj} / g(i,j).$$

Then D' which results from D when the i-th diagonal entry is replaced by $g(i,j)$ and the j-th diagonal entry is replaced by $l(i,j)$ is also a diagonal presentation for G. In particular if α_{ii} and α_{jj} are relatively prime, let we obtain such a diagonal presentation matrix by setting $D'_{jj} = \alpha_{ii} \cdot \alpha_{jj}$, and deleting the i-th row and column (which both contain a single 1).

Proof. To modify D we perform row and column operations only on rows labeled i, j and columns labeled i, j. These operations are induced from operations performed on the 2×2 matrix

$$D_{ab} = \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix},$$

where we simplified the notation. Using the integers f, g given in the formula

$$fa + gb = g(i,j) = \gcd(a,b) \quad ,$$

we may transform D_{ab} into

$$D'_{ab} = \begin{bmatrix} a & g \\ 0 & b \end{bmatrix}.$$

Since $g \mid a$ and $g \mid b$, we may "zero out" the $(1,1)$ and $(2,2)$ entries to obtain

$$\begin{bmatrix} 0 & g \\ -ab' & 0 \end{bmatrix}$$

where $b' = b/g$. Now permuting the variables corresponding to columns i and j and multiplying row j by the multiplicative unit (-1) gives the desired form for D'

$$\begin{bmatrix} g & 0 \\ 0 & h \end{bmatrix},$$

where $h = \operatorname{lcm}(a,b)$. If $g = 1$ we may erase the i-th row and column.

Definition. *The finite abelian group G is said to be of relatively prime decomposition if in any diagonal presentation such as D above, the respective diagonal entries are pairwise relatively prime.*

Proposition 5. *The group of units (invertible elements) U_q in the ring $\mathbb{Z}_q$ is of relatively prime decomposition.*

Proof. U_q is in fact embedded in the finite field $\mathbb{Z}_q$. Therefore, for any integer $s \geq 0$ the equation $x^s = 1$ in U_q can have at most s solutions, by polynomial factorization in the field. But suppose that in some diagonal decomposition as above, U_q has two entries α_{ii} and α_{jj} with a common factor $t > 1$. Let

$$\beta_i = \alpha_{ii}/t$$
$$\beta_j = \alpha_{jj}/t$$

and let ξ_i, ξ_j be the generators corresponding to these columns. Then $\eta_i = \beta_i \xi_i$ has order t and in fact so do $\nu \cdot \eta_i$, $\nu = 1, \ldots, t-1$. Similarly $\eta_j = \beta_j \xi_j$ has order t along with $\nu \cdot \eta_j$, $\nu = 1, \ldots, t-1$. Since $t \geq 2$, we now have $2t - 1 > t$ solutions of $t \cdot x = 0$ in U_q (writing the group operation additively), which violates the fact observed above that the number of these solutions can be at most t.

Observation 1. *The same proof shows that any finite (abelian) group which is embedded in the multiplicative group of a field, must be of relatively prime decomposition. In particular, the multiplicative group Ω_{p^r} of $GF(p^r)$ is of relatively prime decomposition.*

Let us record a few needed facts about abelian (and cyclic) groups.

Fact 4. *In a finite abelian group G, the elements of order p, with the identity, form a subgroup H. In particular, if there are two distinct subgroups of G with order p, H contains at least $p + 1$ elements.*

Proof. Since G is abelian we see that the product of two elements of order p has order p or 1. This property is also preserved under group inverse. Thus the union of all subgroups of G that have order p is a subgroup H. If there were at least two of these subgroups of order p, the resulting union has more than p elements.

Fact 5. *Let p be a prime which divides both α_{ii} and α_{jj} for some $1 \leq i < j \leq k$ in a diagonal presentation D for the group G. Then G has at least two distinct subgroups of order p.*

Proof. As in the proof of Proposition 5, the elements $\{\nu\eta_i\}$, $\nu = 0, \ldots, p-1$ form a subgroup $F_i \subset G$, where $\eta_i = (\alpha_{ii}/p) \cdot \xi_i$. Similarly we form the subgroup $F_j \subset G$ also of order p. If $\eta_i = \mu \cdot \eta_j$, we have, using notation of Proposition 5, $\beta_i \xi_i = \mu \beta_j \cdot \xi_j$. Since $\{\xi_f\}$ independently generate G according to the diagonal presentation, we must in particular have β_i divisible by α_{ii} which is impossible since $p > 1$. Hence $\eta_i \in F_i$ is not in F_j and these two groups are distinct.

If G is a finite abelian group and a an integer, denote by aG the subgroup consisting of all elements $\{a \cdot g\}$ where $g \in G$. Here we use additive notation. The abelian group G is the *direct sum* $H \oplus K$ of subgroups $H \subset G$, $K \subset G$ if any element of G is expressible *uniquely* as the sum of an element of H and an element of K. Equivalently, H and K generate G, and $H \cap K = \{0\}$. Extending the definition of direct sum to finitely many factors, we see that a diagonal presentation expresses a group as the direct sum of finitely many cyclic subgroups.

Observation 2. *Thus Algorithm I asserts that any finite abelian group is a direct sum of cyclic subgroups.*

Fact 6. *If G is the direct sum of H and K, then aG is the direct sum of aH and aK.*

Proof. For an arbitrary $g \in G$, take $g = h + k$ H and K fill G. Then by associativity of group addition, $ag = ah + ak$, so aH and aK fill aG. For uniqueness, H and K intersect in the identity element, and so *a fortiori* do aH and aK.

Next let G be cyclic of order n, $N(G) = n$, $k = \gcd(a, n)$, and $r = \frac{n}{k}$.

Fact 7. *Given an integer a, the order of $H = aG$ equals r.*

Proof. If x is a generator of G, then $a \cdot x$ is a generator of $H = aG$, and we claim so is $k \cdot x$. The integer a is a multiple of k, so $aG \subset \{kx\}$, the cyclic subgroup generated by kx. But also $k = \ell a + mn$ for some integers ℓ, m, so $kG \subset aG$ also and the two subgroups are the same. But for $1 \leq f < r = n/k$, $1 \leq fk < n$ so $fk \cdot x \neq 0$. But $rk \cdot x = nx = 0$, so the order of aG is $r = \frac{n}{\gcd(a,n)}$,

Lemma A. *In a finite cyclic group K, if every element $\neq e$ has the same order, then this order equals the order of the group, and is prime.*

Proof. Let u be a generator of K of order $n = ab$ where $a, b > 1$. Then $v = u^a \neq u$, else u has a smaller order $a - 1$. But now v has a smaller order b which contradicts the hypothesis.

Lemma B. *Any subgroup H of a finite cyclic group G is cyclic.*

Proof. Let G be generated by u, and if H is not the trivial group, H be generated by

$$\beta_1 = u^{a_1}, \ldots, \beta_k = u^{a_k},$$

where $1 \leq a_i \leq p$ for $i = 1, \ldots, k$. Then following the Euclidean algorithm for forming the gcd, we apply the corresponding succession of group multiplications (and inverses) on $\{\beta_j\}$, obtaining the element $\gamma = u^c \in H$, where $1 \leq c \leq p$, and c is $\gcd(a_1, \ldots, a_k)$. Then γ generates H which establishes this well-known Lemma.

Proposition 6. *A finite abelian group G is of relatively prime decomposition if and only if it is cyclic. Thus for q a prime, the group of units U_q has a single generator. Also the multiplicative group Ω_{p^k} of the finite field $GF(p^k)$ is generated by a single element.*

Proof. For the final part, Proposition 5 and Observation 1 showed that $U_q \simeq \Omega_q$ as well as Ω_{p^k} are of relatively prime decomposition, so our result will show that these abelian groups are finite cyclic (have a single generator). Suppose the group G is cyclic and has a primary decomposition D with two diagonal entries m_1 and m_2 that have a common prime factor p. Next we see by Fact 5 that there are at least two distinct subgroups of G of order p. Let $H \subset G$ consist of all elements of order p, which is a subgroup by Lemma A , since G is abelian. Then by Lemma B, H is cyclic of order p. But by Fact 4, H must have more than p elements, which is a contradiction. Hence G must be of relatively prime decomposition.

Conversely, if G is of relatively prime decomposition, choose a diagonal presentation matrix D for G. If D is a 1×1 matrix $[d]$, then G is cyclic of order d. If the order of D is greater than 1, choose two entries on the diagonal a, b. We affect the

matrix D through row and column operations affecting only the rows and columns that contain a and b. By assumption, $\gcd(a,b) = 1$. The part of the matrix of interest looks like

$$D_{ab} = \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}.$$

But by Lemma 0 in the "relatively prime" case, one can change this matrix by allowable transformations into

$$\begin{bmatrix} 1 & 0 \\ 0 & ab \end{bmatrix}.$$

By performing the corresponding transformations on D and deleting the row and column containing the entry 1, we obtain a smaller matrix D' which also gives a diagonal presentation for G, and whose diagonal entries are still relatively prime in pairs. The indicated procedure is applied successively to produce the 1×1 presentation matrix of the given group G as

$$E = \left[\prod_{i=1}^{k} a_{ii} \right].$$

Thus G is cyclic with a single generator of order $\tau = \prod_{i=1}^{k} a_{ii}$, the product of the diagonal entries of D as was to be proved.

Suppose now that we are given a presentation of a finite abelian group G in a diagonal form $D = [a_{ii}]$. We may eliminate diagonal entries equal to 1 by removing the row and column of each such entry. Subsequently if we have the successive division property holding,

$$\alpha_1 \mid \alpha_2 \mid \cdots \mid \alpha_k,$$

we say that D is a *Smith form* matrix. This is a unique invariant: if Smith forms are different, the corresponding abelian groups are not isomorphic. We demonstrate this result as Proposition 7 below.

A *primary decomposition* is a diagonal presentation for G where each of the diagonal entries is a prime power. Suitably ordered, say by increasing prime and increasing power, the primary decomposition P can be made to be a unique invariant for G. The entries of P are the *primary invariants* of G.

Suppose that we have a diagonal presentation D derived from some presentation matrix M for G. Decompose the diagonal entries α_{ii} of D into factors consisting of maximal prime powers. If an entry has more than one distinct primary factor, we may use Lemma 0 to construct a presentation matrix where "on average" the diagonal entries have fewer primary factors than before. To effect this operation requires us each time to augment the matrix with a row and column, and entry 1 on the diagonal. Repeating this finally gives a diagonal matrix with purely primary entries. By performing the suitable ordering of the diagonal entries, we obtain the primary decomposition P. This large matrix can now be transformed into s Smith form matrix. Use the "relatively prime" case of Lemma 0, so that the primary factors are arranged in such a way that the "division property" of the Smith form holds.

Note that the primary decomposition matrix can be much larger as a square matrix than the original diagonal matrix D, or for that matter a given non-diagonal presentation matrix M for G. Furthermore, to obtain a primary decomposition

requires factorization in the ring of integers. One must generate lists of primes, do many trial divisions with numbers not directly computable from the original presentation D. From the standpoint of algorithms, this is computationally intensive, both in terms of number of operations performed and amount of memory required. A moment's thought shows why the Smith form matrix S is not larger than D however (we will show this in a different way in Proposition 7), so it is desirable to construct S directly from D without requiring any data storage outside the matrix itself, any operations other than elementary arithmetic (especially division with remainder), and any branching steps other than "compare with 0".

Such a simpler way to derive the Smith form from a given diagonal presentation exists, which in particular does not require any prime factorization step. This is significant in the case under consideration, of abelian groups, where we must factor in the integers $\mathbb{Z}$. It is also very significant in constructing the Smith form in scenarios of interest to control engineering, where the matrix entries are elements of $F[x]$, that is, polynomials over a given field.

Algorithm II. *Put a diagonal presentation into the Smith form.*

Given a diagonal matrix D, we may perform the transformations indicated by Lemma O by choosing a pair of diagonal indices $1 \leq i < j \leq k$, where k is the order of D. In the i-th diagonal slot place $\gcd(a, b)$ and in the j-th diagonal slot place $\mathrm{lcm}(a, b)$, where a, b are the entry values. Then proceed to the next required operation. In fact we use successively all the pairs of indices, ordered lexicographically. According to Lemma 0, each transformation may be performed by a sequence of row and column operations. Thus on a 3×3 matrix we employ Lemma 0 three times, corresponding to the pairs $(1, 2)$, $(1, 3)$ and $(2, 3)$. Finally superfluous 1's on the diagonal with row and column can be removed. The resulting matrix is easily seen to be in the Smith form.

Note that in Algorithm II use is made only of efficiently computed functions such as the greatest common divisor (according to Euclid's algorithm). In a Smith form presentation

$$D = \begin{bmatrix} z_1 & 0 & & 0 \\ 0 & z_2 & & \\ & & \ddots & \\ 0 & 0 & & z_k \end{bmatrix}.$$

$\{z_i\}$ are called the *invariants* or *Smith invariants* of the finite abelian group G. As a sequence of integers greater than 1, satisfying the "division property", $\{z_i\}$ uniquely characterizes G among abelian groups. Given a presentation matrix D for G, an explicit formula for the $\{z_i\}$ may be given in terms of minor determinants of D, see [3]. In fact, $z_i = \dfrac{\sigma_i}{\sigma_{i-1}}$, where $\sigma_i = \gcd\{\Gamma_i\}$, where Γ_i ranges over the values of all of the $i \times i$ minors of D. Thus σ_i is a generator of the ideal of $i \times i$ minors.

We prove that the Smith form is such a unique invariant for G. The proof is more elementary than the corresponding textbook proof for the primary decomposition. Summing up, Algorithms I and II combine to solve the problem of determining, using only elementary arithmetic operations, bounded memory, and without any lists or random choices, whether two presented finite abelian groups are isomorphic. In particular, given the multiplication table for an abelian group, we can derive without any factorizing step, the unique Smith form for the group.

Proposition 7. *Let H and K be finite abelian groups with Smith form matrices S and T respectively. Then H is isomorphic to K, $H \simeq K$ if and only if $S = T$.*

Proof. Suppose that $S = T$. Consider H a presented by a matrix M_H, such as a group multiplication table. With obvious notation, the Smith form S, respectively T, is derived from M_H, respectively M_K, by allowable operations that certainly preserve the group isomorphism class of the group presented at each stage. By transitivity of isomorphism, we have $H \simeq K$.

Conversely we now suppose that H and K are isomorphic groups. In particular the orders are equal: $Q = N(H) = N(K)$. If the matrix size of the square matrices S and T both equal 1, then $S = T$, since the sole entry equals the group order. Writing $o(S)$ for the matrix size of S, we set $o(S) = s$, $o(T) = t$, one of which is greater than 1. We can put a *linear ordering* $<$ on pairs of matrix sizes by means of the Cartesian (lexicographic) ordering induced from natural ordering of the integers. By an induction hypothesis we assume that for a pair (s', t') strictly smaller than (s, t), that always $o(S) = s'$, $o(T) = t'$ and $H \simeq K$ imply that $S = T$. Alternatively we could take as induction hypothesis that $S = T$ whenever the group involved has order less than Q. Denote by $D = \mathrm{diag}[v_i]$ the $t \times t$ diagonal matrix whose i-th diagonal entry is the i-th component of the vector $(v_1, \ldots, v_t)$.

Let $S = \mathrm{diag}[a_i]$, $T = \mathrm{diag}[b_j]$, $1 \leq i \leq s$, $1 \leq j \leq t$. Of course by definition of Smith form, we have $a_1 > 1$, $b_1 > 1$, as well as $a_i \mid a_{i+1}$ and $b_j \mid b_{j+1}$ for $1 \leq i < s$, $1 \leq j < t$. Note that

$$N(H) = \prod_{i=1}^{s} a_i = \prod_{j=1}^{t} b_j = N(K).$$

Now let $a = a_1$, $b = b_1$. First suppose that a, b are relatively prime. Consider the subgroup $L = aH \subset H$. We have a diagonal presentation and thus a direct sum decomposition of H and K, to which we can apply Fact 6 and Fact 7. From the presentation S we obtain a diagonal presentation for aH, $S' = \mathrm{diag}[c_j]$ with $c_i = \frac{a_i}{a}$. From the presentation T we may construct a diagonal presentation for $L \simeq aK$, $T' = \mathrm{diag}[d_j]$, $d_j = \frac{b_j}{k_j}$, $k_j = \gcd(a, b_j)$. We have for $j = 1, \ldots, t$, $k_j \leq a$, hence $d_j \geq \frac{b_j}{a}$ and also $k_1 = 1 < a$, hence $d_1 > \frac{b_1}{a}$. Therefore,

$$\text{(X)} \quad N(H)\frac{1}{a^s} = \prod_{i=1}^{s} a_i \cdot \frac{1}{a^s} = \prod^{s} c_i = N(aH)$$
$$= N(aK) = \prod^{t} d_j > \prod_{j=1}^{t} b_j \cdot \frac{1}{a^t} = N(H)\frac{1}{a^t},$$

thus $a^s < a^t$, and $t > s$. But by the same argument applied to bK we must conclude that $s > t$. This contradiction implies that a and b cannot be coprime, and have a common factor $g > 1$.

So consider the subgroup $gH \simeq gK$ which has presentations U and V calculated according to Fact 6 and Fact 7. That is, $U = \mathrm{diag}[\frac{a_i}{g}]$ and $V = \mathrm{diag}[\frac{b_j}{g}]$. But the

orders of the groups H and K, and subgroups gH and gK, are equal hence

$$\text{(Y)}\quad N(H)\frac{1}{g^s} = \prod_{i=1}^{s} a_i \cdot \frac{1}{g^s} = N(gH) = N(gK) = \prod_{j=1}^{t} b_j \cdot \frac{1}{g^t} = N(K)\frac{1}{g^t} = N(H)\frac{1}{g^t}.$$

Thus $s = t$ and we conclude that the *size* of the Smith matrix is uniquely determined by the (isomorphism class of the) group. From now on we may use the fact that S and T are square matrices of the same size.

Again examine presentations for the group $L = aH \simeq aK$. The diagonal presentations S, T amount to a direct sum decomposition of H, K, so we apply Fact 6 and Fact 7 to give diagonal presentations of aH and aK. The group aK is generated by $\{a\xi\}$ for all the generators ξ of K (columns of T). According to Fact 7, the induced relations are

$$b_j \longmapsto \frac{b_j}{\gcd(a, b_j)} = d_j.$$

Thus we have a diagonal presentation $T' = \mathrm{diag}[d_j]$ for $L = aH \simeq aK$. Similarly arising from aH is $S' = \mathrm{diag}[c_j]$, the other such diagonal presentation, with $c_i = \frac{a_i}{a}$. Now invoke an assumption to be refuted, that a *does not* divide b. In this case we have with the notation as above that $k_1 < a$, $k_j \le a$, and thus

$$\text{(Z)}\quad N(H)\frac{1}{a^s} = \prod_{i=1}^{s} a_i \cdot \frac{1}{a^s} = N(aH) = N(aK) = \prod_{j=1}^{s} \frac{b_j}{k_j} > \prod^{s} \frac{b_j}{a} = \prod_{j=1}^{t} b_j \cdot \frac{1}{a^s} = N(K)\frac{1}{a^s}.$$

In view of $N(H) = N(K)$, this strict inequality gives a contradiction to $k_1 < a$, so $a \mid b$, and by symmetry considering bH and bK, we have $b \mid a$, hence $a = b$. Now the induced presentations for $L = aH$, bK are

$$S' = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & & \frac{a_f}{a} & & \\ & & & & \ddots & \\ & & & & & \frac{a_s}{a} \end{bmatrix}, \qquad T' = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & & \frac{b_g}{b} & & \\ & & & & \ddots & \\ & & & & & \frac{b_t}{b} \end{bmatrix},$$

where $a_f > a$, $b_g > b$, and $s = t$. Eliminating the rows and columns where a 1 appears on the diagonal results in a Smith form matrix for L in either case. For the pair of matrices, at least one row and column was eliminated, so the induction hypothesis may be applied. This means that the resulting Smith form matrices may be assumed to be identical. We thus obtain $a_f/a = b_g/b, \ldots, a_s/a = b_t/b$. But we saw that $a = b$, so in fact $a_f = b_g, \ldots, a_s = b_t$. We immediately have that $s - f + 1 = t - g + 1$ and because $s = t$, we get $f = g$, so also $a_1 = a = b =$

$b_1, \ldots, a_{f-1} = a = b = b_{g-1}$. Therefore $a_i = b_i$ for $1 \leq i \leq s = t$, and the given Smith forms for isomorphic H and K must be identical as was to be proved.

Starting with an $m \times n$ presentation matrix S for an abelian group, we now have an algorithm which updates values only at the entries of S, so its storage requirements are strictly bounded. The algorithm requires only binary operations in processing those entries. The binary operations used can in principle be limited to addition and subtraction, though for practical coding purposes integer multiplication and division with remainder should be employed. Branching conditions consists only of "compare to zero". Proposition 7 implies that although many rectangular matrices of different sizes may present a finite abelian group G, only one Smith form matrix does so. Arbitrary choices made during Algorithm I may lead to different diagonal presentations, but they finally lead to the same Smith form matrix under Algorithm II.

Complete Invariant for a Module over a Euclidean Ring

Let us try to take the results of the last section a little farther. In particular it is our intention to generalize Proposition 7 to the case of any finitely generated torsion module over a Euclidean ring. That is, Proposition 7 establishes the result that the Smith form depends only on the isomorphism class of a given $\mathbb{Z}$-module (abelian group), and that it is indeed a complete invariant. Only the one (isomorphism class of) abelian group possesses the given Smith form.

Examining the proof of Proposition 7, we denote by Ω the product of the generators z_i of the annihilating $\mathbb{Z}$-ideals in the direct sum decomposition (diagonal presentation). Since the integers $\{z_i\}$ are taken as positive, we have immediately that $\Omega = N(G)$, the order of the group. Since a group isomorphism is *a fortiori* a set isomorphism, Ω must depend on the isomorphism class of G and not on the particular Smith form that was used to represent G.

Now consider the other well-known case of a Euclidean ring and its modules. Given a field $\mathbb{F}$ and $\mathbb{F}^n$ as an $\mathbb{F}[x]$-module, where $\mathbb{F}[x]$ is the ring of polynomials of finite degree. That is, we have the variable x act in standard fashion as a linear transformation T. We know that with respect to a given basis $\{\vec{e}_i\}$, $i = 1, \ldots, n$, the transformation T acts from the left as a matrix $T = [\alpha_{ij}]$, by overloading the notation. In this case the natural relation matrix for the module $M \simeq \mathbb{F}^n$ is given by

$$xI - T = \begin{bmatrix} x - \alpha_{11} & -\alpha_{12} & \cdots & -\alpha_{1n} \\ & x - \alpha_{22} & & \\ \vdots & & \ddots & \vdots \\ & & & x - \alpha_{nn} \end{bmatrix}.$$

Then $\det[xI - T]$ is a generator in $\mathbb{F}[x]$ of the n-th Fitting ideal of $\mathbb{F}[x]$, [3].

Recapitulate that $\mathcal{R} = \mathbb{F}[x]$ is a Euclidean domain and hence has only principal ideals. It is convenient to define the Euclidean height to have the multiplicative property $\hbar(ab) = \hbar(a) \cdot \hbar(b)$ as in [3]. In fact we can define for a general commuative domain $\mathcal{R}$, $\hbar : \mathcal{R}^+ = \mathcal{R} \setminus \{0\} \rightarrow \mathbb{Z}^+$ as a homomorphism of the multiplicative semigroups, and for completeness define $\hbar(0) = 0$. For $\mathcal{R}$ now to be a *Euclidean ring*, the additional axiom is required that allows the Euclidean algorithm to be carried out, which mechanically produces the gcd of two elements. That is, given

a, $b \in \mathcal{R}$, there are q, $r \in \mathcal{R}$ with $\hbar(r) < \hbar(b)$ such that

$$a = q \cdot b + r. \tag{E}$$

Hence for the ring $\mathbb{F}[x]$ the height $\hbar(f(x))$ should not be given as the polynomial degree as in [2] but rather as $2^{\deg(f)}$ as in Keating, [3]. There follows from the semigroup homomorphism property that $\hbar(u) = 1$ where u is a unit. The units in $\mathbb{F}[x]$ are the non-zero scalars ($\alpha \in \mathbb{F}$). Ideals are in one-to-one correspondence with the equivalence classes of elements of $\mathcal{R}$ under *association*, that is, multiplication by a unit. For such a module, a diagonal presentation and hence a Smith presentation can be formed exactly as before by Algorithms I and II. The diagonal entries $\{z_i\}$ are defined up to a multiplicative unit.

We see that the Fitting generator $|xI - T|$ depends only on the module M over $\mathbb{F}[x]$, up to isomorphism. In fact the vector spaces $\mathcal{V} \simeq \mathbb{F}^n$ and $\mathcal{W} \simeq \mathbb{F}^m$ must be linearly isomorphic, since an $\mathbb{F}[x]$-module isomorphism $\psi : \mathcal{V} \to \mathcal{W}$ must preserve addition and scalar multiplication. Thus first of all, $\dim \mathcal{W}$ must equal n also. Let B be a non-singular matrix representing $\psi : \mathcal{V} \to \mathcal{W}$ with respect to the given basis. Then as T respresents the action of x on $\mathcal{V}$, its action on $\mathcal{W}$ is represented by BTB^{-1}. We also see that the isomorphism ψ amounts to a change of basis for the vector space $\mathcal{W}$. The n-th Fitting generator, otherwise known as the characteristic polynomial of T, for the action on $\mathcal{W}$ is then

$$\det[xI - BTB^{-1}] = \det(B \cdot [xI - T] \cdot B^{-1}) = \det B \det[xI - T] \det B^{-1} = \det[xI - T].$$

Thus we recover the fact that the characteristic polynomial is the same regardless of choice of basis. The final Smith form generator $z_n(x)$ is a polynomial divisible by all the other $z_i(x)$, so z_n *annihilates* all of M. No polynomial of lesser degree could annihilate all of $\dfrac{\mathbb{F}[x]}{(z_n)}$, so z_n is the minimal polynomial, from linear algebra.

At this point we can repeat the proof of Proposition 7 to obtain the invariance of the whole Smith form, and all of the divisors z_i, not merely their product Ω. The only difference is that in the proof we compare elements according to the Euclidean height for $\mathbb{F}[x]$, "two power the degree", rather than the height for the ring $\mathbb{Z}$, the absolute value.

The invariance of $\Omega = \sigma_n$ (in other notation above), that is the n-th Fitting generator, was "obvious" in the case of finite abelian groups. Given an $n \times n$ presentation D for an abelian group G, it is not so obvious that $\det D$ is up to sign, equal to the order of G, $N(G)$. But once D has been transformed by elementary operations to any diagonal presentation, it *is* obvious. Similarly in the "linear algebra" situation, with $\Omega(x) = \sigma_n(x)$ being the characteristic polynomial defined up to a non-zero scalar multiple, the invariance is known to students completing a senior-level course in matrices or linear algebra.

In either of these two principal cases, since we have the invariance of Ω, the well-definedness of the sequence of Smith ideals $(z_1), \ldots, (z_n)$ follows as in Proposition 7. Without a doubt, the ways one concluded the invariance of Ω in the $\mathbb{Z}$ case and in the $\mathbb{F}[x]$ case were quite different.

Algorithms I and II apply without any significant change, to any finitely generated *torsion* module M over a Euclidean ring $\mathcal{R}$. Thus for $m \in M$ there is a non-zero $r \in \mathcal{R}$ such that $r \cdot m = 0$. A finite abelian group is a f.g. torsion $\mathbb{Z}$-module and a linear transformation T of a finite dimensional vector space $\mathcal{V}$ over a field $\mathbb{F}$ gives a f.g. torsion $\mathbb{F}[x]$-module.

Thus, a Smith form, or equivalently a sequence of ideals $\mathcal{S} = (z_1), \ldots, (z_k)$ with $z_i \mid z_{i+1}$, $a \leq i < k$ such that

$$M \simeq \frac{\mathcal{R}}{(z_1)} \oplus \cdots \oplus \frac{\mathcal{R}}{(z_1)},$$

and this form is invariant up to the matrix equivalence of P, the presentation matrix that we started with. In fact we have

Theorem M. *The Smith sequence $\mathcal{S}$ depends only on the module type of M (isomorphism class).*

This result is found in [3], an excellent reference for the theory of torsion modules over a Euclidean ring. The proof there of Theorem M utilizes the decomposition of ring elements into powers of prime or irreducible elements. The proof illustrates the interplay between "elementary divisor" decomposition for M, where the direct summands are cyclic with annihilating ideal a "primary" ideal, and the "invariant factor" or Smith form decomposition.

In the present alternative discussion we take the point of view of avoiding the Fundamental Theorem of Arithmetic in the Euclidean ring $\mathcal{R}$, or any explicit use of primary decompositions. We may hope that this is justified in view of the success of Proposition 7 in establishing the complete invariance of the Smith form for the abelian group case and the "linear algebra" case. The only essential data that is now missing from the general Euclidean case is that, up to association, the k-th Fitting generator $\Omega = z_1 \cdots z_k \in \mathcal{R}$ has the same value for any isomorphic $M' \simeq_{\mathcal{R}} M$. Of course Ω can be directly computed from any presentation matrix P of M simply by taking the determinant.

In view of the fact that by Algorithm I, any presentation for M or M' can be reduced to allowable operations to a diagonal presentation of an isomorphic module, without changing the determinant, it is sufficient to prove the invariance of Ω, up to association, in the case of two diagonal presentations. Througout the proof of Theorem M we must be careful to associate Ω to a presentation and not to a module type. If we say for example, that the n-th Fitting generator of the module $\mathcal{M}$ is $\Omega(\mathcal{M})$, we mean that value of Ω at some particular presentation for $\mathcal{M}$ which has been exhibited. We start out with explicit direct sums of cyclic modules (diagonal presentations), and all modules $\mathcal{M}$ that arise will be presented. These presentations are what allows us to compute $\Omega(\mathcal{M})$.

Claim 1. *Let $\mathcal{R}$-modules $\mathcal{U}_1$ and $\mathcal{U}_2$ be explicitly presented by*

$$\mathcal{U}_1 = \frac{\mathcal{R}}{(x_1)} \oplus \cdots \oplus \frac{\mathcal{R}}{(x_s)},$$
$$\mathcal{U}_2 = \frac{\mathcal{R}}{(y_1)} \oplus \cdots \oplus \frac{\mathcal{R}}{(y_t)},$$

and suppose that as modules $\mathcal{U}_1 \simeq \mathcal{U}_2$. Then $\Omega(\mathcal{U}_1) = \prod_{i=1}^{s} x_i = \prod_{i=1}^{t} y_i = \Omega(\mathcal{U}_2)$.

See also Problem 12.6 of [3], where the uniqueness of the elementary divisors (primary factors) is known. In fact to prove Claim 1 it is enough to show the following.

Lemma C. *Given a module homomorphism $\phi : \mathcal{U}_1 \to \mathcal{U}_2$ between two diagonally presented modules as above, which is injective, we have $\Omega(\mathcal{U}_1) \mid \Omega(\mathcal{U}_2)$.*

If we establish Lemma C, Claim 1 is proven since the inverse of an isomorphism ϕ, $\phi^{-1} : \mathcal{U}_2 \to \mathcal{U}_1$ is injective and hence $\Omega(\mathcal{U}_2)$ divides $\Omega(\mathcal{U}_1)$ also, so the two Fitting generators are associate, verifying Claim 1.

The key technical result that allows us to prove Lemma C is the following.

Lemma D. *Let $\mathcal{V} = \dfrac{\mathcal{R}}{\mathfrak{g}}$ be a cyclic module and consider an injective homomorphism*

$$\psi : \mathcal{V} \to \mathcal{U} \simeq \frac{\mathcal{R}}{\mathfrak{b}_1} \oplus \cdots \oplus \frac{\mathcal{R}}{\mathfrak{b}_k}.$$

Then $\mathcal{W} = \dfrac{\mathcal{U}}{\psi(\mathcal{V})}$ is a finitely generated torsion $\mathcal{R}$-module. Furthermore,

1. *$\mathcal{W}$ has a presentation of size $\leq k$,*
2. *the element $g\mathcal{R}$ generating $\mathfrak{g}$ divides $\Omega(\mathcal{U})$ and we can compute $\Omega(\mathcal{W}) = \dfrac{\Omega(\mathcal{U})}{g}$ (division defined up to multiplicative unit).*

Before proving Lemma D, we show how it implies that Lemma C is valid. Let us remark that in what follows we may be a little relaxed with the terminology. Working with principal ideals in a Euclidean ring, it is only human to confuse an ideal and its generating element, or to claim that two elements are equal when they are merely associate. Thus the ideal generated by x should be (x), but a generator of the ideal Δ is sometimes just written as Δ. We strive to be precise in such assertions, but believe that any remaining oversights of this nature should not affect the clarity, or validity, of the main results.

So given $\phi : \mathcal{U}_1 \to \mathcal{U}_2$ injective, consider the restriction ϕ_f to $\frac{\mathcal{R}}{(y_1)}$, the first summand of $\mathcal{U}_1$. Since ϕ_f is injective, we can apply Lemma D. By part (1) of the conclusion, the quotient

$$\mathcal{W} = \frac{\mathcal{U}_2}{\phi_f(\mathcal{U}_1)}$$

by the image of the first summand, itself has t or fewer summands in some diagonal presentation. Let $\hat{\mathcal{U}_1}$ be the complement of $\mathcal{R}(x_1)$ in the direct sum and ϕ_ℓ the restriction of ϕ to $\hat{\mathcal{U}_1}$. Since ϕ is injective, also the composition of ϕ_ℓ with the projection mapping $\pi : \mathcal{U}_2 \to \mathcal{W}$ is injective. We employ induction on the number of summands in $\mathcal{U}_1$, namely s. The induction works since the ground case for $s = 1$ is essentially Lemma D itself. Also by part (1) of Lemma D, a module $\mathcal{W}$ that arises in the construction can always be rewritten (after applying allowable operations) as a diagonal presentation (direct sum) with equal or fewer summands than the previous target module considered. Thus the induction hypothesis allows us to assume that $\Omega(\hat{\mathcal{U}_1}) \mid \Omega(\mathcal{W})$, but by part (2) of Lemma D, $\Omega(\mathcal{U}_2) = \Omega(\mathcal{W}) \cdot x_1$. Also $\Omega(\mathcal{U}_1) = \Omega(\hat{\mathcal{U}_1}) \cdot x_1$. Putting these three formulas together yields $\Omega(\mathcal{U}_1) \mid \Omega(\mathcal{U}_2)$ as needed to verify the conclusion of Lemma C, and also Claim 1.

Proof of Lemma D. Given the injective mapping $\psi : \dfrac{\mathcal{R}}{\mathfrak{g}} \to \mathcal{U} \simeq \dfrac{\mathcal{R}}{\mathfrak{b}_1} \oplus \cdots \oplus \dfrac{\mathcal{R}}{\mathfrak{b}_k}$, we can consider each summand component of $\psi(1)$, the image of the unity 1 of $\dfrac{\mathcal{R}}{\mathfrak{g}}$ in $\mathcal{U}$.

That is, the projection onto the j-th summand of $\psi(1)$ is called $\mathfrak{v}_j$. This element is really a coset of the ideal $\mathfrak{b}_j$ in $\mathcal{R}$. That coset is mapped by multiplication into the ideal $\mathfrak{b}_j$ by some elements $y \in \mathcal{R}$, which in fact form an ideal containing $\mathfrak{b}_j$. Call this ideal $\mathcal{A}(\mathfrak{v}_j)$, the *annihilator* of $\mathfrak{v}_j$.

A set of relations for $\mathcal{W} \simeq \dfrac{\mathcal{U}}{\psi(\mathcal{V})}$ is obtained from any set of relations for $\mathcal{U}$ by adjoining a relation equivalent to the element $\psi(1)$ in $\mathcal{U}$. From the standard diagonal presentation matrix for $\mathcal{U}$, we may now form a presentation matrix of the module $\mathcal{W} = \dfrac{\mathcal{U}}{\psi(\mathcal{V})}$:

$$\mathcal{F} = \begin{matrix} y_1 & 0 & & 0 \\ 0 & y_2 & & \\ & & \ddots & \\ 0 & 0 & & y_k \\ v_1 & v_2 & & v_k \end{matrix},$$

where now the $\{v_i\}$ are representatives of the cosets of the ideals $\mathfrak{b}_i$ obtained from lifting to $\mathcal{R}$ the image $\pi_i \circ \psi(1) = \mathfrak{v}_i$. At this point we notice that we have proven the first conclusion of Lemma D. The presentation matrix $\mathcal{F}$ for the torsion module $\mathcal{W}$ has k columns, and hence any diagonal presentation matrix derived from $\mathcal{F}$ by allowable operations will be square of size k or less. We take steps to put the presentation $\mathcal{F}$ into a diagonal form by means of allowable operations, and succeed at least to the extent of being able to calculate $\Omega(\mathcal{W})$.

Now let (A_j) be the annihilator of v_j in the above sense. If $\mathfrak{v}_j$ were a generator of $\dfrac{\mathcal{R}}{\mathfrak{b}_j}$, the annihilator would equal only $\mathfrak{b}_j$, which is generated by y_j. But the ideal in $\dfrac{\mathcal{R}}{\mathfrak{b}_j}$ generated by v_j may not be this whole quotient ring; it is an ideal annihilated by an ideal A_j strictly containing $\mathfrak{b}_j$. Hence in this case the generator A_j strictly divides y_j.

The lift to $\mathcal{R}$ of the ideal $\mathfrak{v}_j$ can be called Δ_j, which is mapped to $\mathfrak{b}_j$ by multiplication with A_j, which generates the annihilating ideal of $\mathfrak{v}_j$. In fact we have constructively that the ideal Δ_j is generated by an element of the same name which is $\gcd(y_j, v_j)$. Thus we can take $y_j = A_j\Delta_j$ and $v_j = \kappa_j\Delta_j$ where A_j and κ_j are relatively prime. In other words, A_j is the extra multiplier needed to adjoin to v_j in order to put it into $\mathfrak{b}_j$. This discussion now yields a new way of looking at the presentation, namely

$$\mathcal{F} = \begin{matrix} A_1\Delta_1 & 0 & & 0 \\ 0 & A_2\Delta_2 & & \\ & & \ddots & \\ 0 & 0 & & A_k\Delta_k \\ \kappa_1\Delta_1 & \kappa_2\Delta_2 & & \kappa_k\Delta_k \end{matrix},$$

where $\gcd(\kappa_j, \Delta_j) = 1$.

Henceforth we sometimes write for the gcd of a set of elements of $\mathcal{R}$, $\Gamma = (a_1, \ldots, a_s)$ and for the least common multiple (lcm), the expression $\Lambda = [b_1, \ldots, b_t]$. Some justification for this notation is given in the Appendix. Returning to our discussion of the presentation $\mathcal{F}$ we note two facts. Since ψ is a well-defined module

homomorphism we must have that $\mathfrak{g}$, the annihilator for $\mathcal{V}$, must be contained in the annihilator for each component $\mathfrak{v}_i$. Therefore $A_i \mid \mathfrak{g}$ for each $i = 1, \ldots, k$ and hence $\Lambda = \operatorname{lcm}_{i=1}^{k} A_i = [A_1, \ldots, A_k]$ is such that $\Lambda \mid \mathfrak{g}$, or $\mathfrak{g} \subset (\Lambda)$.

On the other hand, the mapping ψ is injective, and we notice that since $A_j v_j = 0 \pmod{\mathfrak{b}_j}$ for each summand, we also have $\Lambda\psi(1) = 0$ in $\mathcal{U}$. By injectivity $\Lambda = 0$ in $\mathcal{V} \simeq \frac{\mathcal{R}}{\mathfrak{g}}$, hence $\mathfrak{g} \mid \Lambda$. Putting the two results together yields, as ideals of $\mathcal{R}$,

$$(\Lambda) = [A_1, \ldots, A_k] = \mathfrak{g}.$$

Thus if we can show that $\Lambda\Omega(\mathcal{W}) = \Omega(\mathcal{U}) = \prod_{i=1}^{k} y_i$ we will have completed the proof of Lemma D. In modifying the presentation $\mathcal{F}$, we arrive at a diagonal presentation similar to that of $\mathcal{U}$ but with certain factors deleted from the "main diagonal". The product of these deleted ring elements generates the "dividing ideal". Step by step with prove that this dividing ideal is just Λ, which finishes the proof. We are now ready to begin a reduction of the presentation $\mathcal{F}$ by means of allowable operations.

Focus on the first columns of $\mathcal{F}$,

$$\begin{array}{lll} A_1\Delta_1 & & \\ & A_2\Delta_2 & \\ \vdots & & \\ \kappa_1\Delta_1 & \kappa_2\Delta_2 & \ldots \end{array},$$

We perform allowable row operations affecting only the first and final ($k+1$-st) rows. Since $\gcd(A_1, \kappa_1) = 1$, by applying the Euclidean algorithm we may find such operations that result in $1 \cdot \Delta_1$ in the $(1,1)$ matrix position, and 0 in the $(k+1, 1)$ position. The resulting matrix portion is

$$\mathcal{F} = \begin{array}{lll} \hat{A}_1\Delta_1 & * & \\ & A_2\Delta_2 & \\ \vdots & & \\ 0 & A_1\kappa_2\Delta_2 & \ldots \end{array},$$

By way of explanation, we write $\hat{A}_1$ in the first row since this factor was eliminated by the operations, but we still want to be reminded of it. The factor A_1 will clearly be a factor in the dividing ideal. The entry $*$ arises through (repeated) adding of a multiple of the $k+1$-st row to the first row. The factor A_1 now occurs in the $(k+1, 2)$ entry for the following reason. The 2×2 submatrix of $\mathcal{F}$ consisting of the first and $k+1$ column, and the first and second rows, is acted upon by the same sequence of allowable row operations. The result is the corresponding submatrix of the modified $\mathcal{F}$. But the determinant of the original 2×2 submatrix equals $A_1\kappa_2\Delta_1\Delta_2$. This determinant is invariant under the sequence of operations, so a factor of A_1 (up to a unit) must appear in the modified matrix at the $(k+1, 2)$ entry.

The sequence of operations applied to $\mathcal{F}$ so far has modified every column of the matrix. New elements represented by $*$ have been adjoined but only to the first row (above the main diagonal). The first entry in the $k+1$-st row has been zeroed out. We now put the remaining submatrix of interest, of rows two and higher, and

columns two and higher, into a form similar to the original $\mathcal{F}$. We illustrate this on column two. The entries now in position $(2,2)$ and $(k+1,2)$ are $A_2\Delta_2$ and $v_2' = A_1\kappa_2\Delta_2$ respectively. We need to find the annihilator of v_2'. This is done exactly as before by using $\rho = \gcd(A_1, A_2)$. Letting $\Delta_2' = \rho\Delta_2$, $\kappa_2' = A_1 \cdot \kappa_2/\rho$, we may now write the transformed $\mathcal{F}$ as

$$\begin{matrix} \hat{A}_1\Delta_1 & * & \\ & \frac{A_2}{(A_1,A_2)}\Delta_2' & \\ \vdots & & \\ 0 & A_1\kappa_2'\Delta_2' & \dots \end{matrix},$$

This is precisely the form that we want, with the new annihilator in the second column equal to $A_2' = A_2/\rho$. We are able the continue this whole process on to the second column since in particular, we have $\gcd(A_2', \kappa_2') = 1$. That is, we may perform the Euclidean algorithm on A_2' and κ_2' through row operations involving the second row and final row. Having said that, it is clear the A_1 is a factor of the dividing ideal, and that the "upper triangular" elements $*$ that arise in no way affect the calculation of the determinant of the diagonal presentation matrix which will ultimately arise from $\mathcal{F}$.

After the above modifications "due to" the annihilating element, our presentation matrix for M has the appearance

$$\mathcal{F}' = \begin{matrix} \hat{A}_1\Delta_1 & * & * \\ 0 & A_2'\Delta_2' & \\ & & \ddots \\ 0 & 0 & A_k'\Delta_k' \\ 0 & \kappa_2'\Delta_2' & \kappa_k'\Delta_k' \end{matrix},$$

Here again, $A_k' = \frac{A_k}{\gcd(A_1, A_k)}$. The truncated matrix $\mathcal{K}$, consisting of $\mathcal{F}'$ with first row and column removed, is now identical in form with $\mathcal{F}$. The same reduction as just performed on $\mathcal{F}$ may now be done on $\mathcal{K}$. Thus an induction hypothesis can be applied to our contention that the "the dividing ideal" (for the case $\mathcal{F}$ the same as $\mathfrak{g}$ as we saw) can be computed as $\Lambda' = [A_2', \dots, A_k']$. Clearly the dividing ideal of $\mathcal{F}$ is A_1 times the dividing ideal of $\mathcal{K}$. To dispose of the case $k = 1$, this follows from the fact that already $\mathcal{F}'$ is diagonal in this case (plus a final "row" of zeroes), so the factor with the hat, $\hat{A}_1 = \hat{A}$ must generate $\mathfrak{g}$, and also generates the dividing ideal. This last is because the Fitting generator is calculated to be $A_1\Delta_1$ before the Euclidean operations, and Δ_1 "after". To carry through the induction step from size $k-1$ to size k (number of columns) of the presentation matrices $\mathcal{F}$ and $\mathcal{K}$ it is now only necessary to prove

$$(\lambda) \qquad \operatorname{lcm}[A_1, \dots, A_k] = A_1 \cdot \operatorname{lcm}[\frac{A_2}{\gcd(A_1, A_2)}, \dots, \frac{A_k}{\gcd(A_1, A_k)}].$$

That this formula in fact holds in a Euclidean ring such as $\mathcal{R}$ is the content of Proposition 9 and is proved in the Appendix.

Finale for Theorem M.

Once the invariance (to module type) of $\Omega(M) = \sigma_n(M)$ was established, we asserted that to prove that "invariant factors" or the Smith form is invariant (and a complete invariant), it was enough to repeat the proof of Proposition 7 with some minor modifications.

Proof of Theorem M. This we now do by referring to that proof without depicting too many formulas. If the Smith forms $\mathcal{S}$ and $\mathcal{T}$ look identical up to units,the modules H and K which they present must be isomorphic, because the direct sums indicated by $\mathcal{S}$ and $\mathcal{T}$. Now suppose conversely that $H \simeq K$ as $\mathcal{R}$-modules, then by Claim 1, $\Omega(H) = \Omega(K)$. If matrices $\mathcal{S}$ and $\mathcal{T}$ both have order 1, then the sole entry in both cases equals $\Omega(H) = \Omega(K)$ up to association. Take as an induction hypothesis that $\mathcal{S} = \mathcal{T}$ whenever $\hbar(\Omega)$ has a smaller value, where $\hbar$ is the Euclidean height in $\mathcal{R}$. Similar to the proof of Proposition 7 we obtain

$$\Omega(H) = \prod_{i=1}^{s} a_i = \prod_{j=1}^{t} b_j = \Omega(K),$$

where a_i, b_j are the diagonal entries of $\mathcal{S}$, $\mathcal{T}$ respectively. As in the previous case (Proposition 7) we first rule out that $a = a_1$ and $b = b_1$ are relatively prime.

Fact 6 still holds for modules over $\mathcal{R}$ as well as for subgroups of an abelian group, and Fact 7 is also valid since it is essentially the case $k = 1$ of Lemma D. In other words, by forming the module aH and presenting it in two different ways we observe from Claim 1 that $\Omega(aH)$ does not change. The entries of the induced diagonal presentation are computed just as in Proposition 7. We get a sequence of equalities, and an inequality similar to (X), but where we apply the Euclidean height to everything. Here we indispensably utilize the supposition that Euclidean height preserves multiplication in $\mathcal{R}$, rather than allowing the vaguer characterization of "height" as for example given in Herstein, [2]. As before the singular inequality $\hbar(k_1) < \hbar(a)$, since k_1 is a unit, implies that $\hbar(d_1) > \hbar(\frac{b_1}{a})$ which leads to a strict inequality

$$\hbar(\Omega(H)) \cdot t\hbar(a) > \hbar(\Omega(K)) \cdot s\hbar(a).$$

This in turn implies that $t > s$, but working from b instead of a gave $x > t$ so we conclude that a and b have a common factor g. Consideration of the isomorphic modules gH and gK yields an equation like (Y), with N replaced by Ω and the height $\hbar$ applied to each quantity. This gives $s\hbar(g) = t\hbar(g)$, and $\hbar(g) \in \mathbb{Z}$, $\hbar(g) > 0$ show that the size of the matrices $\mathcal{S}$ and $\mathcal{T}$ must be equal.

Now given this new information, suppose that b is not a multiple of a. Then specific presentations of aH and aK are available for which Ω can be constructed, and both of those values must be equal by Claim 1. Certainly then $\hbar(\Omega(aH)) = \hbar(\Omega(aK))$. But an explicit calculation shows that this is not true, examining the modification of equation (Z), since here again $\hbar(k_1) < \hbar(a)$. Repeating this argument for $bH \simeq bK$ we arrive at $a = b$ up to units. Now the Smith form matrices for aH, bK respectively, arising from $\mathcal{S}$, $\mathcal{T}$ have Ω values strictly smaller than for $\mathcal{S}$ itself, so that the induction hypothesis may be applied. That is the constructed presentations $\mathcal{S}' = \mathcal{T}'$ are identical up to association, entry by entry. But then also $\mathcal{S} = \mathcal{T}$ similarly to the abelian group case. This completes the proof of Theorem M.

Groups of Order p^2

We complete the study of groups whose order has a small number of factors. Our goal here is to show that the non-abelian case cannot arise, unlike when $kp+1=q$ above. Thus we show that

Proposition 8. *If the order of G is the square of a prime p, then G is abelian. Thus the invariants of G are either (p^2) or (p,p).*

The proof will occupy the rest of the section. If there is $a \in G$ of order p^2, we are done since G will be cyclic, abelian, with invariant (p^2). Hence we may assume that $a \neq e$ implies that $a^p = e$. The center $C(G)$ is a subgroup consisting of all elements that commute with every element of G. Suppose the center is non-trivial, then choose $c \in C(G)$, $c \neq e$. There is an $x \in G$ not a power of c. The centralizer $C(x)$ is also a subgroup consisting of those elements of G that commute with x. But $C(x)$ includes both x and c, and the subgroup generated by them which equals G. Thus x is in $C(G)$ after all. Since all elements can be written as $c^i x^j$, this shows that G is abelian. This is where the invariants are (p,p).

We remain only with the case where G has center $C(G) = \{e\}$. Choose a, r so that r is not a power of a. As before, as a set, $G = \{a^i r^j\}$ where $0 \leq i, j \leq p-1$. Now for $j = 0, \ldots, p-1$, the elements $\{a^j r a^{-j}\}$ are distinct, since otherwise we would have for some r not congruent to 0 mod (p), $a^k r = r a^k$, we again get that the centralizer of r is the whole group so r is in the center which is a contradiction. Thus we have constructed at least p distinct conjugates grg^{-1} of r. On the other hand, if $y = xrx^{-1}$, we have for some k, l, $x = a^k r^l$, so that $y = a^k r a^{-k}$ after all. Thus the conjugacy class of any non-unity element of G has exactly p elements. The conjugacy classes are disjoint and exhaust G, and the identity element e has one element in its conjugacy class. Hence $|G| = mp + 1$, where m is the number of conjugacy classes of non-unity elements. But since $|G| = p^2$, this is impossible, so there must be a non-trivial center. It may be worthwhile to see how all the subgroups of G of order p must be related by conjugacy. This is of course an important part of Sylow theory. But the argument from conjugacy of elements is simple, especially since in the p^2 case, those subgroups do not exist anyway.

Appendix: Formulas in a Euclidean Ring

We list with proof some useful formulas involving gcd, lcm and so forth, culminating in Proposition 9 which was key to establishing Theorem M on the invariance of the Smith form. The proofs are carried out with the same philosophy we adopted for the algorithms and theorems given above. This philosophy is guided by a principle of strict construction of entities used in the proof. Given a pair of elements of the Euclidean domain $\mathcal{R}$, we have the Euclidean algorithm for their gcd which is the epitome of such a construction. It is carried out in far fewer steps than a number linear in the sum of the $\log_2$ of the heights of the elements. There is no repeated branching, probabilistic step or arbitrary choice involved in carrying out the algorithm. One may wonder why use of the Fundamental Theorem of Arithmetic (FTA) is not allowed for our purposes. After all, this theorem is derived by means of the Euclidean algorithm.

One reason for not allowing consideration of primes or primary factors to establish these formulas that involve elements, is that we avoided reliance on prime ideals in our construction of the main invariants of a torsion module. To have

the underlying facts about the elements of $\mathcal{R}$ established by means of the theory of primes would show an inconsistency. Recall that there were advantages not to digress to dealing with "elementary (primary) divisors". In talking about a matrix M, all matrices in subsequent discussion about invariants were no larger than M. Now consider the FTA. It is highly non-constructive in that even over the integers there is no known way to find the factorization much better than by searching over a finite list, and the computational cost is at least order of a high degree polynomial, in the log of the combined heights. If $\mathcal{R}$ is a polynomial ring over an infinite field, this factorization is even harder. That is not to say that in mathematics one should avoid using irreducible constructions at all cost. But the advantages of constructibility and computability remain if we are able to use the axioms of the system in a more basic way. It is also important just to know "if we can do it". Those working in the foundations of different areas of mathematics are certainly motivated to seek proofs of conceptual simplicity, which may sometimes be longer. One reason is certainly to be able to present proofs of sufficient rigor to be affirmed by an automated proof-checker. In the history of mathematics, the "elementary" proof, one that did not incorporate too many advanced results from outside the field of the theorem itself, was sought after. The Prime Number Theorem itself is the best known example of this.

It was by no means clear to the author *ab initio* that all of these essential identities could be established from first principles, without managing a list of prime factors, already provided. Most textbooks for example define the lcm in terms of the prime factorizations of the two elements (numbers) in question. In particular Lemma T, **"Turbo-Euclid"**, which serves to some extent as a substitute for prime factorization, is not completely trivial.

The question remains whether this "prime-free" point of view has any mathematical content. For example, one could imagine a domain that satisfies most of the properties of a true Euclidean ring, but where some elements are not expressible as a finite product of irreducibles. For example, we could allow the Euclidean height values to constitute an ordered set, but not a finite one. We would add as an axiom that for any pair of elements, the Euclidean algorithm terminates in the gcd. Also we would have to postulate that the algorithm in Lemma T, "Turbo-Euclid" also terminates in a finite number of steps. If such a ring exists, it would still imbue its torsion modules with the well-known "direct sum of cyclic" characterization specified in Theorem M. Please see some further discussion of this point in the Postscriptum.

In what follows all Roman letters are elements of $\mathcal{R}^+$. We should not have need for expressions such as $\gcd(0,a)$ though these are conventionally given the value a. Also lcm is best defined on pairs of elements of $\mathcal{R}^+$. Those elements written in capital are suggestive of generators of an ideal, but otherwise there is no distinction.

Fact 8. $((a,b),c) = (a,(b,c))$.

Proof. There are plenty of exceptional "ground" cases, such as where one of the three elements is a unit, where the assertion is obvious. Assume that it holds whenever the sum of the heights of the three elements is lower than that given. Without loss of generality, let a have the greatest height. Then by Euclid there is $r \in \mathcal{R}$ such that $a - rb$ has height less than that of b. By the induction hypothesis we have

$$((a,b),c) = ((a-rb,b),c) = (a-rb,(b,c)) = (a,(b,c))$$

since (b,c) divides b so divides rb. Alternatively the assertion follows from pure propositional logic, since a divisor of a and (b and c), is a divisor of (a and b) and c. By similar construction, all expressions of k given variables, with parentheses consistently placed, are equal to each other, and represent a divisor of all the variables, which is a multiple of any other such divisor. Thus we may write

$$(a_1, \ldots, a_k),$$

which represents (uniquely up to association) an element $x \in \mathcal{R}$ minimal in Euclidean height $\hbar(x)$, where x ranges over all possible non-zero expressions $\alpha_1 a_1 + \cdots + \alpha_k a_k$.

Fact 9. *$x \mid yz$ implies that for some x_1, x_2, $x = x_1 x_2$, $x_1 \mid y$, and $x_2 \mid z$.*

Proof. Let $\gamma = (x,y)$, $x' = \frac{x}{\gamma}$, $y' = \frac{y}{\gamma}$. Then $x' \mid y'z$ and $(x', y') = 1$. Thus for some ring elements k, ℓ

$$kx' + \ell y' = 1$$
$$kx'z + \ell y'z = z,$$

so $x' \mid z$. Thus setting $x_1 = \gamma$, $x_2 = x'$ fulfills the required conclusion.

Fact 10. *If $x \mid z$ and $y \mid z$ then $[x,y] = \dfrac{xy}{(x,y)} \mid z$.*

Proof. For some k, ℓ, $(x,y) = k \cdot x + \ell \cdot y$. Thus we wish to show that

$$x \cdot y \mid z \cdot (kx + \ell y).$$

But xy divides both $z \cdot x$ and $z \cdot y$, hence divides the RHS.

Fact 11. $(a, b \cdot c) \mid (a,b)(a,c)$.

Proof. Let $\Delta = (a, bc)$ so $\Delta \mid b \cdot c$ and $\Delta \mid a$. By Fact 9 choose Δ_1, Δ_2 so that $\Delta = \Delta_1 \Delta_2$, $\Delta_1 \mid b$, $\Delta_2 \mid c$. Each Δ_i also divides a, so $\Delta_1 \mid (a,b)$ and $\Delta_2 \mid (a,c)$ which proves the Fact.

Fact 12. $(b,c) = 1$ *implies that* $(a, bc) = (a,b)(a,c)$.

Proof. By Fact 11, we only need to prove that $(a,b)(a,c) \mid (a, bc)$. Certainly the LHS divides the factor $b \cdot c$ so we need to show that it divides a. Both factors of the LHS (a,b) and (a,c) divide a so we have by Fact 10

$$(\delta) \qquad \frac{(a,b)(a,c)}{((a,b),(a,c))} \mid a.$$

But $((a,b),(a,c)) = (((a,b),a),c)$ by Fact 1 , and similarly

$$(((a,b),a),c) = ((a,b),c) = (a,(b,c)) = 1$$

by hypothesis. Thus from (δ), $(a,b)(a,c)$ divides (a, bc).

Fact 13.

(i) $(az, bz) = z(a, b)$.
(ii) *If* $(y, z) = 1$, *then* $(xy, z) = (x, z)$.

Proof. For (i) we note that both the LHS and RHS are equal to an expression $k \cdot az + \ell \cdot bz$ which is minimal for height among such expressions. Thus they must be equal up to association. For (ii) we have by Fact 11 that $(xy, z) \mid (x, z)(y, z) = (x, z)$. But in general $(x, z) \mid (xy, z)$ since a divisor k of x and z is also a divisor of xy and z.

Fact 14. *We have*

$$\operatorname{lcm}[A_1, \ldots, A_n] = [A_1 [\ldots [A_{n-1}, A_n] \ldots] = [\ldots [A_1, A_2] \ldots] A_n].$$

Proof. The case $n = 1$ is an easy consequence of the fact that $a \mid b \mid a$ implies that a and b are associate. We establish the first equality, the second following from the first by looking at it in a mirror. In fact up to association, all bracket products of $A_1, \ldots, A_n$ in any consistent combination are equal. By an induction hypothesis we have to prove that $[A_1, \ldots, A_n] = [A_1, [A_2, \ldots, A_n]] = [A_1, Y] = \dfrac{AY}{(A, Y)}$. Since by definition of lcm, $A_2, \ldots, A_n$ divide Y, thus $A_1, \ldots, A_n$ all divide $Z = [A_1, Y]$. Suppose that $A_1, \ldots, A_n$ all divide W, then by the induction hypothesis $Y \mid W$ and by Fact 10 $[A_1, Y] \mid W$. Thus $Z = [A_1, Y] = [A_1 [\ldots [A_{n-1}, A_n] \ldots]$ fulfills the characterization of lcm and must equal $\operatorname{lcm}[A_1, \ldots, A_n]$ up to multiplicative unit.

Lemma T "Turbo-Euclid". *Let* $b, c \in \mathcal{R}$ *and* $\Delta = (b, c)$ *the* gcd. *Then we may write* $\Delta = \Delta_1 \cdot \Delta_2$, *where* $b = b'\Delta_1$, $(b', \Delta_1) = 1$ *and* $c = c'\Delta_2$, $(c', \Delta_2) = 1$. *Furthermore* $(\Delta_1, \Delta_2) = 1$ *and* $(b', c') = 1$.

Proof. Let the total height of the word $b \cdot c$ be $\theta = \hbar(b \cdot c) = \hbar(b) \cdot \hbar(c)$. If every element in the hypothesis is a unit, the result we wish for is easy, so we may proceed inductively, assuming that the result holds for pairs b, c with a lesser value of θ. The induction is well-founded since θ takes positive integer values. Now set $b_1 = b$, $c_1 = c$, $\beta = \frac{b_1}{\Delta}$, $\gamma = \frac{c_1}{\Delta}$. Then $\gcd(\beta, \gamma) = 1$ since a common factor would lead to Δ strictly dividing a common factor of b_1 and c_1. Then let $x = \gcd(\beta, \Delta)$, $y = \gcd(\gamma, \Delta)$ which are also relatively prime and each divide both b_1 and c_1, which implies that $xy \mid b_1$ and $xy \mid c_1$. Therefore we may set $b_2 = \frac{b_1}{xy}$, $c_2 = \frac{c_1}{xy}$. The same steps can now be applied to b_2 and c_2.

This algorithm will continue to diminish the total height of $b_i \cdot c_i$ as $i = 1, \ldots$ increases, until both x and y are equal to 1 (equal to a unit). In fact we may stop the algorithm when either $x = 1$ or $y = 1$; consider without loss of generality the case $x = 1$. Let $\Delta_1 = \Delta$ and $\Delta_2 = 1$. Then $b' = \beta$, and $(\beta, \Delta) = 1$ implies $(b', \Delta_1) = 1$, where the other conclusions in this case are trivially satisfied.

Thus we may assume by induction hypothesis that $\Gamma = \gcd(b_2, c_2)$ can be split in the prescribed manner into $\Gamma = \Gamma_1 \cdot \Gamma_2$. By hypothesis $(\Gamma_1, \Gamma_2) = 1$ and $(\frac{b_2}{\Gamma_1}, \Gamma_1) = 1$. The latter is equivalent to $(\frac{b_2}{\Gamma}, \Gamma_1) = 1$ by Fact 13(ii), since $\Gamma = \Gamma_1 \cdot \Gamma_2$ and the factors are relatively prime.

Now

$$\Gamma = (\frac{b_1}{xy}, \frac{c_1}{xy}) = \frac{1}{xy} \cdot (b_1, c_1) = \frac{1}{xy} \Delta,$$

using Fact 13(i). Furthermore $\beta_2 = b_2/\Gamma = \left(\frac{b_1}{xy}\right) \div \left(\frac{\Delta}{xy}\right) = \frac{b_1}{\Delta} = \beta$. Since Γ_1 is relatively prime to β_2, as followed immediately by Fact 13(ii) from the hypothesis, we have $(\beta, \Gamma_1) = 1$. *A fortiori* $(x, \Gamma_1) = 1$ since x is a divisor of β, and by the symmetrical argument, we have $(\gamma, \Gamma_2) = 1$ and $(y, \Gamma_2) = 1$.

Let $\Delta_1 = \Gamma_1 \cdot y$ and $\Delta_2 = \Gamma_2 \cdot x$. In view of the "relatively prime" relations established, it follows immediately that $(\Delta_1, \Delta_2) = 1$. Furthermore,

$$b = \beta \cdot \Delta = \beta\, xy\Gamma = (\beta y \Gamma_2)(x\Gamma_1)$$

equals $\beta\Delta_2\Delta_1$. We see two factors in parentheses, each consisting of two or three subfactors. We showed that each subfactor is relatively prime to each subfactor in the other factor. For example, β is relatively prime to x and to Γ_1. Thus Δ_1 is relatively prime to $b' = \frac{b_1}{\Delta_1}$ using Fact 13(i) as necessary, thus fulfilling the conclusion. The proof of **"Turbo-Euclid"** is completed by noting the symmetrical argument to all of the above showing that Δ_2 is relatively prime to $c' = \frac{c_1}{\Delta_2}$. The fact that $(b', c') = 1$ follows from $b' = \beta\Delta_2$, $c' = \gamma\Delta_1$, $(\beta, \Delta) = (\gamma, \Delta) = 1$ and $(\Delta_1, \Delta_2) = 1$.

Fact 15. *For all* $a, b, c \in \mathcal{R}^+$,

$$(a, [b, c]) = [(a, b), (a, c)].$$

Proof. Apply Lemma T **"Turbo-Euclid"** to the elements b, c and use the same notation as in the proof above. Note that $(a, b') = (a, b)/(a, \Delta_1)$ and $(a, c') = (a, c)/(a, \Delta_2)$. Using Facts 8 and 12, we have

$$\begin{aligned}(a, [b, c]) = (a, \frac{bc}{\Delta}) = (a, b'\, c') = (a, b')(a, c') \\ = \frac{(a, b)}{(a, \Delta_1)} \cdot \frac{(a, c)}{(a, \Delta_2)} = \frac{(a, b)(a, c)}{(a, (b, c))} \\ = \frac{(a, b)(a, c)}{((a, b), (a, c))} = [(a, b), (a, c)].\end{aligned}$$

as was to be proved.

Fact 16. $(a, \mathrm{lcm}[b_1, \ldots, b_k]) = \mathrm{lcm}[(a, b_1), \ldots, (a, b_k)]$.

Proof. By Fact 14 we may rewrite the lcm in the LHS as a sequence of least common multiples of pairs of elements, nested to the right. That is, the LHS equals $(a, [b_1, [b_2, [\ldots, [b_{k-1}, b_k] \ldots])$. This expression by use of Fact 15 now equals $[(a, b_1), (a, [b_2, \ldots, b_k])]$. This process can be repeated a finite number of times on the inner gcd so eventually we obtain a nested sequence of brackets of the form $[(a, b_1), [(a, b_2), [\ldots, [(a, b_{k-1}, (a, b_k)] \ldots]$ which by Fact 14 is equal to the RHS of the conclusion.

Fact 17. $[[a, b], c] = [a, [b, c]]$.

Proof. Notice that this result is a consequence of Fact 14. However, we demonstrate it again using Fact 15. With this formula involving three variables, we easily

can recover the part of Fact 14 that indicates that all consistently "bracketed" expressions involving k variables $b_1, \ldots, b_k$, are equal. Looking at the RHS we note

$$[a,[b,c]] = \frac{a[b,c]}{(a,[b,c])} = \frac{a \cdot bc}{(a,[b,c])(b,c)}.$$

By applying Fact 15 to the denominator of the last expression we arrive at

$$\frac{abc \cdot ((a,b),(a,c))}{(a,b)(a,c)(b,c)} = \frac{abc \cdot (a,b,c)}{(a,b)(a,c)(b,c)}$$

using Fact 8. But the last expression is invariant under all permuations of a, b, c, hence we have for example $[a,[b,c]] = [c,[a,b]]$ which is the same as the desired conclusion.

Fact 18.

$$(\frac{b}{(a,b)}, \frac{c}{(a,c)}) = \frac{(b,c)}{(a,b,c)}.$$

Proof. "Ground instances" for an induction, where one of the three variables is a unit, are very easy to establish. Otherwise, we wish to decompose either b or c according to Lemma T. For instance we would like $b = \Delta_1 b'$ to be a non-trivial splitting, using the notation of Lemma T. If $\Delta_1 = b$, the splitting would be trivial, but then $\Delta = b$ and $b \mid c$. On the other hand, Δ_1 could equal 1, when we look at c. Then possibly $\Delta_2 = \Delta = 1$ else we are covered. Thus the two exceptional cases are $b \mid c$ (similarly $c \mid b$), and $(b,c) = 1$.

Consider the case $b \mid c$ or $kb = c$. Then the LHS is equal to $b/(a,b)$ invoking Fact 8 as necessary. Working on the RHS, we will be done when we can prove that

$$\frac{b}{(a,b)} \mid \frac{c}{(a,c)}.$$

But this is equivalent to $b(a,kb) \mid kb \cdot (a,b)$ which follows immediately from Fact 11.

Supposing $(b,c) = 1$, the RHS is equal to 1. Examining the LHS, its value would be a factor m common to both $b/(a,b)$ and $c/(a,c)$, but that would be a factor common to b and c, hence a unit.

Having handled special cases, we may without loss of generality take Δ_2 equal neither to 1 nor to c. First note that $(a,b,c) = (a,(b,\Delta_2 c')) = (a,(b,\Delta_2)(b,c')) = (a,(b,\Delta_2))(a,(b,c'))$ using Fact 12. Thus

$$\frac{(b,c)}{(a,b,c)} = \frac{(b,\Delta_2)(b,c')}{(a,(b,\Delta_2))(a,(b,c'))}.$$

The RHS in this expression is the product of two quotients each similar to the LHS. Induction on $\hbar(abc)$ can be applied to each quotient. Thus by induction hypothesis, $\frac{(b,c')}{(a,b,c')} = (\frac{b}{(a,b)}, \frac{c}{(a,c')})$, with a similar expression holding, where c' is replaced througout by Δ_2. Multiplying the two expressions together and applying Facts 8 and 11 as required, yields

$$(\frac{b}{(a,b)}, \frac{\Delta_2}{(a,\Delta_2)}) \cdot (\frac{b}{(a,b)}, \frac{c'}{(a,c')}) = (\frac{b}{(a,b)}, \frac{\Delta_2 c'}{(a,\Delta_2)(a,c')}) = (\frac{b}{(a,b)}, \frac{c}{(a,c)})$$

which completes the proof.

The following formula is included for completeness and is not used in the sequel.

Fact 19.

$$\left(\frac{B_1}{(A,B_1)},\ldots,\frac{B_k}{(A,B_k)}\right)=\frac{(B_1,\ldots,B_k)}{(A,B_1,\ldots,B_k)}.$$

Proof. A suitable ground case is just Fact 18. For $k>2$ we have

$$\left(\frac{B_1}{(A,B_1)},\ldots,\frac{B_k}{(A,B_k)}\right)=\left(\frac{B_1}{(A,B_1)},(\frac{B_2}{(A,B_2)},\ldots,\frac{B_k}{(A,B_k)})\right)$$

from Fact 8, which in turn equals, using an induction hypothesis for smaller k, $\left(\frac{B_1}{(A,B_1)},\frac{(B_2,\ldots,B_k)}{(A,(B_2,\ldots,B_k))}\right)$. Letting $A=A$, $B=B_1$, $C=(B_2,\ldots,B_k)$ and applying Fact 18 again yields

$$\frac{(B_1(B_2\ldots B_k))}{(A,(B_1,(B_2\ldots B_k)))},$$

which equals the desired RHS in the statement of Fact 19.

Fact 20. *We have the following Product Formula for least common multiple.*

$$[B_1,\ldots,B_k]=\left[\frac{B_1}{(A,B_1)},\ldots,\frac{B_k}{(A,B_k)}\right]\cdot[(A,B_1),\ldots,(A,B_k)].$$

Proof. The statement is true if $k=1$, we assume also that it holds for values less than our given k. Taking brackets nested to the right, multiplying top and bottom by (A,B_1), and expanding $[B_2,\ldots,B_k]$ according to the induction hypothesis yields

$$\begin{aligned}[B_1,\ldots,B_k]&=\frac{B_1\cdot[B_2,\ldots,B_k]}{(B_1,[B_2\ldots B_k])}\\&=\frac{B_1}{(A,B_1)}\cdot(A,B_1)\left[\frac{B_2}{(A,B_2)},\ldots,\frac{B_k}{(A,B_k)}\right]\cdot\frac{[(A,B_2),\ldots,(A,B_k)]}{(B_1,[B_2\ldots B_k])}.\end{aligned}$$

Using definitional formulas for the lcm yields

$$\left[\frac{B_1}{(A,B_1)},\ldots,\frac{B_k}{(A,B_k)}\right]\cdot\left(\frac{B_1}{(A,B_1)},\left[\frac{B_2}{(A,B_2)},\ldots,\frac{B_k}{(A,B_k)}\right]\right)\cdot\frac{[(A,B_1)\ldots(A,B_k)]\cdot((A,B_1),[(A,B_2)\ldots(A,B_k)])}{(B_1,[B_2\ldots B_k])}.$$

Comparing this expression with the desired RHS of the conclusion, we need only prove the following formula.

$$\text{(W)}\quad\left(\frac{B_1}{(A,B_1)},\left[\frac{B_2}{(A,B_2)},\ldots,\frac{B_k}{(A,B_k)}\right]\right)\cdot((A,B_1),[(A,B_2)\ldots(A,B_k)])=(B_1,[B_2\ldots B_k]).$$

Using Fact 16 the LHS of (W) equals

$$\left[(\frac{B_1}{(A,B_1)},\frac{B_2}{(A,B_2)}),\ldots,(\frac{B_1}{(A,B_1)},\frac{B_k}{(A,B_k)})\right]\cdot[(A,B_1,B_2),(A,B_1,B_3),\ldots,(A,B_1,B_k)],$$

and now applying Fact 18 in the first lcm factor and Fact 8 in the second factor gives

$$\left[\frac{(B_1,B_2)}{(A,(B_1,B_2))},\frac{(B_1,B_3)}{(A,(B_1,B_3))},\ldots,\frac{(B_1,B_k)}{(A,(B_1,B_k))}\right][(A,(B_1,B_2))\ldots(A,(B_1,B_k))].$$

Setting $F_1=(B_1,B_2),\ldots,F_{k-1}=(B_1,B_k)$ we see that by Fact 20 for $k-1$, this last expression equals $[F_1,\ldots,F_{k-1}]=[(B_1,B_2),(B_1,B_3),\ldots,(B_1,B_k)]$ which by Fact 16 is equal to the RHS of (W) which was all that was required to complete the proof.

Proposition 9.

$$A\cdot\left[\frac{B_1}{(A,B_1)},\ldots,\frac{B_k}{(A,B_k)}\right]=[A,B_1,\ldots,B_k].$$

Proof. By Fact 16 we obtain

$$A\cdot\left[\frac{B_1}{(A,B_1)},\ldots,\frac{B_k}{(A,B_k)}\right]=A\cdot\left[\frac{B_1}{(A,B_1)},\ldots,\frac{B_k}{(A,B_k)}\right]\frac{[(A,B_1)\ldots(A,B_k)]}{(A,[B_1\ldots B_k])}.$$

Utilizing Fact 20 and the formula $[x,y](x,y)=xy$, this yields immediately

$$\frac{A\cdot[B_1,\ldots,B_k]}{(A,[B_1,\ldots,B_k])}=[A,B_1,\ldots,B_k],$$

which completes the proof of Proposition 9 and Theorem M.

Postscriptum: Non-Principal Rings

The question was already raised concerning the "new content" of an approach to finitely generated torsion modules over a suitable commutative domain $\mathcal{R}$. There may be non-Euclidean domains which are amenable to application of Propositions 7 and 9, as an alternative to other approaches such as restricting the underlying ring of the module $\mathcal{M}$ under consideration.

One could simply extend the definition of the height $\hbar:\mathcal{R}^+\to\mathbb{Q}^+$, with the homomorphism property on semigroups, value 1 at a unit, and the Euclid property (division with remainder). Since gcd and lcm are generally taken over finite collections of elements, it is often not difficult to prove that the Euclidean algorithm for the gcd, and the **"Turbo-Euclid"** algorithm both terminate with the required answer. Then by Algorithm I and II, and by Theorem M, we have $\mathcal{M}$ decomposable into a Smith form, which form characterizes the module type.

The first example could be $\mathcal{T}=\mathbb{F}\{X^\alpha\}$ consisting of finite sums with coefficients in a field $\mathbb{F}$, of powers of X where α ranges over powers of two, $2^n, n\in\mathbb{Z}$. The union of the principal ideals generated by these powers, is not principal, so the ring $\mathcal{T}$ is not a PID. But a pseudo-Euclidean height $\hbar(x)$ can be defined by taking 2 to a power, the largest "degree" in a sum x of rational powers of X. The required properties for "height" are satisfied, and one sees that the Euclid and Turbo-Euclid algorithms terminate with the correct answer, since a common divisor of all terms in the elements concerned exists. Thus Y can be found which is a rational power

of X, such that all the elements we work with are (integer power) polynomials in Y. Thus we revert to the proper Euclidean situation.

Let us formalize this result, applied to the slightly more sophisticated ring $\mathcal{V} = \mathbb{F}\{X^\beta\}$ where β is allowed to be any positive rational number. Defining the "height" exactly as in the previous case, we see that without invoking prime decomposition on elements or ideals (that strictly speaking does not hold in $\mathcal{V}$), Algorithms I and II, and Theorem M apply.

Thus consider the class $\mathcal{S}$ of Smith sequences (sequences of ideals of $\mathcal{V}$, or elements defined up to association, having the division property), and $\mathcal{C}$ the class of finitely generated, torsion module types over $\mathcal{R}$. Let $\Theta(s)$ associate to $(s_1, \ldots, s_k) \in \mathcal{S}$ the module type of the direct sum

$$\frac{\mathcal{R}}{(s_1)} \oplus \cdots \oplus \frac{\mathcal{R}}{(s_k)} .$$

Then the conclusion is that Θ is one-to-one and onto. Its inverse Ψ is well defined on $\mathcal{C}$ and is computed for a module M by forming a presentation P of M, and applying Algorithms I and II to obtain a Smith sequence. In other words we have

Theorem V. *Let $\mathcal{V}$ be the non-principal domain above. Two finitely generated, presented torsion modules give by Ψ above the same Smith sequence, if and only if there is a module isomorphism between them. Thus if their Smith sequences are identical up to units, the modules have the same type over $\mathcal{V}$. Conversely, two modules that may be inequivalently presented, but have the same module type, have the same Smith sequence.*

References

[1] James H. McKay, *Another Proof of Cauchy's Group Theorem*, American Math. Monthly **66** (1959), 119.
[2] I.N. Herstein, *Topics in Algebra, 2nd Ed.*, J. Wiley and Sons, New York, 1975.
[3] M.E. Keating, *A First Course in Module Theory*, Imperial College Press, London, 1998.

Ballston Station, Virginia