8000 Day 4, Group Actions, Simplicity of Icos, Sylow, Jordan Holder.

We continue to study finite groups. To study non abelian ones, we try as with abelian groups to decompose them into products composed of smaller subgroups. This is not always possible, and even to attempt it we need to prove the existence of smaller subgroups. A finite abelian group G has a subgroup of order n for every n that divides #G. This is not true for non abelian groups, but it is true for prime power factors p^s dividing #G. To find these subgroups we could look for non trivial homomorphisms, but the kernel of a homomorphism is a normal subgroup, and subgroups of non abelian groups may not be normal. Worse, some non abelian groups have no proper normal subgroups, i.e. they are "simple". A homomorphism from a simple group G is thus either injective or constant.

We cannot have a product decomposition of such a group, since a product KxH admits two projections, to K and to H, whose kernels are the normal subgroups {e}xH and Kx{e}, which intersect in the identity. It cannot be a "semi direct product" Kx_cH since that requires K to be normal and K to intersect H only in the identity, nor even be an extension of a group H by a group K, i.e. there is no exact sequence {e}-->K-->G-->H-->{e}, since that requires K to be normal.

Thus we need another tool to study general groups, "group actions", a refinement of the technique of homomorphisms.

Definition: A group G acts on a set S if there is a map GxS-->S taking (g,x) to gx, such that g(hx) = (gh)x for all g,h in G, and all x in S, and ex = x for all x in S.

This is equivalent to a homomorphism G-->Bij(S) of G into the group of bijections of S, taking g to the bijection (x-->gx) and allows us to study G by looking at how it moves points of S around.

The key concepts are **orbit**, **stabilizer**, and the counting principle
#(G) = #(orbit)#(stabilizer).

More precisely:

Defn: If G acts on S, the orbit O(y) of a point y in S is the image of the map G x {y}-->S, i.e. O(y) = {gy: all g in G}.

Defn: If y is in S, Stab(y) = {g in G: gy = y}.

Cosets and conjugacy classes come in as follows:

Lemma: If y is in S, and gy = z, then Stab(z) = gStab(y)g^{-1}, is a conjugate of Stab(y), and the set {elements in H taking y to z} = the coset (g.Stab(y)), and its conjugate hg.Stab(y)h^{-1} = {elements in H taking h(y) to h(z)}
proof: exercise:

Counting principle: For any y in S, #(G) = #O(y).#Stab(y).
proof: Since every element of G takes y to some element of the orbit of y, G is the disjoint union, over all z in O(y), of the sets {all h in H: hy = z}. Since each of these is a coset of Stab(y), and since multiplication by g is a bijection Stab(y)-->gStab(y), each of these cosets has the same cardinality as Stab(y). **QED.**

Lemma: Every subgroup H of G is a stabilizer for some action.
proof: Let G act on left cosets of H by left translation. I.e. x takes yH to (xy)H. Then H is the stabilizer of the coset eH = H. **QED.**

Thus stabilizers for actions can be used to study all subgroups.
Corollary(LaGrange): For every subgroup H of G, #(H) divides #(G).
proof: The counting principle says #(G) = #(H).#(cosets of H in G). **QED.**

Note: Being in the same orbit is an equivalence relation on S, so an action partitions S into disjoint orbits, each orbit having cardinality dividing #(G).

Def: A fixed point is a point y of S such that Stab(y) = G, i.e. O(y) = {y}.

Corollary: If S is finite, and #(G) = p^r, where p is prime, then #(S) is congruent modulo p, to the number of fixed points.
proof: S is the disjoint union of orbits, and each orbit has cardinality divisible by p, except the singleton orbits. **QED.**

Example of a simple group: G = Icos = rotation group of a regular icosahedron. G acts on the points of the icoshedron, in particular on the vertices, which form one orbit of 12 points. Since each vertex is fixed by exactly 5 rotations, #(G) = (5)(12) = 60. This agrees with the orbit of 20 faces, each fixed by 3 rotations, and the orbit of the 30 edges, each fixed by two rotations.

The 20 elements of order 3 fixing the 10 pairs of opposite faces, the 24 elements of order 5 fixing the 6 pairs of opposite vertices, and the 15 elements of order 2 fixing the 15 pairs of opposite edges, give all 59 non trivial elements of G.

Since the stabilizers of all vertices are conjugate, a normal subgroup containing one element of order 5 contains all, and similarly for the other orders. Hence a normal subgroup K of G has order = 1 + some or none of the integers 15, 20, 24. But the only divisors of 60 those sums form are 1 and 60. Hence G has no proper normal subgroups, so is simple.

Next we use actions to produce stabilizer subgroups of prime power orders.
Theorem(Sylow): Let #(G) = mp^r where p does not divide m.
1) There exist subgroups of G of order p^r.
2) All subgroups of order p^r are conjugate to one another,
3) The number of subgroups of order p^r divides m, and is congruent to 1 modulo p.
proof: Suppose G acts on a set S such that p does not divide #(S). S is a disjoint union of orbits, so there is an orbit O(x) whose order is not divisible by p. By the counting principle p^r divides #(Stab(x)). So if we can find such an action where #(Stab(x)) $\leq p^r$, we would be done.

Since G is an arbitrary group, the only thing G acts on is G itself, by translation, and conjugation. But G has order divisible by p. We might consider subgroups of G, but we do not know how many there are! So we consider subsets of G, with G acting by translation. If a subgroup H stabilizes a non empty set T, then for any y in T, translation is an injection H-->T taking g in H to gy in T. So H is no larger than T. Thus if we let G act on subsets of size p^r, then the stabilizers will have cardinality $\leq p^r$ as desired.

So we hope the number of such subsets is not divisible by p. Of course the set S of subsets of g of size p^r, has order = $\binom{mp^r}{p^r} = \frac{(mp^r)(mp^r-1)\ldots(mp^r-k)\ldots(mp^r-[p^r-1])}{(p^r)(p^r-1)\ldots(p^r-k)\ldots(p^r-[p^r-1])}$. In this fraction every factor in the top of form (mp^r-k), is divisible by p^s, $s \leq r$, if and only if k is, if and only if the factor (p^r-k) in the bottom is. Thus every factor of p occurring in the top is canceled by a factor from the bottom. Hence this binomial coefficient is not divisible by p, and thus the stabilizer of any subset in an orbit not divisible by p, gives a subgroup of G of order p^r. **QED**

Lemma: If H,K are subgroups of G and H lies in N(K), then the set of products HK is a subgroup of G, and HK/K ≈ H/(HmeetK).
proof: exercise.

To count the number of subgroups $P_1,\ldots,P_n$, of order p^r, (called Sylow p - subgroups, or p^r - subgroups) let P_1 act by conjugation on all of them. We claim P_1 fixes only P_1. To prove it, if P_1 fixes P_j, then P_1 lies in the "normalizer" $N(P_j) = \{g \text{ in } G \text{ such that } g^{-1}P_jg = P_j\}$. Then P_1P_j is a subgroup of G, and $(P_1P_j)/P_j \approx P_1/(P_1 \text{meet} P_j)$. Since the latter quotient group has order dividing $\#(P_1) = p^r$, it follows that $\#(P_1P_j)$ is a power of p. Since P_1P_j contains P_1, whose order is already the largest possible power of p for a subgroup of G, hence $P_1 = P_j$. Thus the action of P_1 on the set S of Sylow p subgroups, has exactly one fixed point. By the counting principle above for p-groups, #(S) is congruent to 1, mod p.

Now let G act on S by conjugation. The G- orbit of P_j contains the P1 orbit of P_j. Thus the G orbits are unions of P_1 orbits, and all the P_1 orbits except $\{P_1\}$, have order divisible by p. So the G orbit containing P_1 has order congruent to 1 mod p, while the others are divisible by p. But the normalizer of any P_j in G contains Pj. The order of the G orbit of P_j equals the index of that normalizer, hence divides m, so cannot be divisible by p. Thus there is only one G orbit, i.e. all P_j are conjugate. Since the order of each orbit divides m, and there is only one orbit, #(S) divides m. **QED.**

Cor: A group G of order 24 cannot be simple.
proof: $24 = 2^3.3$, so the number k of subgroups of order 8, divides 3, hence k = 1, or 3. If k = 1, the unique subgroup of order 3 is normal, if k = 3, we get a transitive action of G by conjugation on the 3 subgroups, hence a homomorphism G-->S(3), which must have a non trivial kernel, since #S(3) = 6 < 24 = #(G).

Cor: Every simple group G of order 60 is isomorphic to A(5).
proof: First we want to find a non constant homomorphism G-->S(5).
Since $60 = 2^2.3\cdot5$, by Sylow there are 1, 3, 5, or 15, sylow 4-subgroups, and 1, 4, 10, or 20, sylow 3-subgroups, and 1, 6, or 12, sylow 5-subgroups. G is simple, hence has trivial center, so cannot have a conjugacy class of one non identity element, and a transitive action on a set of n < 5 elements gives non constant homomorphism to a group S(n) of order less than 60. So there are either 5 or 15 subgroups of order 4; 10 of order 3; and 6 of order 5. This gives 20 elements of order 3, and 24 elements of order 5. So we focus on the groups of order 4.

If there are 5 of them, since G acts transitively on them by conjugation, we have our non constant map G-->S(5). If there are 15, they cannot all intersect trivially, since there are only 15 elements left in the union of all the 4-subgroups. Hence some pair of distinct 4 groups contain a common element x, necessarily of order 2.

Then the normalizer N(x) of x is a subgroup which contains two distinct sylow subgroups of order 4. Thus #(N(x)) = 4n for some n > 1, and #(N(x)) divides 60. Hence #(N(x)) = 12, 20 or 60. Hence the index of N(x), i.e. the order of the class of elements conujugate to x, has order ≤ 5. Since G acts transitively on this class, it has order 5, and again we have our non constant map π:G-->S(5).

The map π is injective since the kernel is a normal subgroup smaller than G. Moreover if S(5)-->{±1} is the "sign map" (discussed below), the composition G-->S(5)-->{±1}, must have non trivial kernel in G. Since the only non trivial normal subgroup of G is G itself, the image of the map G-->S(5) lies in A(5) = kernel(sign map). Hence $G \approx A(5)$. **QED.**

Challenge: Consider groups of order 168. Try to find a simple one, and prove it is unique. Then prove there are no other simple groups of order < 168, or even < 360, (except abelian ones of prime order).

Exercise: Extend Sylow's theorem, by showing the following:

i) If p is prime and p^s divides #G, then G has a subgroup of order p^s.
[hint: It suffices to look inside a sylow p-subgroup. Prove the center of a p-group is always non trivial by looking at conjugacy classes. I.e. elements of the center define conjugacy classes with only one element. All non trivial conjugacy classes are divisible by p. So how many singleton classes must exist? Then you can mod out by the center and use induction.]

ii) If G has a subgroup H of order p^s where p is prime, prove H is contained in a sylow p-subgroup. [hint: the proof we gave above for the number of sylow groups showed that when a p-group acts on all sylow p-subgroups by conjugation, it must be contained in any fixed subgroup.]

Exercise(i) If a group G acts non trivially on set S of n elements, (some element of G moves some element of S), and if #(G) does not divide n!, then G has a non trivial normal subgroup K, such that #(G/K) divides n!.
(ii) A group G of order 36 has a normal subgroup K such that 3|#(K). [Hint: G acts on the cosets of its Sylow subgroups by translation.]

Decomposing groups as "products" of subgroups.
Direct products:

Now that we have a good supply of subgroups in any group G, we ask when G decomposes as a product of some of these subgroups. We define a direct product of groups exactly as before:
Def. H x K = {all pairs (h,k) with h in H, k in K} and (x,y)(h,k) = (xh,yk).

Non abelian products only have half the mapping properties of abelian ones:
Lemma: The projections HxK-->H and HxK-->K are homomorphisms, and if f:G-->H and g:G-->K are any two homomorphisms, there is a unique homomorphism G-->HxK, whose

compositions G-->HxK-->H, and G-->HxK-->K equal f and g. **proof: exercise.**

This does not help us to decompose G, because if H,K are subgroups of G, we only have inclusion maps H-->G and K-->G. In the non abelian case, these do not define a map H x K-->G. This is why it is harder to decompose G as a product. The image of such a map would be the set of products of elements of H and K, but these products usually do not even define a subgroup of G unless at least one of H or K is normal.

Exercise: If H,K are subgroups of G and H lies in the normalizer of K, then HK is a subgroup of G, and HK/K ≈ H/(HmeetK).

To define a map out of a product we need some commutativity. We identify H, K with the subgroups H x {e}, and {e} x K in H x K. Then H and K intersect only in {e} = {(e_H,e_K)}, and every element of H commutes with every element of K, i.e. (h,e)(e,k) = (h,k) = (e,k)(h,e). Thus both H and K are normal subgroups of H x K. Conversely, if normal subgroups H, K of a group G intersect only in {e}, they commute with each other since for x in H, y in K, we have $x(yx^{-1}y^{-1}) = (xyx^{-1})y^{-1}$, belongs both to H and K. Hence $xy(x^{-1}y^{-1}) = e$, so xy = yx.

This much commutativity is enough to define a map out of a product.
Proposition: If f:H-->G and g:K-->G are group maps then f(H) and g(K) are subgroups of G. If the elements of these image subgroups commute with each other, i.e. if f(x)g(y) = g(y)f(x) for every x in H, y in K, then the map (f x g):H x K-->G with (f x g)(s,t) = f(s)g(t) is a homomorphism whose restrictions to H,K are f, g respectively.
proof: With this definition, (fxg)(u,v).(fxg)(s,t) = f(u)g(v)f(s)g(t) = f(u)f(s)g(v)g(t) = f(us)g(vt) = (fxg)(us,vt) = (fxg)((u,v).(s,t)). **QED.**

Cor: If H,K are normal subgroups of G and HmeetK = {e}, there is an injective homomorphism HxK-->G sending (h,k) to hk, whose image is HK.
proof: We have just proved the image groups H,K commute, so this is a homomorphism. If hk = e, then h^{-1} = k, so belongs to both H and K, hence k = e = h, proving injectivity. The image is obviously HK. **QED.**

Cor: If H,K are normal in G, HK = G and HmeetK = {e}, then G ≈ HxK.

Examples: A group of order 15 has sylow subgroups H,K of orders 3,5, which are unique, since 1 is the only factor of 5 congruent to 1 mod 3, and also the only factor of 3 congruent to 1 mod 5. Thus both H, K are normal, intersect only in {e}, so G ≈ Z/3 x Z/5 ≈ Z/(15). **QED.**

This example generalizes as follows. If #G = pq, with p,q primes, the sylow subgroups H, K have orders p, q. If p > q, the number of sylow p-subgroups divides q and has form 1, p+1, .., hence equals 1. So the sylow subgroup of the larger prime is always normal. The number of q - sylow subgroups has form nq+1 and divides p, so since p is prime it equals p, so nq = p-1, and q divides p-1. Thus we have:

Proposition: If #G = pq where p>q are primes, and q is not a factor of p -1, then G is cyclic.

proof: As above, both sylow subgroups are normal, so G ≈ Z/p x Z/q ≈ Z/(pq). **QED.**

E.g., all groups of orders, 15, 35, 65, 77, 91, 85,139, 95, 133,.... are cyclic.

What about groups of order p^2?

Proposition: All groups of order p^2 are abelian. There are only 2 of them, Z/p^2 and Z/p x Z/p.
Lemma: A p - group always has a non trivial center.
proof: This uses the orbit formula in the following form: If N(x) = {y: yx=xy} = the normalizer of x, then N(x) is a subgroup, and its index is the order of the conjugacy class of x. Hence #G = sum over one element x from each conjugacy class, of the indices of the N(x). In particular, since an element is in the center Z(G) if and only if its normalizer is G with index 1, we have:

The class equation: #G = #Z(G) + summation IndexN(x), for one x in each non trivial conjugacy class.

proof of lemma: For a p - group G, these non trivial indices are all powers of p, as is #G, hence so is #Z(G). I.e. #Z(G) is divisible by p, so the center contains more than just {e}. **QED lemma.**

proof of proposition:
If x is any element of any group, the normalizer of x always contains both x and the center Z(G). If x is in the center then N(x) = G. If not, then N(x) is strictly larger than Z(G). Since in a p group, #Z(G) is at least p, then for every x ,#N(x) is at least p^2. But that means for every x, N(x) = G. Hence every x is in Z(G). **QED.Prop.**

We now know all groups of order 4, 9, 25, 49, 121,..., and may ask about groups of order pq where p > q and q is a factor of p-1, like #G = 6, or 21, or 2p, for p odd. As above, these are the cases where only one of the two sylow subgroups need be normal. So what happens in that case? I.e. how does the "product" group HK look then? We need another tool.

Semi - direct products
If H,K are subgroups of G and only K is normal, the products kh still form a subgroup KH, but the multiplication is more complicated. If we understand H and K, we need to know how to multiply products of form (xs)(yt) where x,y are in K, s,t are in H. If s,y did commute, then (xs)(yt) would equal xyst, but sy may not commute, but the extent to which they do not commute is given by conjugation. Thus sy may not equal ys, i.e. sys^{-1} may not equal y, but it does equal $c_S(y)$ where c_S:K-->K is conjugation by s.

I.e. if we know the automorphism c_S:K-->K, then $sys^{-1} = c_S(y)$, so $sy = c_S(y)s$. Thus $xsyt = x(sy)t = x(c_S(y)s)t = (x.c_S(y))(st)$. Thus if c:H-->Aut(K) is the homomorphism taking each s to c_S = conjugation by s, the product (xs)(yt) is given by $(x.c_S(y))(s.t)$. This tells us how to define a twisted product, called the semi direct product of K and H, with twisting given by a homomorphism c:H-->Aut(K).

Defn: Let H,K be groups and let c:H-->Aut(K) be a homomorphism. Then define multiplication

on the cartesian product KxH by setting (x,s).(y,t) = (x.c_s(y), st). Denote the resulting semi direct product by Kx_cH.

Exercise: With definitions as above, prove:
(i) The semi direct product Kx_cH is a group.
(ii) The subsets K' = {(k,e) for all k in K}, and H' = {(e,h) for all h in H} are subgroups of G isomorphic to K, H respectively, and K' is normal.
(iii) The action of H on K via c becomes the conjugation action of H' on K', i.e. if k' = (k,e), h' = (e,h), then h'k'h'$^{-1}$ = (c(h)(k))' = (h(k),e).
(iv) H' is normal in Kx_cH if and only if c is the trivial homomorphism.
(v) If H, K are subgroups of a group G, K is normal, and we define c:H-->Aut(K) to be conjugation of K by H, then letting f(k,h) = kh, defines a homomorphism f:Kx_cH-->G, which is surjective if G = KH, and injective if KmeetH = {e}.

Proposition: If G has order 2p where p is an odd prime, there is exactly on non abelian group of order 2p, the dihedral group D_p.
proof: The subgroup K of order p is normal, so we have an isomorphism G ≈ (Z/p) x_c (Z/2), where c:Z/2-->Aut(Z/p) is a non trivial homomorphism. Since Aut(Z/p) ≈ (Z/p)* ≈ (Z/(p-1)), there is only one element of order 2 in Aut(Z/p), hence only one on trivial map c, hence one non abelian group. Since D_p is non abelian of order 2p, this is it. **QED.**
This classifies all groups of orders 6, 10, 14.

Next we show homomorphisms c:H-->Aut(K) that differ by an automorphism of H, define isomorphic semi direct products.
Proposition: Let H, K be groups, c:H-->Aut(K) a homomorphism, g:H-->H an automorphism of H, and define c':H-->Aut(K) by c' = cg^{-1}. Then the map f:Kx_cH-->K$x_{c'}$H defined by f(k,h) = (k,g(h)), is an isomorphism.
Proof: f is a bijective function, with inverse f^{-1}(k,h) = (k,g^{-1}(h)), so we check the homomorphism property. If (k,h), (k_1,h_1) are in Kx_cH, their product is (k,h)\(k_1,h_1) = (k.c(h)(k_1),hh_1), whose image is f(k.c(h)(k_1), hh_1) = (k.c(h)(k_1), g(hh_1)).
On the other hand the two images of (k,h) and (k_1,h_1) are f(k,h) = (k,g(h)) and f(k_1,h_1) = (k_1, g(h_1)), hence the product of the images is (k,g(h)).(k_1, g(h_1)) = (kc'(g(h))(k_1), g(h)g(h_1)). Since c'g = c, and g is a homomorphism, thus indeed f((k,h).(k_1,h_1)) = (k.c(h)(k_1), g(hh_1)) = (kc'(g(h))(k_1), g(h)g(h_1)) = f(k,h).f(k_1,h_1).
QED.

Exercise: i) If p-1 = mq, there are exactly q-1 non constant maps c:(Z/q)-->Z/(p-1), taking [1] to some multiple of [m].
ii) Aut(Z/p) ≈ Z/(p-1).
iii) If p-1 = mq, all non constant maps c:Z/q-->Aut(Z/p) define isomorphic semi direct products (Z/p) x_c (Z/q).
iv) If p-1 = mq, there is exactly one non abelian group of order pq.

Classifying groups whose order has more than 2 factors is more work.
Theorem: There are exactly 2 non abelian groups of order 8, up to isomorphism, Hamilton's unit quaternions, and D_4 = Isom(square).

Proof: #G = 8 = 2^3, and G not cyclic, so all elements have order 1,2, or 4.

Lemma: Two elements of order 2 in a group commute if and only if their product has order 2.
proof: If x,y, and xy have order 2, then (xy)(xy) = e, so xy = $(xy)^{-1} = y^{-1}x^{-1}$ = yx, since x,y have order 2. The other direction is even easier. **QED.**

Hence G has elements of all orders 1,2, and 4.
case 1) Assume there is only one element of order 2, hence 6 elements of order 4. Then let x be an element of order 4, and y another element of order 4, with y different from both x and x^{-1}. The subgroup <x> has index 2, hence is normal. Since G = <x>.<y>, and <x> ≈ <y> ≈ Z/4, G must be the image of a surjective map from a non trivial semidirect product Z/4 x_c Z/4, defined by a non constant homomorphism c:Z/4-->Aut(Z/4) ≈ Z/2. There is only one such map, hence only one such non trivial s.d.p. Z/4 x_c Z/4.

Now for the map Z/4 x_c Z/4-->G. It is multiplication, (or exponentiation in our notation) hence maps {0}xZ/4--><y> isomorphically ([0,n]--> y^n), and maps Z/4 x {0}--><x> isomorphically ([n,0]--> x^n). Since there is only one element of order 2 in G, the elements x^2 = y^2 are the same, so the element [2,2] of Z/4 x_c Z/4, must be the unique non trivial element of the kernel. Hence G ≈ [Z/4 x_c Z/4]/{(2,2)}, is also uniquely determined. So there is only one non abelian group of order 8 with a unique element of order 2. Note that Hamilton's quaternions do satisfy this description, hence this is the quaternion group.

case 2) Assume there are more than one element of order 2. There are still some elements of order 4, so let x have order 4, hence x^2 is the unique element of order 2 in the subgroup <x>. then choose another element of order 2, say y, different from x^2. Then <x> is normal and the subgroup <x>.<y> = G, so G ≈ <x> x_c <y> ≈ (Z/4)x_c(Z/2), defined by the unique non trivial map c:Z/2-->Aut(Z/4). So there is only one non abelian group of order 8 with more than one element of order 2, which must be D_4 = Isom(square).

Theorem: There are 3 non abelian groups of order 12, up to isomorphism.
proof: #G = 12 = 2^2.3, so there are sylow subgroups H,K of orders 3,4. If there are 4 subgroups of order 3, hence 8 elements of order 3, there are only 4 elements left to form one group of order 4, so the sylow 4-subgroup is unique and normal. Hence at least one of the sylow subgroups is normal. If both sylow subgroups H,K are normal, G ≈ HxK, hence G is abelian. So if G is non abelian, only one sylow subgroup is normal.

Since HK = G, we have in all cases an isomorphic map Kx_cH-->G where c:H-->Aut(K) is a non constant homomorphism. (The constant homomorphism defines the trivial direct product, which is abelian.) If the 4-subgroup is normal, we have c:Z/3-->Aut(K), where K is either Z/4 or Z/2 x Z/2. Since the only homomorphism Z/3-->Aut(Z/4) ≈ Z/2 is constant, K must be Z/2 x Z/2. Then Aut(Z/2 x Z/2) ≈ S(3) has 2 elements of order 3 so there are two non constant maps c:(Z/3)-->Aut(K). Since one can show that Aut(Z/3) acts on the set of the resulting semi direct

products by isomorphisms, and since Aut(Z/3) ≈ Z/2, the two non cobnstant maps Z/3-->S(3) yield isomorphic groups Kx_cH.

Thus there is only one non abelian group G ≈ (Z/2 x Z/2) x_c (Z/3) of order 12, with normal sylow 4 - group. In fact the group Tet = Isom(tetrahedron) has order 12, and 4 distinct subgroups of order 3, so must be this group. The action on the 4 vertices also embeds this group as A(4) in S(4), since that sub group is generated in S(4) by the 8 elements of order 3.

If K = Z/3 is the normal subgroup, and H is the sylow 4-subgroup, we have a map H-->Aut(K) ≈ Aut(Z/3) = {±Id} ≈ Z/2. If H ≈ Z/4 there is only one non trivial map, taking [1] to -Id. So there is only one non abelian group of order 12 with Z/3 as normal subgroup, and having a subgroup isomorphic to Z/4, i.e. one non trivial semi direct product (Z/3) x_c (Z/4).
I have not run across this group in geometry.

If K = Z/3 is the normal subgroup, and H ≈ Z/2 x Z/2 is the sylow 4-subgroup, then c:(Z/2 x Z/2)-->Aut(Z/3) = $(Z/3)^*$ ≈ Z/2, so there are three non constant maps, each taking two of the vectors (1,0), (0,1), (1,1) to 1, and taking the other vector to 0. But again Aut(Z/2 x Z/2) ≈ S(3) acts transitively on these maps. Hence all three resulting semi direct products are isomorphic, so there is really only one non abelian semi direct product of form (Z/3) x_c (Z/2 x Z/2). Since the dihedral group D_6 = Isom(hexagon) has order 12, seven elements of order 2, two elements of order 6, and two elements of order 3, it must be this group.

Theorem:
G is solvable iff all subgroups and quotient groups of G are solvable iff there is one normal subgroup K such that both K and G/K are solvable.
proof:

I. Assume K is normal in G, and that both K and G/K are solvable. Thus we have normal series $K = K_0 > K_1 > .... > K_n = \{e\}$, and $G/K = H_0 > H_1 > .... > H_m = \{[e]\}$, and all quotients K_i/K_{i+1} and H_j/H_{j+1} are abelian. Then define a normal series for G by augmenting that for K, by the pull back of that for G/K. I.e. let $G_j = f^{-1}(H_j)$ where f:G-->G/K is the natural projection. Since the inverse image of a normal group is also normal, all G_j are normal. Hence $G = G_0 > G_1 > ....> G_m = K = K_0 > .....> K_n = \{e\}$ is a normal series for G. The K_i/K_{i+1} are still abelian, $G_m/K_0 = \{e\}$ is abelian, and for j < m, we have $G_j/G_{j+1} \approx (G_j/K)/(G_{j+1}/K) \approx H_j/H_{j+1}$ is abelian. That proves G solvable.

II. Next assume G solvable with abelian normal series $G = G_0 > G_1 > .... > G_m = \{e\}$, and let H be any subgroup. Define H_i = HmeetG$_i$. Then H_{i+1} is not necessarily normal in G, but it is normal in H_i. I.e. conjugating an element y of H_{i+1} = HmeetG$_i$+1 by an element x of H_i = HmeetG$_i$ is conjugating by an element of G_i, and G_{i+1} is normal in G_i. Hence xyx^{-1} lies in G_{i+1}. But x also lies in H, as does y, and H is normal in H, so xyx^{-1} also lies in H. I.e. for all x in H_i and y in H_{i+1}, xyx^{-1} lies in HmeetG$_{i+1}$ = H_{i+1}.

Now H_i/H_{i+1} = (HmeetG$_i$)/(HmeetG$_{i+1}$), so if we map HmeetG$_i$ into G_i/G_{i+1}, the kernel is precisely (HmeetG$_{i+1}$). Hence H_i/H_{i+1} is isomorphic to a subgroup of G_i/G_{i+1}, hence is also abelian. Thus H is solvable.

III. Assume G is solvable with abelian normal series $G = G_0 > G_1 > ....> G_m = \{e\}$, K is normal in G and consider G/K. Define $H_i = (KG_i)/K$. Since the class of elements of K are trivial,

each class in H_i can be represented as [x] for some x in G_i, and similarly each [y] in H_{i+1} can be represented as [y] for some y in G_{i+1}. Thus $[x][y][x^{-1}] = [xyx^{-1}]$ is in H_{i+1}, since xyx^{-1} is in G_{i+1}. Hence H_{i+1} is normal in H_i.

Now consider $H_i/H_{i+1} = (KG_i/K)/(KG_{i+1}/K) \approx (KG_i)/(KG_{i+1})$. Then map $G_i \to KG_i \to (KG_i)/(KG_{i+1})$. The composition map f is onto since again every class [y] in the quotient can be represented by an element y of G_i. Then since the subgroup G_{i+1} of G_i goes to zero under this composition, there is an induced map $[f]:(G_i/G_{i+1}) \to (H_i/H_{i+1})$ which is still surjective. Since H_i/H_{i+1} is thus isomorphic to a quotient of the abelian group G_i/G_{i+1} modded out by the kernel of [f], the quotient H_i/H_{i+1} is also abelian. **QED.**

Composition series

Notice that if G is a cyclic group Z/(mn) of order mn, then G has a cyclic subgroup generated by [m] of order n, whose quotient is cyclic of order m. Hence a cyclic group of order $n = \prod p_i^{r_i}$ has a maximal, non redundant, normal series whose quotients are of prime order, and equal to the prime factors of n. Thus every maximal non redundant normal series for G has the same quotients, up to isomorphism, but possibly in a different order. That this also holds for non abelian groups is called the Jiordan Hoklder theorem.

Definition: A <u>composition series</u> for a group G is a normal series $G = G_0 > G_1 > \ldots > G_m = \{e\}$, in which every quotient group G_i/G_{i+1} is simple but not trivial, (thus a maximal, non redundant, normal series).

Theorem: **(Jordan - Holder)** If a finite group G has two composition series, then they have the same length, and after some renumbering of the quotients.the two sequences of simple quotients are the same, up to isomorphism.

proof: By induction on the order of G, prime order being trivial. Let
$G > G_1 > \ldots > G_m = \{e\}$, and $G > H_1 > \ldots > H_n = \{e\}$, be composition series for G.

case I. $G_1 = H_1$. Then we are done by induction, since the groups $G_1 = H_1$ have smaller order, so their compositiion series are the same length, and have isomorphic quotients, in some order.

case II. G_1 and H_1 are different. Then both G_1, H_1 are maximal proper normal subgroups of G, so their product G_1H_1 is normal and larger than either, hence $G_1H_1 = G$. We want to construct another composition series for G to reduce to case I. Look at $G_1 \text{meet} H_1$. This is not equal to either G_1 or H_1 and is normal in both, so call it K_2, and construct a composition series for $K_2 > K_3 > \ldots > K_s$.

Then we have two new composition series for G: $G > G_1 > K_2 > \ldots > K_s$, and $G > H_1 > K_2 > \ldots > K_s$. To check this, we only have to show that both G_1/K_2 and H_1/K_2 are simple and non trivial. But $G_1/K_2 = G_1/(G_1 \text{meet} H_1) \approx G_1H_1/H_1 = G/H_1$, is simple. Same for H_1/K_2.

Now case I tells us that m = s, and the composition series
(A) $G > G_1 > \ldots > G_m$ and (B) $G > G_1 > K_2 > \ldots > K_s$, have isomorphic quotients. Also n = s, and the series (C) $G > H_1 > K_2 > \ldots > K_s$, and (D) $G > H_1 > \ldots > H_n$ have isomorphic quotients. Since $G_1/K_2 \approx G/H_1$ and $H_1/K_2 \approx G/G_1$, we see series (B) and (C) also have isomorphic quotients. Hence the same holds for series (A) and (D), as desired. **QED.**

Corollary: A group G is solvable if and only if in every composition series for G, all the simple quotients are cyclic of prime order. [Necessarily the orders of the quotients is the sequence of primes in the prime factorization of #G].

Corollary: Prime factorization of integers is unique, up to order of factors.
proof: A prime factorization of n gives a composition series for Z/n. **QED.**

Free groups and free products of groups.

We noted that given two maps g:G-->K, h:H-->K of groups, setting (gh)(x,y) = g(x)h(y) may not define a map GxH-->K, since elements of G commmute with elements of H, but their images in K may not commute. Since we have no restriction on the elements of K, in order to construct a group from which the maps g,h, do always define a map to K, we must allow no commutativity in our "product" at all. Let G = H = Z, the simplest case. Call a the generator of the first copy of Z, and b the generator of the second copy. Since the only relations we allow are those forced by the group laws, we must allow abababab and $a^2bab^{-3}ab$, and so on, all to be different elements. So we define a "word" constructed from the letters a,b, to be any finite sequence of powers of a and b, e.g. $a^{r1}b^{s1}a^{r2}b^{s2}$..... The exponents can be any integers. The sequence where all powers are zero, is the identity element, and words are multiplied by juxtaposition. When we juxtapose two words, the powers of a and b may not alternate, so we combine adjacent powers of the same letter. The trivial word has only zero exponents. A non trivial word is <u>reduced</u> if it has no adjacent powers of the same letter and no zero exponents. We also consider the trivial word to be reduced and write it as e.

Clearly inverses exist, and the trivial word is the identity since $e = x^0 = y^0$. Associativity is not so easy, but there is a nice proof in Mike Artin's book, which I copy.

Artin calls a <u>word</u> a finite sequence of the elements a,b,a^{-1},b^{-1}, and a reduction of a word is obtained by a cancellation of some adjacent pair, of form xx^{-1}, or by a sequence of such cancellations. A reduced word is one in which no cancellations are possible. The main point is that starting from any word and performing cancellations, there is only one possible reduced result. This is true if the word is already reduced, for example if it has length zero or one, so use induction on the length of the word. If a word is not reduced it must contain a pair of form $\mathbf{xx^{-1}}$. If we cancel this pair first, the induction hypothesis says there is only one possible reduced result for this word. If we perform some other sequence of cancellations, and eventually cancel this \pair, we might as well have canceled it first, and the same result holds. If on the iother hand we can one of these elements but not both, we must do it as follows: by cancelling the first (i.e. leftmost) two of $x^{-1}\mathbf{xx^{-1}}$, or the last (i.e. rightmost) two of $\mathbf{xx^{-1}}x$. Either way, the result is the same as if we had canceled the original pair, so the argument in that case holds. **QED.**

Definition: The set of reduced words in the letters {a,b}, with the operation of juxtaposition, is the <u>free</u> <u>group</u> on those two letters. The empty word is the identity, written $e = a^0 = b^0$. We shorten the notation by using higher exponents for adjacent occurrences of the same letter.

Exercise: Associativity follows easily from the (obvious) associativity on the set of unreduced words, by the uniqueness result above for reduction.

Definition: The free product of two copies of Z is defined as the free group Fr(a,b) on two letters. It is easy to see that any two group maps f:Z-->K and g:Z-->K define a unique group map (fxg):Fr(a,b)-->K.

There is one plausible result that is not too hard.
Theorem: If G = Fr(a,b) is the free group on two letters, and G' is the commutator subgroup, the quotient G/G' ≈ Z x Z, the free abelian group on two letters. [**hint:** prove they have the same mapping property.]

Remark: The same construction, with the same proof, defines a free group on any set of letters, and proves the existence of a free product of any number of copies of the group Z. It follows that every group G is the image of a homomorphism from a free group Fr(S)-->G, but it is hard to make much use of this. I.e. unlike the abelian case, the free groups, even on two letters, are very complicated and hard to understand.

Theorem: 1) Every free group on any finite or countable number of letters, is a subgroup of Fr(a,b).
2) Conversely, every subgroup of Fr(a,b) is isomorphic to a free group on a finite or countable set of letters. [look at π_1(figure eight).]
"proof": If X(n) = the "figure eight" with n loops, then X(n), n ≥ 2, is a covering space of X(2), and $\pi_1(X(n)) \approx Fr(a_1,...,a_n)$, same for X(infinity). This proves 1), by the homotopy lifting property. For 2) given any subgroup of $\pi_1(X(2))$, it defines a covering space Y whose π_1 is that subgroup. But the figure eight is a graph, and every covering space of a graph is again a graph (1 dimensional complex), hence homotopic to a wedge of circles, so π_1 is again free. **qed.**

8000 Day 5 Field extensions and groups of automorphisms.

Introduction: Galois theory is concerned with the self mappings of a field, i.e. automorphisms of a field E, that are specified on some subfield, for example that equal the identity on some subfield k. We want to see how to construct such automorphisms, to count how many there are, and to compute their exact fixed field. If E has finite vector dimension n over k, we will see there are at most n automorphisms of E that fix k pointwise. The reason is simple. It will turn out that such maps must send roots in E of polynomials with coefficients in k, to other roots in E of these polynomials. It follows that the number of such automorphisms will be determined by the number of distinct roots such polynomials have in E, which is always bounded by the degree of the polynomials. Galois theory is most useful when the number of distinct roots of an irreducible polynomial equals its degree. This is the "separable" case.

Since the vector dimension of a field extension can be computed in terms of the degree of the polynomials satisfied by generators for the extension, it will follow that the number of automorphisms is also bounded by the vector dimension of the extension, and that equality holds in the separable case. The proofs proceed by carrying out the "simple" case first, then using induction to deduce the result for any sequence of simple extensions, i.e. any finite extension.

First we review a few technical facts about extensions that are probably familiar from math 6000.

Def1: If k is a subfield of E, then E is a vector space over k, and we write [E:k] for the k dimension of E, also called the degree of E over k.

Def 2: If k is a subfield of a field E, an element a of E is algebraic over k, iff a is a root of some non zero polynomial f in k[X].

Lemma 3: If k is a subfield of a field E, and a in E is algebraic over k, there is a unique monic irreducible polynomial in k[X] satisfied by a, i.e. the unique monic polynomial in k[X] of lowest degree with a as root. This is called the minimal polynomial of a over k.
proof: If a is algebraic over k, the evaluation map k[X]-->E, sending g to g(a) has non trivial kernel of form (f), inducing an injection k[X]/(f)-->E. So (f) is a maximal, prime ideal generated by a unique monic irreducible polynomial, the monic polynomial of lowest degree in the kernel. **QED.**

Definitions:
4. The ring generated by a over k = k[a] = the k-algebra generated by a.
If k is a subfield of E, and a is an element of E, k[a] is the intersection of all subrings of E that contain both a and k. Concretely k[a] = {f(a): for all polynomials f in k[X]}.

5. The field generated by a over k = k(a) = the intersection of all subfields of E that contain both a and k. Concretely k(a) = {f(a)/g(a): f,g are in k[X], and g(a) ≠ 0}.

Theorem 6: If k is a subfield of E, and a is an element of E, then TFAE:
i) a is algebraic over k.
ii) k[a] = k(a).
iii) k(a) has finite (vector) dimension over k.
iv) k[a] has finite (vector) dimension over k.
v) k[a] is contained in a finite dimensional k - subspace of E.
vi) k(a) ≈ k[X]/(f) where f is an irreducible polynomial in k[X].
In particular, if a is algebraic over k, the k dimension of k(a) equals the degree of the minimal polynomial of a over k.
proof: [sketch]: If a is algebraic over k, then for some n, a^n is a linear combination of lower powers of a. This implies every power a^m with m>n, is also a linear combination of powers of a less than n. Hence the dimension of k[a] over k is finite. If we take a monic dependency relation $c_0+c_1a+c_2a^2+...+a^n$ for the smallest possible power a^n, then a satisfies the monic polynomial $f(X) = c_0+c_1X+c_2X^2+...+X^n$, but no polynomial of lower degree, so f is irreducible, and the map k[X]/(f)-->k[a] is injective and surjective. Since k[X]/(f) is a field, so is k[a] = k(a). If k[a] is contained in some finite dimensional k subspace of E, say of dimension n, then $1,a,a^2,...,a^n$ are dependent/k, and a dependency relation gives a polynomial satisfied by a, so a is algebraic. The k - dimension of k[X]/(f) is deg(f) = n, and a basis is given by [1], [X],....,$[X^{n-1}]$. If a is not algebraic over k, then the infinite sequence of powers $\{1,a,a^2,...,a^n,...\}$ is linearly independent over k. **QED.**

Lemma 7: If k is a subfield of E, E a subfield of F, then [F:k] = [F:E][E:k].

proof: If $x_1,...,x_n$ is a k basis for E, and $y_1,...,y_m$ an E basis for F, then $\{x_iy_j\}$, $1\leq i\leq n$, $1\leq j\leq m$, is a k basis for F. This is trivial to check by changing the order of summation. E.g. if z lies in F, then there exist constants $b_1,...,b_m$ in E such that $z = b_1y_1+...+b_my_m$. But each b_j lies in E, so there exist constants a_{ij} such that each $b_j = a_{1j}x1+...+a_{nj}x_n$. Hence $z = b_1y_1+...+b_jy_j+....b_my_m = ...+(a_{1j}x1+...+a_{nj}x_n)y_j+..... = ...+a_{ij}(x_iy_j)+.....$

Thus the products $\{x_iy_j\}$ span F over k. In particular, if both [F:E] and [E:k] are finite, so is [F:k]. You should check that the products $\{x_iy_j\}$ are independent over k, as this is a favorite little prelim question. **QED.**

Cor 8: If E is a field containing k, the subset of E consisting of elements which are algebraic over k, is a subfield of E.
proof: If a,b in E are algebraic over k, we must show that a+b, ab, a/b, are also algebraic over k. But k(a) has finite dimension over k, and since k(b) also has finite dimension over k, it has even smaller dimension over k(a). Thus both [k(a,b):k(a)] and [k(a):k] are finite. Hence also [k(a,b):k] is finite, and so all field combinations of a,b, such as a+b, ab, a/b, etc,.. , belong to finite dimensional subfields of E, hence are algebraic over k. **QED.**

Note 9: All finite dimensional extensions of k are algebraic, but not all algebraic extensions are finite dimensional. We know simple algebraic extensions are finite dimensional. Hence finitely generated (as fields) algebraic extensions, are finite dimensional (as vector spaces).

Cor 10: A field extension of k is finite dimensional (as a vector space), if and only if it is both finitely generated (as a field) and algebraic/k.

Next we show that a given field k has lots of algebraic extensions, and in particular any family of k polynomials has roots in a suitable extension.
Lemma 11: If g is any polynomial over the field k, and f an irreducible factor of g, there is a field extension of k in which f, and hence g has a root.
Proof: Consider the quotient ring k[X]/(f) which is a field since f is irreducible. I.e. k[X] is a pid so an irreducible f generates a maximal ideal, and thus k[X]/(f) is a field. The map k-->k[X]/(f) is injective since no element of k is divisible by f, so we may regard k as a subfield of k[X]/(f) by identifying elements of k with their images in k[X]/(f).

The class [X] of the variable X is the desired root of f. I.e. since the map k[X]-->k[X]/(f) taking f(X) to [f(X)] = f([X]) is a homomorphism, we have f([X]) = [f(X)] = [0]. **QED.**

By induction we can adjoin a full set of roots of g.
Lemma 12: If g is any polynomial over the field k, there is a field extension E of k in which g factors into linear factors, and we say that g has a full set of roots, or splits, in E.
proof: Keep adjoining one root at a time of irreducible factors of g until all irreducible factors have degree one. **QED.**

Def 13: A "splitting field" for a polynomial f in k[X], ios a field E containing k, such that f splits into linear factors in E[X], and E is generated as a field by the roots of f in E.

Proposition 14: If f is a polynomial in k[X], there is a splitting field for f.

proof: Find a field as above in which f splits into linear factors, and take the subfield generated by the roots of f. **QED.**

Proposition 15: If k is any field, there is an extension E of k in which every irreducible polynomial in k[X] has a root.
proof: Let S be a set of variables X_f, one for each irreducible polynomial f in k[X], and let k[S] be the polynomial ring in the variables X_f. For each irreducible polynomial f(X) in k[X], consider the polynomial $f(X_f)$ in k[S]. Let T be the set of all these polynomials in k[S]. We claim T does not generate the unit ideal of k[S]. E.g. if there were a finite linear combination $\sum g_i f_i(X_{fi}) = 1$, then by extending k to a field where all of the polynomials f_i have roots, and substituting in these roots, we would have 0 = 1.

Hence there is some maximal ideal M containing the set T. Then E = k[S]/M is a field, and in this quotient field every class $f([X_f]) = [f(X_f)] = [0]$. I.e. $[X_f]$ is a root of f. Thus every irreducible polynomial f in k[X] has a root in E. **QED.**

Proposition 16: If k is any field, there is an extension E of k in which every polynomial in k[X] splits completely into linear factors.
proof: Let E_1 be a field containing k, in which every irreducible polynomial in k[X] has a root. Then let E_2 be an extension of E_1 where every irreducible polynomial in $E_1[X]$ has a root. Etc..... Then take the union E of all these fields. If f is a polynomial of degree n in k[X]. Then f splits completely into linear factors in $E_n[X]$, hence also in E. **QED.**

Def 17: A field E such that every polynomial in E[X] has a root in E, hence splits completely, is called algebraically closed. An algebraically closed field which is algebraic over a subfield k, is called an algebraic closure of k.

Theorem 18: Any field k has an algebraic closure.
proof: Let F be the field constructed above, in which every polynomial in k[X] splits into linear factors, and let E be the subfield generated by the roots of all polynomials over k. Then E is algebraic over k. Let g be a polynomial in E[X], and let a be a root of g in an extension field of E. Then a is algebraic over E, and E is algebraic over k, so a is algebraic over k. Thus a is a root of some polynomial in k[X] and hence a is already contained in E. **QED.**

Extending homomorphisms.

Given a (homomorphic) map, f:k-->k', of fields which are subfields of larger fields E, E', we want to know when it is possible to extend f to a map f':E-->E'. Of course field maps are injective whenever they exist.

Field map extensions are always done one generator at a time, and for that "simple" case we use the fundamental isomorphism $F(a) \approx F[X]/(g)$ where g is the minimal F polynomial of a, to tell us how to extend the map.

Theorem 19: Assume f:k-->k' is a map of subfields of E, E'.
i) f extends uniquely to a ring map k[X]-->k'[X], by applying f to the coefficients of each polynomial.

Let a in E have minimal polynomial g in k[X], mapping to g' in k'[X].
2) f extends to a map k(a)-->E' with f(a) = a', iff a' is a root of g'.
3) The number of extensions of f to f':k(a)-->E' equals the number of distinct roots in E' of g', and is at most the degree of g = $\dim_k(k(a))$.

proof: i) We check (g+h)' = g'+h' and (gh)' = g'h'. The coefficient of X^n in g+h is the sum a_n+b_n where a,b, are the coefficients in g,h. And the coefficient of X^n in (g+h)' = f(a_n+b_n) = $(a_n+b_n)'$ = $(a_n)'+(b_n)'$ = the sum of the coefficents of X^n in g' and h'. Hence (g+h)' = g'+h'.

The coefficient of X^n in gh is the sum of the products a_ib_j over all i,j, with i+j = n. f applied to this sum is the sum of the corresponding products $a_i'b_j'$, which is the corresponding coefficient in g'h'. Hence (gh)' = g'h'.

ii) Applying f to g(a) = 0, gives g'(a') = 0, so the condition is necessary. If g'(a')=0 for some a' in E', we get a map k'[X]-->E' sending X to a', inducing a map k'[X]/(g')-->E' sending [X] to a', hence by composition a map k(a)-->k[X]/(g)-->k'[X]/(g')-->E', sending a to a'.

iii) Each different choice for f(a) gives a different map k(a)-->E', so the number of maps equals the number of roots a' of g' in E'. The number of roots of g' in the field E' never exceeds the degree of g', which equals the degree of g, which equals the k dimension of k(a). **QED.**

Cor 20: In the setting above, the number of extensions of f:k-->k', to f':k(a)-->E' is $\leq \dim_k(k(a))$, and equals this number if g' has deg(g') distinct roots in E', equivalently if g' factors into distinct linear factors in E'[X].

Note 21: Even if g' has d = deg(g') distinct roots in E', and there exist d distinct maps k'-->E' extending f, all of these maps may have the same image in E'. I.e. if $\{a_1', a_2', \ldots, a_d'\}$ are the distinct roots of g' in E', and k-->k' is an isomorphism, then all the fields k'(a_i') are isomorphic, but it may or may not happen that they are equal as subfields of E'. We call such isomorphic subfields "conjugate" over k', and when they are all equal we say they are all normal. We will say more about normal extensions later, and show how in this situation they define normal subgroups of the group of automorphisms of E'.

The argument for extending maps from k to k(a), lets us extend maps to any algebraic extension, if the target field has enough roots.

We have to be a little careful about the hypotheses, since if k has more than one generator, it does not suffice to require just one root in E' for each of their minimal polynomials. For instance, suppose g is an irreducible polynomial in k[X] with distinct roots a,b and E = k(a,b), while E' = k(a) ≠ k(a,b). Then a,b, have the same minimal polynomial over k, namely g, so in both cases, the minimal polynomial of a and b does have a root in E'. Still there is no extension of the identity map k-->k to a map k(a)-->k(a,b).

So we use a little stronger hypothesis on E' in the next result.
Theorem 22: Assume E = $k(a_1,\ldots a_n)$ is finitely generated and algebraic, i.e. finite dimensional over k, f:k-->k' a field map and E' a field containing k'. Assume further for every generator a_i of E, that g_i' factors completely into linear factors in E'[X], where g_i' in k'[X] is the image under f of the

minimal polynomial g_i of a_i, i.e. that E' contains a splitting field for the polynomial $\prod g_i'$.

Then there exist extensions of f to f':E-->E'. The number of such extensions is at most equal to the dimension [E:k], and equals this dimension if $\prod g_i'$ factors into distinct linear factors in E'[X].

proof: (existence of extensions): From the argument in the simple case we get an extension to $k(a_1)$-->E'. Now we want to extend further to $k(a_1, a_2)$-->E'. There is only a tiny difference from the simple case. We regard this as a simple extension $k_1(a_2)$, of $k(a_1) = k_1$. But our hypothesis only says the minimal polynomial g_2 of a_2 over k corresponds to a polynomial g_2' in E'[X] that splits completely into linear factors. We need to know about the minimal polynomial h_2 of a_2 over the field k_1.

The point is that h_2 is always a factor of g_2. I.e. both h_2 and g_2 have coefficients in k_1, and h_2 is irreducible there. Since a_2 is a root of both polynomials, h_2 must divide g_2 in $k_1[X]$. Thus the corresponding polynomial h_2' in E'[X] is a factor of g_2', and since g_2' factors in E'[X] completely into linear factors, some of those same factors give a factorization of h_2' into linear factors.

Extending successively over each simple extension, we eventually get an extension to E-->E', by induction on the dimension [E:k].

(number of extensions): By the argument in the simple case, i.e. by the corollary above, at each stage the number of extensions from k_i-->E' to k_{i+1}-->E' is at most equal to the dimension $[k_{i+1}:k_i]$, and equals that dimension if h_i' factors in E'[X] into distinct linear factors. But the number of extensions of f from k-->k' to f'E-->E' is the product of the number at each stage, and the dimension of successive field extensions is also multiplicative, i.e. $[E:k] = \{E:k_{n-1}][.....][k_1:k]$. Thus the number of extensions is at most equal to [E:k], and equals this dimension when every polynomial h_i' factors into distinct linear factors in E'[X]. But since the linear factors of h_i' are a subset of the linear factors of g_i', if $\prod g_i'$ factors into distinct linear factors in E'[X], then so does every h_i'. **QED.**

Next we use Zorn's lemma in the infinite dimensional case. The moral is that there are only three steps to the argument: **i)** the simple case; **ii)** the observation that the minimal polynomial over a larger extension is a factor of the minimal polynomial for a smaller one, which allows you to repeat the simple case more than once; **iii)** using Zorn, i.e. "transfinite induction", to find a maximal extension.

Theorem 23: Assume E is an algebraic extension of k, f:k-->k' is a field map and E' a field containing k'. Assume further for every element a of E, that g' factors completely into linear factors in E'[X], where g' in k'[X] is the image under f of the minimal polynomial g of a. This is true for example if E' is algebraically closed. Then there exist extensions of f to f':E-->E'.

proof: Consider the set of all partial extensions of f to F-->E' where F is a field intermediate between k and E. These form a partially ordered set where g > h if g extends h. (Each such partial extension is a map from a subset of E to E', hence its graph is a subset of ExE'. The collection of all such partial extensions is a subset of the set of all subsets of ExE', in particular it is a set, to which we can try to apply Zorn's lemma.)

Given any totally ordered collection or "chain" of partial extensions, they define a partial extension on the union of their domains, and this is an upper bound for the chain. Since the

hypothesis of Zorn is thus satisfied, there exist maximal partial extensions. We claim any maximal partial extension is actually defined on all of E.

If g:F-->E' is an extension of f that is not defined on all of E, there is some element a of E that is not in F. Then a is algebraic over k hence also over F, and satisfies an irreducible polynomial over F which is a factor of its irreducible polynomial over k. The corresponding polynomial over E' thus factors completely into linear factors as argued above, and we get an extension of f to F(a)-->E'. This shows any extension of f whose domain is not all of E cannot be maximal, hence any maximal extension is defined on all of E. **QED.**

Cor 24: If E is an algebraic closure of k, then every algebraic extensiion of k has an isomorphic copy inside of E. In particular, any two algebraic closures of k are isomorphic.
proof: The theorem gives an isomorphism from any algebraic etension F of k to a subfield F' of E. If F is algebraically closed then F' is also algebraically closed. E is algebraic over k, hence over F', so E = F'. **QED.**

Remark 25: We may speak of "the" algebraic closure E of a field k, although the elements of E are not really unique, since E has alots of automorphisms exchanging one element with another. E.g. the element i in the complex numbers is not intrinsically determined in the algebraic closure of R, only the pair {i,-i} is intrinsically determined by R.

Separability

The phenomenon of polynomials with multiple roots also bears examination, since it affects the number of extensions of a homomorphism.
Def/Exercise 26: A polynomial g in k[X] is separable iff gcd(g, dg/dx) = 1 in k[X], iff g, dg/dx have no common root in any extension of k, iff f has no multiple linear factor in any extension of k, iff in some extension of k, g has deg(g) distinct roots, hence factors into distinct linear factors.

Def 27: An algebraic element a of a field E containing k is separable over k iff its minimal polynomial in k[X] is separable.

An element separable over k is separable over any larger field F, since its minimal F-poynomial divides its minimal k -polynomial.

Theorem 28: If E is a field containing k, the elements of E that are separable over k form a subfield of the field of algebraic elements.
proof: Let a,b, be elements of E that are separable over k. We claim every element of k(a,b) is separable over k. Let c be any element of k(a,b), and let f,g,h, be the minimal polynomials of a,b,c over k, with f,g separable. Let E' be a splitting field for (fgh) over k. Since f,g both factor into distinct linear factors in E', there are exactly [k(a,b):k] extensions of the inclusion map k-->E' to maps of k(a,b)-->E'.

If c is not separable over k, there are fewer than deg(h) extensions of k-->E' to k(c)-->E', hence fewer than [k(a,b,c):k] extensions of k-->E' to k(a,b,c)-->E'. But since k(a,b,c) = k(a,b) this is a contradiction. So c is separable over k. **QED.**

Def 29: The subfield of separable elements in the algebraic closure of k is called the separable algebraic closure of k.

Cor 30: If $E = k(a_1,...a_n)$ is a finite separable extension of k, g is a polynomial in k[X] satisfied by the generators a_i, and E' contains a splitting field for g, there exist exactly [E:k] maps E-->E', extending the identity on k.

Cor 31: If f is a separable polynomial over k, E a splitting field for f, there are exactly [E:k] automorphisms of E which extend the identity on k.

Remark 32: All algebraic extensions are separable in characteristic zero. All algebraic extensions of finite fields are also separable. A field is called perfect if all its algebraic extensions are separable.

Def 33: The "fixed field" of a set S of automorphisms of a field E is the subfield of E of elements left identically fixed by every element of S. Automorphisms of E leaving a subfield k fixed are called k automorphisms.

Def 34: If E is any algebraic extensiion of k, the group of k automorphisms of E is called $Gal_k(E)$ = the galois group of E over k.

Cor 35: If k is any field, f a separable polynomial over k, and E a splitting field for f, the fixed field of $G = Gal_k(E)$ equals k.
proof: If G fixed a larger subfield F of E containing k, there would be more than [E:F] automorphisms of E fixing F, contradicting Theorem 2. **QED.**

Remark 36: DF reserves the use of term galois group for the situation in Cor 35, a splitting field extension of a separable polynomial.

Cor 37: If E is a splitting field for a separable polynomial f over k, F an intermediate field, the fixed field of the subgroup $Gal_F(E)$ of $Gal_k(E)$, is F.
proof: E is also a splitting field for f over F. **QED.**

Hence if E is a splitting field of a finite separable polynomial over k, the map from subgroups of $G = Gal_k(E)$ to their fixed fields between k and E, is surjective. I.e. every field between k and E is the fixed field of some subgroup of $Gal_k(E)$.

Theorem 38: If f is a separable polynomial of degree n over k, with splitting field E, the group $Gal_k(E)$ is isomorphic to a subgroup of the symmetric group S(n). In particular it has finite order dividing n!
proof: Every k automorphism of E permutes the roots of f, hence G acts on the set of n roots. But the roots generate E, so this action determines the element of G, hence the map G-->S(n) is injective. **QED.**

Cor 39: If E is the splitting field of a separable polynomial over k, there are only a finite number of intermediate fields between k and E.

It remains to show in this setting that no two subgroups of $Gal_k(E)$ can fix exactly the

same subfield of E, so that the number of subgroups and intermediate fields is the same. I.e. it could conceivably occur that a subgroup H fixes exactly F, but that actually also more elements of G also fix F, so the full subgroup fixing F is larger than H. This never happens.

Theorem 40: If G is a finite group of automorphisms of a field E, and k the subfield fixed by G, then [E:k] = #G.

proof: We know #G ≤ [E:k], so we want to show any subset $x_1,...,x_n$ of E with n > #G is dependent over k. This a matter of solving linear equations. If $f_1,...,f_m$ are the elements of G, the m by n matrix $[f_j(x_i)]$ over E has a non zero solution vector $c = (c_1,...,c_n)$ in E^n, since n > m. Thus for all i, **(*)** $\sum_j c_j f_i(x_j) = 0$. Choose c with as few non zero entries as possible. There must be at least two ≠ 0 entries, since we may assume all $x_j \neq 0$, else dependency is obvious, and the f_i are automorphisms of E. By reordering the x's, we may assume $c_1 \neq 0$, and multiplying through by c_1^{-1} we may assume $c_1 = 1$ belongs to k.

We claim actually all c_j belong to k, yielding a k - linear relation among the elements $(f_i(x_1),....f_i(x_j),...,f_i(x_n))$ for every i. Since one of the f's is the identity, this will give a k linear relation among the (x_j) as claimed.

If some $c_r \neq c_1$ is not in k, then since k is the fixed field of G, some $f_s \neq id$ does not fix c_r. We can renumber so that $f_r(c_r) \neq c_r$. Then apply f_r to the system of equations getting $\sum_j f_r(c_j f_i(x_j)) = \sum_j f_r(c_j)\, f_r f_i(x_j) = 0$ for all i. But as f_i runs over the group G, the product $f_r f_i$ does the same. So we can say also that

(**) $\sum_j f_r(c_j)\, f_i(x_j) = 0$ for all i. Now if we subtract the two systems of equations (*) and (**), we get $\sum_j [f_r(c_j)-cj]\, f_i(x_j) = 0$ for all i. In this system, since c_1 is in k, $f_r(c_1)-c_1 = 0$, but $f_r(c_r)-c_r \neq 0$. Hence we have a non zero solution vector $(,...,f_r(c_j)-c_j,....)$ in the kernel of the matrix, but with fewer non zero entries than before, a contradiction. **QED.**

Cor 41: If E is a finite dimensional extension of a field k, **TFAE:**

i) $\#Gal_k(E) = [E:k]$.

ii) E is the splitting field of a separable polynomial.

iii) k is the fixed field of some finite group of automorphisms of E.

proof: We proved ii) implies i) and iii), and that iii) implies i) above. So it suffices to prove i) implies ii). By the arguments used above, i) implies the minimal polynomial of every generator of E over k splits into distinct linear factors in E, so E is a splitting field for the product of the minimal polynomials of a finite set of generators for E over k, where each of these polynomials is separable. It remains to show their product may be assumed separable. If any two of these polynomials are equal we may omit one. So we want to rule out that two distinct separable irreducible polynomials over k share a root in E. But any element of E has a unique irreducible minimal polynomial over k, so this is impossible. **QED.**

Def 42: If the conditions in 41. hold, call E a <u>finite</u> <u>galois</u> extension of k.

Cor 43: If E is a finite galois extension of k, different subgroups of $G = Gal_k(E)$ have different fixed fields.

proof: If E is galois over k, and H is a subgroup of G, let F be the fixed field of H, and K the

subgroup of all elements of G fixing F. By the previous corollary, #H = [E:F]. H is a subgroup of K. But theorem 2 shows that #K ≤ [E:F] = #H. Thus H = K. **QED.**

We have proved the following.

Th 44: Fundamental theorem of galois theory:
If E is a finite galois extension of k,
(i) the "galois" group G of k automorphisms of E has order #(G) = $\dim_k(E)$;
(ii) the fixed field of G is precisely k;
(iii) the correspondence between subgroups of G and the subfields of E they leave pointwise fixed, is a bijection between subgroups and intermediate fields; e.g. the number of intermediate fields is finite.

It is often useful to know that a finite separable extension is actually simple. This is called the theorem of the "primitive element".

Lemma 45: If G is a finite subgroup of units in a field k, then G is cyclic.
proof: G is finite abelian hence a product of cyclic groups whose orders divide each other. If G is not cyclic, there is an integer 0 < d < #G such that every element of G satisfies the polynomial $X^d - 1 = 0$. But in a field, this polynomial can have no more than d roots. So G is cyclic. **QED.**

Cor 46: If k is a finite field, every finite extension E is simple.
proof: The group of units in E is cyclic, so any generator of this group, also generates the field extension. **QED.**

Lemma 47: If k is an infinite field, a finitely generated algebraic extension E with only finitely many subfields is simple.
proof: Let a,b be among the generators of E. Consider the subfields of E generated over k by the elements a + tb, for all t in k. Since only finitely many subfields exist, and k is infinite, for some s,t in k, k(a+sb) = k(a+tb). This field also contains (a+sb)-(a+tb) = (s-t)b, where s-t ≠ 0 in k. Hence k(a+sb) also contains b, and hence also a, so k(a,b) = k(a+sb). This means we have reduced the number of generators by one. Continuing, we get down to one generator.

Another way to see this is to note that subfields are k vector subspaces, and over an infinite field, no finite union of proper subspaces can fill up a vector space. Then any element of E not lying in one of the finite number of proper subfields generates E over k. **QED.**

Cor 48 (theorem of the "primitive element")**:** A finite separable extension E of a field k, has only finitely many intermediate fields, hence is simple.
proof: Enlarge E to F by adding in all roots of the product of the minimal polynomials of the generators of E over k. F is a splitting field of a separable polynomial, hence there are only a finite number of intermediate fields between k and F, hence also between k and E. **QED.**

Remark 49: (i) One sometimes needs the stronger result that if all but one of a finite set of algebraic field generators is separable, the extension is simple. [**sketch**: If a is separable over k, but b is not, let a= $a_1, a_2,...,a_n$, and b = $b_1,...,b_m$, be the distinct roots of their minimal polynomials f,g. Choose c any element of k different from all quotients $(b_i-b)/(a-a_j)$ for all i and all j > 1, let e = b+ca, and set h(X) = g(e-cX). Then h(a) = g(b) = 0, so a is a common root of f,h.

By choice of c, $e-ca_j = b + c(a-a_j) \neq b_i$ for any i, and any $j > 1$. The only common root of f,h, is X=a, and f has no multiple roots, so their gcd is X-a. Since f,h, both have coefficients in k(e), their gcd does too, hence a belongs to k(e). Thus k(e) = k(a,b). Continue by induction.]

(ii) Note that in a galois extension E/k of finite dimension d, an element c of E generates E over k iff its minimal polynomial over k has degree d, if and only if its orbit under the galois group $Gal_k(E)$ has d elements, iff the only element of G fixing c is the identity. In general an element c of E generates a field of dimension equal to the degree of its minimal polynomial = the number of distinct elements of its Galois orbit.

Normal extensions

A galois extension is one whose elements have separable minimal polynomials, and such that those polynomials split in the given extension. These two conditions are both important, the separability and the splitting condition. We have discussed separability, and next we discuss the splitting condition, called normality.

Caution: Some books, e.g. the classic notre dame lectures of E. Artin himself, use the word "normal" for what we call "galois".

This concerns the question of when different maps E-->E' extending f:k-->k', actually map onto different subfields of E'. For this discussion we suppose f:k-->k' is an isomorphism, i.e. surjective as well as injective.

Example 50: X^3-2. The polynomial $g(X) = X^3-2$ is irreducible in Q[X] by Eisenstein. Let the three roots be a, b, c, with a the real root, and b = wa, $c = w^2a$, where $w = e^{2\pi i/3}$ is a primitive (complex) cube root of 1. These roots are distinct. Indeed an irreducible polynomial over any field containing Q (i.e. characteristic zero) always has distinct roots.

Let f:Q-->C be the inclusion of Q into the complex field C. Then there are exactly three extensions of f to maps f':Q(a)-->C. The simplest is the inclusion with f'(a) = a, but f'(a) = b, and f'(a) = c also define field maps. We claim the three images Q(a), Q(b), Q(c), are all distinct subfields of C, although they are isomorphic over Q.

Certainly Q(a) does not equal either Q(b) or Q(c), since Q(a) contains only real numbers and b,c, are both complex. But also Q(b) and Q(c) are different fields, since if b,c were in the same field, then c/b = w also would be there, and then b/w = a would be there. It follows too that Q(a,b) = Q(b,c) = Q(a,c) = Q(a,b,c).

In this situation, we say the distinct fields Q(a), Q(b), Q(c) are conjugate subfields of C over Q, and that none of them are normal over Q.

Def: If k is a subfield of F, and $\{E_i\}$ are intermediate fields, we say the fields $\{E_i\}$ are conjugate over k, if for each pair i,j there is an isomorphism from E_i to E_j that equals the identity on k. Thus intermediate fields between k and E are "conjugate over k", if they are k isomorphic.

Def 51: An extension E of k is <u>normal</u> if in every larger field F, E is conjugate over k only to itself.

Useful fact 52: If an extension E of k is normal, then for every larger field F, every k

automorphism of F restricts to a k automorphism of E.

Thus there are two concepts controlling the number of k automorphisms an extension E can have: if the extension is not normal, some of the k isomorphisms from E into a larger field F might not map into E, hence some k isomorphisms from E to F will not be automorphisms of E. And if E is not separable over k, even if it is normal, there just are not the maximum number of different maps at all.

Since there are fewer maps in the absence of separability, some books introduce a concept called "separable degree" which measures the number of maps. This is possibly smaller than the degree = dimension of the extension, but it also behaves multiplicatively. We will not use this concept, nor the dual concept of inseparable degree. [See Lang, Algebra.]

Cor 53: If E is a finite galois extension of k, and F an intermediate field normal over k, then restriction of automorphisms defines a surjective group map $Gal_k(E)\text{--->}Gal_k(F)$ with kernel $Gal_F(E)$.

proof: By the extension theory of homomorphisms above, the converse of the useful fact is also true, i.e. every k automorphism of the normal extension E extends to a k automorphism of F. **QED.**

This lets us refine the fundamental theorem of galois theory:

Cor 54: If E is finite and galois over k, in the correspondence between subgroups of $Gal_k(E)$ and intermediate fields, conjugate subgroups correspond to conjugate subfields, hence normal subgroups have normal fixed fields, and vice versa.

proof: This is familiar from our study of group actions and isotropy subgroups, i.e. that isotropy subgroups for elements of the same orbit are conjugate. I.e. let F be the fixed field of H and L the fixed field of K. Then if H and K are conjugate subgroups, say $g^{-1}Hg = K$, then K fixes $g^{-1}F$, and H fixes gL, so L and F are conjugate fields. Vice versa, if F = gL are conjugate fields, then $g^{-1}Hg$ fixes $L = g^{-1}F$, so $K = g^{-1}Hg$ are conjugate groups. In particular, normal subgroups correspond to normal extensions. **QED.**

Theorem 55: An extension E is normal over k if and only if it is a splitting field for a family of k polynomials, if and only if it contains a splitting field of the minimal polynomial of each of its elements.

Thus a finite extension E of k is normal if and only if it is a splitting field for one polynomial (not necessarily separable) over k.

proof: If some element has a minimal polynomial which does not split in E, then our technique of extending homomorphisms allows the construction of a k homomorphism from E into some larger field, that sends that element outside of E, so E would not be normal. Since such maps are determined by what they do to any set of generators, there can be no set of generators whose minimal polynomials split in E. Thus if E is normal, the minimal k polynomials of all elements split in E, which is equivalent to E being a splitting field for some family of k plynomials. Moreover, if all minimal polynomials split, then no map of E into a larger field over k can map E outside of E, since every element must go to another root of its miknimal polynomial. **QED.**

Lemma 56: An algebraic extension E of k lies in a "smallest" normal extension.

proof: In an algebraic closure for E, take the field generated by all subfields k - conjugate to E. **QED.**

Def 57: The smallest normal extension of k containing a given extension field E, as in 56, is called the normal closure of E/k.

Example 58: To construct a normal extension that is not separable, assume k has characteristic p, and there is an element b of k that has no pth root in k. Then we adjoin a pth root of b, say c, to get k(c) = E. In characteristic p, $(X-c)^p = X^p-c^p = X^p-b$, so c is a p-fold root of X^p-b which thus has no other roots. Hence E is a non trivial splitting field over k, but every k automorphism of E sends c to c, so the only k automorphism of E is id. Thus our non trivial normal extension has Galois group G = {id}.

In practice, fields in which an element has no pth root are not so immediate to find. In a finite field k of characteristic p, the pth power map from k to k is injective by the odd fact that $(a-b)^p = a^p-b^p$, hence also surjective, so all elements have pth roots. But if we introduce a variable T, and form the rational function field k(T), where say k = Z/p, then T has no pth root. Equivalently, the extension k(T) is normal but not separable over the subfield field $k(T^p)$.

Remark 57: If E is a finite and normal over k, there are two natural subfields, the fixed field F of $Gal_k(E)$, which may be larger than k, and the field L of all elements of E which are separable over k. Then L is galois over k, and E is galois over F. Moreover the restriction map $Gal_k(E) = Gal_F(E) \to Gal_k(L)$ is an isomorphism. Presumably $dim_F(E) = dim_k(L) = \#Gal_k(E)$ = "separable degree" of E over k, and $dim_k(F) = dim_L(E)$ = "inseparable degree" of E over k.

58: Remark on the infinite dimensional case: A general "galois" extension is one which is both normal and separable. Although we study their galois groups only in the finite case, the infinite ones are also of great interest. In that case there are more subgroups than intermediate fields, in the sense that different subgroups can have the same fixed field, so the statement of the fundamental theorem must be modified. There is a natural topology on the galois group such that intermediate fields correspond bijectively with closed subgroups. (In the finite case the topology is discrete and all subgroups are closed.) Much current research concerns the galois group of the algebraic closure of Q, over Q.

Galois' theorem on solvability of polynomials by radicals

We will prove next that in characteristic zero, a polynomial which is "solvable by radicals" has a solvable galois group, in the sense of having an abelian normal tower. Hence we can give a simple example of a polynomial that is not solvable by radicals.

Lemma: The galois group of a polynomial is isomorphic to a subgroup of permutations of its distinct roots. If the polynomial, is irreducible, the subgroup of permutations is transitive on the roots.

Cor: If an irreducible polynomial over Q has prime degree p, and exactly two non real roots, its Galois group is isomorphic to S(p).

Def: A primitive nth root of 1 (or of "unity"), is an element w of a field such that w^n = 1, but no smaller power of w equals 1.

Lemma: If char(k) = 0, for every n > 0, there is an extension of k which contains a primitive nth root of 1.

Theorem 1: If char(k) = 0, the Galois group G of X^n -1 over k is isomorphic to a subgroup of the multiplicative group (Z/n)*, hence G is abelian.

Rmk: If k = Q, then G ≈ (Z/n)*, as we will show later.

Theorem 2: If k is a field containing a primitive nth root of 1, and c is an element of k, the galois group G of X^n-c is isomorphic to a subgroup of the additive group Z/n, hence G is abelian.
Rmk: It is NOT always true that G equals Z/n, even if k is the splitting field of X^{n-1} over Q.

Theorem 3: If ch(k) = 0, and w is a primitive nth root of 1, and if $k_0 = k(w)$, $k_i = k(w,a_1,...,a_i)$, where for all i = 1,...,m, $a_i^{r_i} = b_i$ is some element of k_{i-1}, and r_i divides n, then the galois group of $k_m = k(w,a_1,...,a_m)$ over k is solvable.

Def: A radical extension E of k, is one obtained by successively adjoining radicals of elements already obtained. I.e. $E = k(a_1,...,a_m)$ where for each i, some positive integral power of a_i lies in the field $k(a_1,...,a_{i-1})$.

Def: A polynomial f in k[X] is "solvable by radicals" if its splitting field lies in some radical extension of k.

Theorem 4: If k has characteristic zero, and f in k[X] is solvable by radicals, then the galois group of f is a solvable group.

Rmk: The converse is true as well, also in characteristic zero.

Cor 5: The polynomial $f = X^5-80X + 2$ in Q[X], is not solvable by radicals.
proof: The derivative $5X^4 - 80$, has two real roots 2,-2, so the graph has two critical points, (-2, 130), and (2, -126). Since f is monic of odd degree, it has thus exactly 3 real roots, and 2 non real roots. The galois group is therefore isomorphic to S(5), which has a non solvable subgroup A(5) ≈ Icos. **QED.**

We show next the galois group of X^n-1 over Q is isomorphic to (Z/n)*. it suffices to prove the following.
Lemma: All primitive nth roots of 1 over Q have the same minimal polynomial.
proof: Define $f_n(X) = \prod(X-w)$ as w ranges over all primitive nth roots of 1 in the complex field. Then $X^n -1 = \prod f_d(X)$ for all factors d of n, since every element of the multiplicative group of roots of X^n-1 has a unique order dividing n; i.e. each nth root of 1 is a primitive dth root for some

unique d diving n. Since X^n-1 is a primitive polynomial and monic with coefficients in Z, and the same holds for f_1 = (X-1), induction proves the same for all d, by Gauss' lemma. Since every primitive nth root of 1 is obtained from one of them by repeatedly raising it to prime powers not dividing n, it suffices for the lemma to show if $f_n(w) = 0$, then $f_n(w^p) = 0$, when p is any prime not dividing n.

If this is false, let f be the minimal polynomial of w, and g that of w^p. If $f \neq g$, then being irreducible f,g have no common roots and fg divides X^n - 1. But g(w^p) = 0, so f(X) divides g(X^p). Mod p then, f(X) divides g(X^p) = (g(X))^p, so mod p every irreducible factor of f is a factor of g. Then fg has repeated factors and repeated roots mod p. But fg is a factor of X^n - 1, which has no repeated roots mod p by the derivative test. I.e. since p does not divide n, the polynomials X^n - 1 and nX^(n-1) have no common roots mod p.

Thus in fact f = g and f(w^p) = 0. Repeating this for all primes not dividing n, we see every power of w relatively prime to n, hence every primitive root of 1, satisfies f. Hence $f = f_n$, which thus is irreducible of order = euler's phi function of n.

To summarize this somewhat tricky proof, if p does not divide n, then X^n - 1 has distinct roots mod p, but if the primitive roots w and w^p have different minimal polynomials f,g, then fg is a factor of X^n - 1, and we can force their roots to come together mod p, contradiction. **QED.**

The Discriminant of a polynomial

Every separable irreducible cubic polynomial has galois group S(3) of order 6, or A(3) of order 3. The difference is whether or not adjoining one root also gives all three roots or not. This is not so easy to tell, but there is a computational method.

Assume we transform the polynomial by translating the roots so that their sum is zero. then we have $f(X) = X^3+pX+q$, where p,q are elements of k. Note that if a,b,c, are the roots of f, then p = -(ab+ac+bc), and q = abc. These functions of a,b,c, are left fixed by the full group S(3) of permutations of the roots. We want to know whether the galois group G is this full group S(3) or some proper subgroup, so it helps to know if G leaves some other less symmetric combinations of a,b,c, fixed as well. Consider this fully symmetric expression $(a-b)^2(a-c)^2(b-c)^2 = D$, the "discriminant" of f. Its square root $\partial = (a-b)(a-c)(b-c)$ changes sign when the roots are permuted by a transposition, hence it is invariant precisely under the alternating group A(3). Hence if ∂ lies in k, and thus is left fixed by G, then G must be equal to A(3), but not otherwise. But how to tell whether ∂ lies in k? There is a simple expression for D in terms of p,q, and then we can check whether D has a square root in k or not.

Namely $D = -4p^3-27q^2$. To see this, one can use various clumsy but constructive proofs of the fundamental theorem of symmetric functions, or the quick trick shown me by Dr Varley: Since D is homogeneous of degree 6 in a,b,c, it has form $D = ap^3 + bq^2$, and we can solve for a,b, by computing D on any two polynomials, like (X-1)(X-1)(X+2), and X(X+1)(X-1).

Examples: X^3 -12X+2, is irreducible by eisenstein, and has discriminant $D = 4(12)^3 - 27(4) = 4(3)^4(21)$ is not a rational square, so the galois group is S(3). X^3-3X+1 is irreducible since it

has no rational roots, and D = 108-27 = 81, a square in Q. Hence the galois group is A(3) ≈ Z/3.

More complicated criteria of this sort exist for quartics as well. More useful general criteria use specialization mod p, for various primes. Next we show all finite groups occur as Galois groups, and all finite abelian groups occur over the base field Q.

Cyclotomic Polynomals

Consider the field C of complex numbers, and the non zero complex numbers of finite multiplicative order. We have shown in the notes on Galois theory that all elements of the same finite order have the same irreducible polynomial F_n over Q. Conversely, any two roots of the same irreducible factor F_n of a polynomial of form X^n - 1, are conjugate under some Q automorphism of their splitting field, hence have the same multiplicative order. It follows that the complex roots of the Q polynomial F_n, are precisely those complex numbers of order n.

We call a complex number of finite order n, a "primitive nth root of 1". Thus the primitive nth roots of 1 form the set of roots of a distinguished Q- irreducible factor F_n of X^n-1. That Q-irreducible polynomial F_n is called the "nth cyclotomic polynomial". Remember that F_n does not have degree n, but F_n has degree phi(n) where phi is the euler function. Since every root of X^n -1 has order dividing n, it follows that X^n -1 factors over Q as the product of the cyclotomic polynomials F_d, for d|n.

For example: X^6 -1 = (X^3-1)(X^3+1) = (X-1)(X^2+X+1)(X+1)(X^2-X+1), so F_1 = X-1, F_2 = X+1, F_3 = X^2+X+1, and F_6 = X^2-X+1, and X^6 -1 = $F_1F_2F_3F_6$. In particular the spliting field of X^6 -1 is generated by the roots of F_6, hence has degree 3.

Lemma: For all n ≥ 1, F_n is monic, has coefficients in Z, and irreducible over Q.
Proof: We know the product of the F_d over all d|n, equals X^n -1, which has content 1. It follows from the multiplicativity of content that the product of the primitive versions of each of these F_d also equals X^n -1. Since X^n -1 is also monic, "the" primitive version of each F_d may also be assumed monic. Since each F_d is also monic, it follows that each F_d is already primitive, and in particular has coefficients in Z. We already know these polynomials are irreducible over Q. **QED.**

Thus one way to obtain F_n is to factor X^n - 1 over Z into irreducible factors, and take for F_n the unique irreducible factor that does not divide any polynomial of form X^d - 1 for any d dividing n. Or if one proceeds recursively, one only need exclude at each stage the polynomials found in previous steps.

What about these same polynomials mod p? Since they have integer coefficients we can reduce them mod p and call them the cyclotomic polynomials mod p, but they may not be irreducible, and their roots may not have order n, i.e. may not be primitive nth roots of 1.

The distinct roots of X^n - 1 mod p are a finite subgroup of a field, hence still form a cyclic group but that group may have order some d dividing n, and less than n. There will exist a

primitive nth root of 1 if and only if that cyclic group has order n, if and only if the polynomial X^n - 1 is separable, if and only if it has distinct roots in the algebraic closure of Z/p.

If a primitive nth root exists, there will always be the same number of them, namely the number of generators of the cyclic group Z/n, equal to the degree of the cyclotomic polynomial F_n, so its roots mod p will still be the primitive nth roots mod p if those exist. But even if they exist, F_n may not be irreducible, so different primitive nth roots may have different minimal polynomials over Z/p. Thus mod p, F_n may be reducible, and its irreducible factors may or may not be irreducible factors of some X^d -1, where d|n.

It follows that there do exist primitive nth roots in the algebraic closure of Z/p if and only if X^n -1 is separable mod p, if and only if p does not divide n. Moreover when p does not divide n, an element of the algebraic closure of Z/p is a primitive nth root if and only if it is a mod p root of the nth cyclotomic polynomial.

Application: Finding primes in the arithmetic series n+1, 2n+1, 3n+1,….

Now when is there a primitive nth root of 1 not just in the algebraic closure of Z/p, but in Z/p itself? This requires that p not divide n, and then that some mod p root of F_n actually live in Z/p. So you need to know F_n, over Z, and then you need there to be an integer a such that p divides $F_n(a)$. Hence if p > n, then p cannot divide n, so then Z/p contains a primitive nth root of 1, if and only if p divides some value $F_n(a)$ of F_n at some integer a.

Thus if p > n, and if p divides some value $F_n(a)$ for some integer a, then Z/p contains a primitive nth root of 1, i.e. an element of multiplicative order, namely the class of a in Z/p. Since the multiplicative subgroup of Z/p has order p-1, n divides n-1.

This lets us study one case of the famous theorem of Dirichlet, namely finding primes congruent to 1, mod n. I.e. given n, to prove there exist primes p such that n divides p-1, we look for primes p > n, and such that F_n has a root in Z/p. Now F_n has a root in Z/p iff p divides $F_n(a)$ for some integer a. Hence we look for primes p greater than n and dividing one of the integers in the sequence:

(*) $F_n(1), F_n(2), F_n(3),...., F_n(s),.....$

Theorem: Given any integer $n \geq 1$, there are infinitely many primes p dividing integers in the sequence (*) above.

I.e. given n, let p belong to the set T iff p is prime and there is a natural number s such that p divides $F_n(s)$. Then T is an infinite set of primes.

Cor: Given any $n \geq 1$, there are infinitely many primes p congruent to 1 mod n. E.g. if n = 10, there are infinitely many primes ending in 1.
proof: The subset of T consisting of primes p > n is also an infinite set

of primes, and for any prime p in this set, we know (Z/p)* contains a primitive nth root of 1. **QED** Corollary.

Proof of theorem: This is a generalization of the proof in Euclid that there are infinitely many prime factors of integers in the sequence of all natural numbers. We assume there are only finitely many and cook up an element x of the sequence which is congruent to 1 modulo all the given primes. Since x has some prime factors, those prime factors are not among the primes already found. The proof works for any integral polynomial of degree ≥ 1, not just the F_n.

Let f be any polynomial of degree ≥ 1, with integer coefficients, and assume that the primes $p_1,...,p_r$ occur as factors of the integers in the sequence f(1), f(2), f(3),..... Since f has degree ≥ 1, it has only a finite number of roots, so let a in Z not be a root, i.e. assume that $f(a) = b \neq 0$. Then consider the integer polynomial $f(a + bp_1p_2....p_rX) = g(X)$. The constant term of g, is $g(0) = f(a) = b$, and by the binomial theorem, also every non constant term of g contains the factor $bp_1p_2....p_rX$, hence is also divisible by b. Thus we can divide g by b to get a polynomial $h(X) = b^{(-1)}g$, with integer coefficients and constant term 1.

Now plug in any integer s, for X, and note that we get one term (the constant term of h) equal to 1, plus a sum of terms (the non constant terms) each containing the factor $bp_1p_2....p_r$,. Hence for all s, the number h(s) is congruent to 1 mod every pj.

Since deg(h) = deg(f) ≥ 1, there are only a finite number of natural numbers s such that $h(s) = \pm 1$, so let s be such that $h(s) \neq \pm 1$. Then h(s) has a prime factor different from any pj. Then also $bh(s) = g(s) = f(a + bp_1p_2....p_rs)$ has a prime factor different from any pj. I.e. this integer value of f has a new prime factor. **QED** theorem.

Application: Every finite abelian group is a Galois group/Q.
Lemma: If G is a finite abelian group, there is a surjective group homomorphism $(Z/n)^* \rightarrow G$, for some integer n, i.e. quotients of the unit groups (Z/n)* exhaust all finite abelian groups.
Proof: We know G is a finite product of cyclic groups of some orders $n_1,....,n_r$. Choose $p_1 > n_1$, such that p_1 divides some value $F_{n1}(a_1)$ for some integer a_1. Then for n_2, choose $p_2 > p_1$, such that p_2 divides some value $F_{n2}(a_2)$,..... etc.
Let $n = \prod p_i$ and consider $(Z/n)^* \approx \prod (Z/p_i)^*$. Since each n_i divides p_i-1, there exist surjections $(Z/p_i)^* \rightarrow Z/n_i$ for every I, and hence a surjection $(Z/n)^* \rightarrow \prod Z/n_i \approx G$. **QED.**

Corollary: Every finite abelian group G occurs as the Galois group of some subfield of a cyclotomic field extension of Q.
Proof: Let G be any finite abelian group. Since (Z/n)* is the Galois group of the cyclotomic splitting field E of $X^n - 1$, if F is the fixed field of the kernel of the surjection $(Z/n)^* \rightarrow G$, then by the fundamental theorem of Galois theory, G is the Galois group of the subfield F of E. **QED.**

It is known as well that finite abelian Galois groups over Q cannot occur in any other way

than as above, i.e. if a finite extension E of Q has abelian Galois group, then E is a subfield of some cyclotomic extension.

It is an open question whether all finite groups occur as Galois groups over Q. Since every finite group G occurs as a subgroup of some permutation group S(n), and it can be shown that these groups do occur, it follows that every finite group G does occur as a Galois group over some finite extension of Q, but not necessarily over Q itself.

We show next the "general" equation of degree n over a field, has Galois group S(n), and hence any finite subgroup H of S(n) occurs as the Galois group of the root field of the general equation over the fixed subfield of H. Hence every finite group can be a Galois group.

Transcendence sets and Galois groups of general equations

Definition: If c is an element of a field extension of k, but is not algebraic over k, we call c "transcendental" over k. This means a non zero polynomial expression in c over k is never equal to zero, i.e. distinct powers of c are k linearly independent.

A collection of elements of a field extension of k are called independent transcendentals over k, if any set of distinct monomials in these elements is k linearly independent, i.e. no non zero k polynomial expression in a finite number of these elements is ever equal to zero.

A collection S of elements of a field extension E of k is a generating transcendence set for E/k if E is algebraic over the intermediate field k(S) generated by the set S over k. A transcendence basis for E/k is a generating set of independent transcendentals for E/k.

A maximal independent set of transcendentals, or a minimal generating set of transcendentals, is a transcendence basis. By Zorn's lemma, transcendence bases exist.

If a set of transcendentals is not independent, some k polynomial in them equals zero. Hence some transcendental is algebraic over the field generated over k by the others, hence that transcendental can be eliminated and the remainder still generate. Thus any finite generating transcendence set can be reduced to a transcendence basis.

Theorem: Any two finite transcendence bases for E/k have the same number of elements, called the transcendence degree of E/k.

Proof: If we precede an ordered set T of m generating transcendentals by an element from a set S of n independent transcendental ones, then the resulting set can be reduced by removing the first element such that the elements up to that one in the ordering are dependent. This is necessarily one of the elements of T, and the remaining set will still generate.

Then we adjoin another element from S independent of the first at the beginning, and again the set becomes dependent. Again we can eliminate the first element such that it, together with the elements preceding it, forms a dependent set, again necessarily an element of T.

Since the set becomes dependent everytime we adjoin an element of S, hence we can eliminate one element of T for each element of S. Thus there must be as many elements of T as elements of S.

Since any finite generating set has as many elements as any independent set, if there is a

finite generating set, then every independent set is finite, in particular there exist finite transcendence bases, and any two transcendence bases have the same number of elements. **QED.**

Example: If k is a field, and X a variable, then {X} is a transcendence basis for the field k(X) of rational functions over k. If X1,…,Xn are independent variables, then {X1,…,Xn} is a transcendence basis for k(X1,…,Xn) over k. Thus variables and transendentals are algebraically the same.

If X1,…,Xn are any variables, the elementary symmetric functions of these variables are: $s1 = \sum Xi$, $s2 = \sum i<j\ XiXj$, $s3 = \sum i<j<k\ XiXjXk$,…., $sn = \prod Xi$.

Cor: If X1,…,Xn are independent transcendentals over k, and if s1,…sn are the elementary symmetric functions of the Xi, then all Xi are algebraic over k(s1,…,sn). Hence the si are also independent transcendentals/k.
Proof: Since $\prod(t-Xi)$ is a polynomial in t of degree n, with each of the Xi as a root, and its coeficients are ± si, the set {si} is a generating set of n transcendentals for k(X1,…,Xn), so the cor. follows from the previous Theorem. **QED.**

Cor: The galois group of k(X1,…,Xn) over k(s1,..,sn) is S(n).

Every permutation of the Xi leaves k(s1,..,sn) fixed. So S(n) is a subgroup of Gal. But the cardinality of gal is at most n! which is a bound for the degree of the splitting field of the general polynomial. Since S(n) has cardinality n!, thus S(n) = Gal.