AN INTRODUCTORY LECTURE ON EULER SYSTEMS

BARRY MAZUR

(these are just some unedited notes I wrote for myself to prepare for my lecture at the Arizona Winter School, 03/02/01)

Preview

Our group will be giving four hour lectures, as the schedule indicates, as follows:

1. Introduction to Euler Systems and Kolyvagin Systems. (B.M.)

2. L-functions and applications of Euler systems to ideal class groups (ascending cyclotomic towers over $\mathbf{Q}$). (T.W.)

3. Student presentation: The "Heegner point" Euler System and applications to the Selmer groups of elliptic curves (ascending anti-cyclotomic towers over quadratic imaginary fields).

4. Student presentation: "Kato's Euler System" and applications to the Selmer groups of elliptic curves (ascending cyclotomic towers over $\mathbf{Q}$).

Here is an "anchor problem" towards which much of the work we are to describe is directed. Fix E an elliptic curve over $\mathbf{Q}$. So, modular. Let K be a number field. We wish to study the "basic arithmetic" of E over K. That is, we want to understand the structure of these objects:

• The *Mordell-Weil group* $E(K)$ of K-rational points on E, and

• The *Shafarevich-Tate group* $\mathrm{Sha}(K, E)$ of isomorphism classes of locally trivial *E-curves* over K.

[By an E-**curve** over K we mean a pair (C, ι) where C is a proper smooth curve defined over K and ι is an isomorphism between the jacobian of C and E, the isomorphism being over K.]

Now experience has led us to realize

1. *(that cohomological methods apply:)* We can use cohomological methods if we study both $E(K)$ and $\mathrm{Sha}(K, E)$ at the same time. That is, for each positive integer n we have *the Classical Selmer group* $S(K, E; \mathbf{Z}/n\mathbf{Z})$ which fits into an exact sequence

$$0 \to E(K)/n \cdot E(K) \to S(K, E; \mathbf{Z}/n\mathbf{Z}) \to \mathrm{Sha}(K, E)[n] \to 0,$$

and the Selmer group is directly expressible in terms of one-dimensional Galois cohomology over K. We will be defining a more general kind of Selmer group in a moment.

2. *(that varying the groundfield sometimes helps:)* There is an advantage to studying Mordell-Weil groups, Shafarevich-Tate groups, and Selmer groups for a large class $\mathcal{L}$ of number fields which are abelian Galois extensions of a "base" number field K (" all at once") rather than for just a single number field K. Here are some standard choices of $\mathcal{L}$:

- (p-*cyclotomic extensions of the rational number field.*) When $K = \mathbf{Q}$ we may take $\mathcal{L}$ to be the class of all abelian extensions of $\mathbf{Q}$; or we may fix a prime number p and take the class of *all* p-abelian extensions of $\mathbf{Q}$, restricting to the p-primary components of the Selmer groups; or (as in classical Iwasawa theory) we might take $\mathcal{L}$ to be the class of all p-abelian extensions of $\mathbf{Q}$ unramified outside p. Much work has been done in studying the asymptotics of Mordell-Weil groups, Shafarevich-Tate groups, and Selmer groups ascending this tower . [**But:** On preparing these notes, I did a double-take when I realized that what I have just written is not really true! See Appendix **A** below.]

- (*anti-cyclotomic extensions of the quadratic imaginary fields.*) When K is a quadratic imaginary field and take $\mathcal{L}$ to be the class of (all, or just those with ramification restricted to the primes dividing p) abelian extensions of K which are Galois extensions of $\mathbf{Q}$ and such that the conjugation action of the nontrivial element of $\mathrm{Gal}(K/\mathbf{Q})$ on their Galois group is via multiplication by -1.

3. *(the principle that you can't get something for nothing:)* There is a powerful method of bounding (from above) the "size" of the Selmer group attached to a given representation over K. This method requires *constructing* certain (*"Euler"*) systems of elements in the Selmer groups attached to the dual representation over each of the number fields in the given "large" class $\mathcal{L}$. The "Heegner" Euler system in the anti-cyclotomic context, and the "Kato" Euler system in the cyclotomic context.

4. *(that L-functions "control" Euler systems which "control" Selmer groups:)* Tom Weston will be explaining this in his lecture, but let me point out that here is where the real power lies. Cohomological methods are pretty good at ferreting out information modulo a single number n, or equivalently modulo powers of prime numbers p but only for finitely many prime numbers p "at a time". But a *single* special value of an L-function can, at times, by its connection to an Euler System, bound from above the size of the relevant p-Selmer groups for *all* (or at least for all but a finite number of) prime numbers p.

Put the flow-chart here.

MENU

Here are the topics to be discussed today. Much of it is already "traditional" and in the literature. See especially Karl Rubin's book:

[R] *Euler Systems*, Annals of Mathematics Studies, Princeton University Press (2000)

To the extent that I will get to any "new" material today, it represents joint work with Karl Rubin.

1. Selmer groups. How playing off local duality against global duality gives a mechanism for bounding Selmer groups.

2. Euler Limits.

3. General bounds.

4. Bounds governed by L-functions

5. The combinatorial rigidity of Kolyvagin systems (see my article with Karl Rubin *Kolyvagin systems* URL:)

1. Selmer groups. How playing off local duality against global duality gives a mechanism for bounding Selmer groups.

If K is a field, $\bar{K}/K$ denotes a choice of separable algebraic closure and $G_K := \mathrm{Gal}(\bar{K}/K)$ its Galois group. Let T be a finite abelian group with continuous G_K action, and $H^*(K,T) := H^*(G_K,T)$ cohomology computed with continuous cochains. Let $T^* := \mathrm{Hom}(T, \mathrm{G}_m) = \mathrm{Hom}(T,\mu)$ be the Cartier dual of T, and

$$T \times T^* \to \mu$$

the duality pairing. This pairing induces a bilinear pairing (via cup-product)

$$H^1(K,T) \times H^1(K,T^*) \to H^2(K,\mu)$$

which looks quite different when we take K to be a global field, or a local field. The mechanism we are about to describe will play one against the other (the global against the local).

Suppose, then, that K is a number field, and K_v is some completion of K (non-archimedean or archimedean). We have the global-to-local restriction mappings

$$H^1(K,T) \longrightarrow \prod_v H^1(K_v,T),$$

(denoting by $\prod_v h_v$ the image of the global cohomology class h)
and

$$H^1(K,T^*) \longrightarrow \prod_v H^1(K_v,T^*).$$

Let us consider how duality enters the story, beginning with the local situation. Let v be a place of K. We have that the cup pairing

$$H^1(K_v,T) \times H^1(K_v,T^*) \to H^2(K_v,\mu) \subset \mathbf{Q}/\mathbf{Z}$$

is a perfect pairing. Hence the cohomology class h_v may be identified with the functional "cupping with h_v", $h_v : H^1(K_v, T^*) \to H^2(K, \mu)$.

In contrast, Global Class Field Theory tells us that if we compose the global (cup-product) pairing

$$H^1(K, T) \times H^1(K, T^*) \to H^2(K, \mu)$$

with the homomorphism $H^2(K, \mu) \to \mathbf{Q}/\mathbf{Z}$ which is given by summing local invariants, we get a beautiful bilinear pairing

$$H^1(K, T) \times H^1(K, T^*) \to \mathbf{Q}/\mathbf{Z}$$

which has the virtue of vanishing identically.

How can we make this disparity work for us?

We wish to impose *local conditions* on the restrictions of global cohomology classes to K_v for places v. To prepare for this let us simply call a **local condition** $\mathcal{F}$ **at** v (for T) any choice of subgroup, which we denote

$$H^1_{\mathcal{F}}(K_v, T) \subset H^1(K_v, T).$$

By the **singular cohomology** for such a local condition $\mathcal{F}$, which we will just denote $H^1_{\mathcal{S}}(K_v, T)$, we mean the quotient group

$$H^1_{\mathcal{S}}(K_v, T) := H^1(K_v, T)/H^1_{\mathcal{F}}(K_v, T).$$

We get then an exact sequence,

$$0 \longrightarrow H^1_{\mathcal{F}}(K_v, T) \longrightarrow H^1(K_v, T) \longrightarrow H^1_{\mathcal{S}}(K_v, T) \longrightarrow 0. \tag{1}$$

Given such a local condition at v, Tate Duality allows us to stipulate a "dual" local condition at v (for T^*), namely, $H^1_{\mathcal{F}^*}(K_v, T^*) \subset H^1(K_v, T^*)$ is defined to be the annihilator subgroup of $H^1_{\mathcal{F}}(K_v, T)$ under the Tate pairing, and the dualization of the above exact sequence yields

$$0 \longrightarrow H^1_{\mathcal{F}^*}(K_v, T^*) \longrightarrow H^1(K_v, T^*) \longrightarrow H^1_{\mathcal{S}*}(K_v, T^*) \longrightarrow 0. \tag{2}$$

The natural choice. If K_v is nonarchimedean let $\mathcal{O}_v$ be the ring of integers in K_v, $\mathbf{F}_v$ its residue field, and $K_v^{unr} \subset \bar{K}_v$ the maximal unramified subfield of $\bar{K}_v$. Let $\mathcal{I}_v$ denote the inertia group $\mathrm{Gal}(\bar{K}_v/K_v^{unr})$, and $G_{\mathbf{F}_v} := \mathrm{Gal}(K_v^{unr}/K_v)$. These groups fit into the exact sequence

$$\{1\} \longrightarrow \mathcal{I}_v \longrightarrow G_{K_v} \longrightarrow G_{\mathbf{F}_v} \longrightarrow \{1\}. \tag{3}$$

Note that if $\bar{\mathbf{F}}_v$ is an algebraic closure of $\mathbf{F}_v$, then $G_{\mathbf{F}_v} \cong \mathrm{Gal}(\bar{\mathbf{F}}_v/\mathbf{F}_v) \cong \hat{\mathbf{Z}}$ (the latter isomorphism sending the Frobenius automorphism in $\mathrm{Gal}(\bar{\mathbf{F}}_v/\mathbf{F}_v)$, $x \mapsto x^{|\mathbf{F}_v|}$, to $1 \in \hat{\mathbf{Z}}$).

The vanishing of $H^2(G_{\mathbf{F}_v}, T^{\mathcal{I}_v})$ yields the canonical exact sequence

$$0 \longrightarrow H^1(G_{\mathbf{F}_v}, T^{\mathcal{I}_v}) \longrightarrow H^1(K_v, T) \longrightarrow H^1(\mathcal{I}_v, T)^{G_{\mathbf{F}_v}} \longrightarrow 0. \tag{4}$$

Now the above exact sequence presents a "natural" choice of local condition for any nonarchimedean v; namely we *could* take $H^1_{\mathcal{F}}(K_v, T)$ to be equal to $H^1(\mathbf{F}_v, T^{I_v}) \subset H^1(K_v, T)$. In this case, $H^1_{\mathcal{S}}(K_v, T) = H^1(\mathcal{I}_v, T)^{G_{\mathbf{F}_v}}$ and If T

is unramified for at v the"natural choices" for T and for T^* are dual under Poitou-Tate duality. In particular, we may identify elements of $H^1_{\mathcal{S}}(K_v, T)$ with linear functionals on $H^1_{\mathcal{F}^*}(K_v, T^*)$.

Let us return to the global situation and say that a **Selmer structure** $\mathcal{F}$ on a (finite) G_K-module T is a local condition $\mathcal{F}_v$ at all v which is the "natural choice" for almost all v. The dual of Selmer structure for T is a Selmer structure for T^*. Note that if we are given a Selmer structure for T the global-to-local mapping $H^1(K,T) \to H^1_{\mathcal{S}_v}(K_v, T)$ vanishes for almost all v, and therefore we have a well-defined global-to-local homomorphism

$$H^1(K,T) \to \bigoplus_v H^1_{\mathcal{S}}(K_v, T).$$

By the **Selmer group** $\mathrm{Sel}_{\mathcal{F}}(K,T)$ associated to $(T, \mathcal{F})$ let us mean the kernel of the above homomorphism. So we have an exact sequence:

$$0 \longrightarrow \mathrm{Sel}_{\mathcal{F}}(K,T) \longrightarrow H^1(K,T) \longrightarrow \bigoplus_v H^1_{\mathcal{S}}(K_v, T). \tag{5}$$

and, dually,

$$0 \longrightarrow \mathrm{Sel}_{\mathcal{F}^*}(K,T^*) \longrightarrow H^1(K,T^*) \longrightarrow \bigoplus_v H^1_{\mathcal{S}^*}(K_v, T). \tag{6}$$

We can now say what the basic mechanism is which allows global cohomology to bound Selmer groups: given any global cohomology class $h \in H^1(K,T)$ consider its image,

$$\bigoplus_v h_{\mathcal{S}_v} \in \bigoplus_v H^1_{\mathcal{S}_v}(K_v, T),$$

and note that since the global duality mapping as displayed above is zero, we get a "semi-local" relation satisfied by any class $\sigma \in \mathrm{Sel}_{\mathcal{F}}(K,T)$. namely,

$$\sum_v h_{\mathcal{S}_v}(\sigma_v) = 0.$$

Given enough of these relations, we can completely describe $\mathrm{Sel}_{\mathcal{F}}(K,T)$. in *good cases.*

2. Euler Limits. We recall here the notion of Euler systems [R] (with some minor modifications). Let R be a complete noetherian local ring with finite residue field of characteristic p. Let K be a number field, $\bar{K}/K$ an algebraic closure, and F be an "intermediate field"; i.e., a field F such that $K \subset F \subset \bar{K}$, so $G_F \subset G_K$ is the subgroup which fixes the elements of F; if F/K is Galois, let $G(F/K) \cong G_K/G_F$ be its Galois group. If M is a compact R-module equipped with continuous G_K-action let $H(F,M) := H^1_{\mathcal{F}}(G_F, M)$, i.e., $H(F,M)$ is the R-module of one-dimensional Galois cohomology of M over F, computed with continuous cochains, and with some chosen Selmer structure $\mathcal{F}$. We assume that $\mathcal{F}$ puts no condition on the primes dividing p. If F/K is Galois, then $H(F,M)$ is naturally an $R[[G(F/K)]]$-module. Note that besides the covariant functoriality of $H(F,M)$ in the pair (F,M) we have a contravariant functoriality given by the corestriction, or norm, mappings $\nu_{F'/F} : H(F',M) \to H(F,M)$ for finite (intermediate) field extensions F'/F. More generally, we might think of allowing $H(F,M)$ in the discussion below to stand for any "decent" functor from pairs (F, M) to the category of R-modules which admits "norms" (i.e., corestrictions).

Let T be a *free* R-module of finite rank equipped with a continuous R-linear G_K-action and which is unramified outside a finite set of primes of K. Let $T^* := Hom_{cont}(T, \mathbf{Z}_p(1))$ denote the dual G_K-module. For each prime q of K fix $\phi_q^{-1} \in G_K$, a choice of Frobenius-inverse element at q. If q is a prime of K unramified in the action of G_K on T^*, form:

$$P_q(X) := det(1 - X\phi_q^{-1}|T^*) \in R[X],$$

the characteristic polynomial of ϕ_q^{-1} acting on the dual module, $T^* := Hom_{cont}(T, \mathbf{Z}_p(1))$. This is well-defined since q is unramified in the action of G_K on T^*.

Fix $\mathcal{N}$, a set of primes containing all ramified primes of T and the primes of K lying above p. If F'/F is an intermediate Galois extension, let $\Sigma(F'/F; \mathcal{N})$ denote the (finite) set of primes of K not in $\mathcal{N}$ which are unramified in the extension F/K and ramified in the extension F'/K. Define:

$$P(F'/F; \mathcal{N}) := \prod_{q \in \Sigma(F'/F; \mathcal{N})} P_q(\phi_q^{-1}) \in R[G_K],$$

noting that this element in the group ring depends upon the choices of Frobenius elements, and on a choice of ordering of the factors in this product. However, since the action of $R[G_K]$ on $H(F, T)$ factors through the quotient ring $R[[G(F/K)]]$, and the primes q contributing to the above product are all unramified in F/K, we see that the natural action of $P(F'/F; \mathcal{N})$ on $H(F, T)$ provides us with a unique R-endomorphism

$$P(F'/F; \mathcal{N}) : H(F, T) \to H(F, T)$$

independent of choice of Frobenius elements if F/K is an abelian Galois extension. Now let L/K be any "intermediate" abelian Galois extension.

Definition. By the $\mathcal{N}$-**Euler (projective) limit** of the system

$$\{H(F, T)\}_{K \subset F \subset L}$$

we mean the $R[[G(L/K)]]$-module of systems of elements

$$\{c_F \in H(F, T)\}_{K \subset F \subset L}$$

satisfying the following compatibility for finite intermediate extensions $F \subset F'$:

$$\nu_{F'/F} \cdot c_{F'} = P(F'/F; \mathcal{N}) \cdot c_F.$$

Provisional notation for this $R[[G(L/K)]]$-module could be

$$\text{E.S.}(L/K, T) := \text{EulerLim}_{F \to L} H(F, T)$$

when the choice of $\mathcal{N}$ is understood. We will refer to this as the $R[[G(L/K)]]$-module of **Euler systems for** $(L/K, T; \mathcal{N})$ noting, however, that the term *Euler systems* is reserved in [R] (Defn. 2.1.1) for the more restricted situation where L contains all the ray class fields over K relative to primes not dividing $\mathcal{N}$ and it contains a $\mathbf{Z}_p$-extension in which no (finite) prime of K splits completely.

Comments. If L/K is unramified outside $\mathcal{N}$, then the "Euler limit" is just the standard inverse limit compiled via norms. In particular, this is the case if L/K

is a $\mathbf{Z}_p$ extension. A nontrivial element in $\mathrm{E.S.}(L/K,T)$ corresponds to a large number of cohomology classes all compatible in this Euler limit way, something of a *hyper-* universal norm! We will restrict attention, below, to p-abelian extensions L/K. [**Note:** For some remarks and queries about universal norms, see Appendix **B** below.]

In our present case, let Γ denote the quotient of the compact p-abelian group $G(L/K)$ by its torsion subgroup. Let K_∞/K be the fixed subextension of L/K under the torsion subgroup of $G(L/K)$. Then we have a natural surjection $G(L/K) \to G(K_\infty/K)$. Putting $\Gamma := G(K_\infty/K)$ we have $G(K_\infty/K) \cong \mathbf{Z}_\nu^d$ for some non-negative integer ν (which we can call the $\mathbf{Z}_p$-rank of L/K). Put $\Lambda := R[[\Gamma]]$.

Let us denote:

$$H_\infty(F,T) := \text{proj.lim.}_{F\to K_\infty} H(F,T).$$

We have the natural homomorphism of Λ-modules,

$$\mathrm{E.S.}(L/K,T)\otimes_{R[[G(L/K)]]}\Lambda \longrightarrow \mathrm{E.S.}(K_\infty/K,T) = H_\infty(F,T).$$

Example. Let R, T be as above with $K=\mathbf{Q}$ and take L/K to be the maximal p-abelian extension of $\mathbf{Q}$. Let $\mathcal{N}$ denote the set containing the ramified primes for T and the prime number p. Let $\mathbf{Q}(\infty)/\mathbf{Q}$ be the (cyclotomic) $\mathbf{Z}_p$-extension and $\mathbf{Q}(n)\subset\mathbf{Q}(\infty)$ the subfield of degree p^n over $\mathbf{Q}$. It follows from the "weak Leopoldt Conjecture" that

$$\mathrm{rank}_\Lambda H_\infty(\mathbf{Q},T) = d^-,$$

where d^- is the dimension of the minus eigenspace of the complex conjugation involution acting on T. Also very reasonable hypotheses guarantee that $H_\infty(\mathbf{Q},T)$ has no Λ-torsion.

3. General Bounds.

Here are two "hypotheses" and it would be very good to establish them quite generally.

Hypothesis A. The kernel of

$$\mathrm{E.S.}(L/\mathbf{Q},T)\otimes_{R[[G(L/K)]]}\Lambda \to H_\infty(\mathbf{Q},T)$$

is a Λ-torsion module.

Hypothesis B. If $d^-=1$ the characteristic ideal of the cokernel of

$$\mathrm{E.S.}(L/\mathbf{Q},T)\otimes_{R[[G(L/K)]]}\Lambda \to H_\infty(\mathbf{Q},T)$$

is equal to the characteristic ideal of the Selmer group of the (Cartier) dual Galois representation T^* with dual Selmer structure $\mathcal{F}^*$.

As for **B.** one has that under very general hypotheses the characteristic ideal of the Selmer group of the (Cartier) dual Galois representation T^* with dual Selmer structure $\mathcal{F}^*$ *divides* the characteristic ideal of the cokernel of

$$\mathrm{E.S.}(L/\mathbf{Q},T)\otimes_{R[[G(L/K)]]}\Lambda \to H_\infty(\mathbf{Q},T).$$

But we can establish the full strength of **B** at present only in very few instances. (For example, when $T = \mathbf{Z}_p(1) \otimes \chi$ where χ an even nontrivial character of finite order. **Query:** Can one show hypothesis **A** in this case?)

4. Bounds governed by L-functions

The example for which we have the most complete information, and which might serve as a template for what we might try to get in other cases is given by taking $K = \mathbf{Q}$ and $T = \mathbf{Z}_p(1) \otimes \chi$ where χ an even nontrivial character of finite order. Modify the Selmer structure on T by putting the *natural* local condition at p; and form

$$H_{\infty,\mathcal{S}}(\mathbf{Q}_p, T) := \text{proj.Lim}_{n\to\infty} H_{\mathcal{S}}(\mathbf{Q}_p(n), T).$$

Either of the two standard proofs of the classical *main conjecture* establishes the fact that the Λ module $\text{E.S.}(L/\mathbf{Q}, T) \otimes_{R[[G(L/K)]]} \Lambda$ is generated by the *cyclotomic Euler system* and the characteristic ideal of the cokernel

$$\text{E.S.}(L/\mathbf{Q}, T) \otimes_{R[[G(L/K)]]} \Lambda \longrightarrow H_{\infty,\mathcal{S}}(\mathbf{Q}, T),$$

is generated by the Leopoldt-Kubota L-function $L_p(\chi, s)$ viewed as element of Λ, which is the characteristic ideal of the Iwasawa module constructed from of p-primary components of ideal class groups of layers of the p-cyclotomic tower.

In the case of elliptic curves E defined over $\mathbf{Q}$, working with the Heegner Euler System over an anti-cyclotomic tower over a quadratic imaginary field, or using Kato's Euler System over the cyclotomic tower over $\mathbf{Q}$ one presently has *divisibility results* (i.e., the characteristic ideal of the appropriate Selmer group *divides* the ideal generated by corresponding p-adic L function (cf. Chapter 3 of [R]), and this alone is enough to establish striking information about Mordell-Weil and Sha, as the later lectures will explain, but in both cases we still await a general "main conjecture."

5. The combinatorial rigidity of Kolyvagin systems.

One of Kolyvagin's many original insights is to make "maximal" use of the Euler Systems of cohomology (these cohomology classes exist in various field extensions of the base) by astutely *descending* these classes to get cohomology classes over the base field, so as to be able to apply the duality methods of section **1.** The collection of classes one gets over the base have a tight structure (see my article with Karl Rubin "Kolyvagin Systems"). If, for example, one's ring of scalars $R = k$ is a finite field, and the Galois representation T satisfies certain reasonably general hypotheses, we show that the system of cohomology classes can be thought of as a section of a linear system of one-dimensional k-vector spaces over a certain connected subgraph of the multiplicative graph of natural numbers. It follows (in this situation) that just by the combinatorial constraints that these Kolyvagin cohomology classes satisfy, Kolyvagin systems, if they exist, are uniquely determined up to normalization. Moreover, Ben Howard has recently demonstrated that the linear systems in question have no monodromy (i.e., are constant) and therefore that the Kolyvagin systems of cohomology classes do exist, irrespective of whether a corresponding Euler system exists.

To give a flavor of the combinatorial nature of these Kolyvagin systems, let me illustrate what it boils down to in the case of an elliptic curve E over $\mathbf{Q}$. For simplicity, make the
Irrelevant hypothesis: $L(E,1) \neq 0$ and $E(\mathbf{Q})$ is a finite group of order relatively prime to p.

Remark. $L(E,1) \neq 0$ implies that $E(\mathbf{Q})$ is a finite group. I am making the hypothesis above only because it simplifes some of the terminology and statements of the propositions below. We begin by discussing E-curves, i.e., torsors over E, these being given by cohomology classes in $H^1(G_{\mathbf{Q}}, E)$. An E-curve is **split** at a prime ℓ (resp., at "infinity") if it has a point over $\mathbf{Q}_\ell$ (resp., over $\mathbf{R}$). The Shafarevich-Tate group of E, Sha(E), is nothing more than the group of $\mathbf{Q}$-isomorphism classes of everywhere split E-curves. We let Sha$'(E)$ denote the group of E-curves which are split at all places different from p. The **order** of an E-**curve** is its order as a cohomology class. If C is an E-curve of order p^ν, then C is (thanks to our irrelevant assumption) also given by a unique element in $c = c(C) \in H^1(G_{\mathbf{Q}}, E[p^\nu])$. For a prime number ℓ dividing $p^\nu - 1$, say that an E-curve C is **transverse** at ℓ if its cohomology class $c(C) \in H^1(G_{\mathbf{Q}}, E[p^\nu])$ goes to zero under the natural homomorphism

$$H^1(G_{\mathbf{Q}}, E[p^\nu]) \to H^1(G_{\mathbf{Q}_\ell(\mu_\ell)}, E[p^\nu]).$$

The condition of transversality is *stronger* than the requirement that the E-curve C is split over $\mathbf{Q}_\ell(\mu_\ell)$ but *weaker* than the requirement that C is split over $\mathbf{Q}_\ell$. Let $\mathcal{N}_\nu$ be the set of positive squarefree integers n which are divisible only by primes congruent to 1 mod p^ν. For $n \in \mathcal{N}_\nu$ let $H(n; \mathbf{Z}/p^\nu\mathbf{Z})$ denote the $\mathbf{Z}/p^\nu\mathbf{Z}$-module of E-curves which are transverse at all divisors of n and split at all primes not dividing pn. Thus (given our *irrelevant assumption*) $H(1; \mathbf{Z}/p^\nu\mathbf{Z}) = \mathrm{Sha}'(E)[p^\nu]$, the kernel of multiplication by p^ν in Sha$(E)'$.

Fix $\nu \geq 1$. Form the graph X_ν whose vertices are the integers $n \in \mathcal{N}_\nu$ and whose edges are in one:one correspondence with pairs of vertices $n, n\ell$ in $\mathcal{N}_\nu$, these vertices being its endpoints.

Definition: A **simplicial sheaf** on X_ν is ....

To illustrate things let us restrict attention to $\nu = 1$. One constructs a canonical sheaf $\mathcal{S}$ on X whose stalk at a vertex n is given by $\mathcal{S}(n) := H(n; \mathbf{Z}/p\mathbf{Z}) \otimes W(n)$, where $W(n)$ is the $\mathbf{F}_p$ vector space of dimension one given by $W(n) := \otimes_{\ell \mid n} \mathbf{F}_\ell^* \otimes \mathbf{F}_p$. One then restricts this sheaf to the subgraph $X' \subset X$ on which each stalk is of dimension one. A Kolyvagin system in this context is simply a trivialization of this subsheaf.

Appendix A. The ℓ-asymptotics of Sha as one ascends a p-cyclotomic tower. The p-adic "main conjecture" for elliptic curves packages much that one might want to understand about p-asymptotics of Sha as one ascends a p-cyclotomic tower, but I have never heard, or read, any mention of the perfectly natural companion question alluded to in the title of this section, when $\ell \neq p$. The natural guess here, is to follow the lead of Larry Washington's 1978 Inventiones article where he proves that if $\ell \neq p$, k is any abelian number field, k_n/k the n-th layer of the p-cyclotomic $\mathbf{Z_p}$-extension of k, $(n = 1, 2, \dots)$, and ℓ^{e_n} the exact power of ℓ dividing the class number of k_n. then e_n is constant for n sufficiently large. If we allow ourselves to be influenced by that, and by the "standard analogy" between ideal class groups and Sha, a first guess might be that if $\ell \neq p$, and E is an elliptic curve over k such that the G_k-representation on $E[\ell]$ is absolutely irreducible, then the order of the ℓ-primary component of $\mathrm{Sha}(E/k_n)$ is constant for n sufficiently large. It would be even more interesting if there were counter-examples to this first guess. Using modular symbols, how hard would it be to get data on this? I also wonder whether people have considered the analogous questions for arithmetic K-groups.

Appendix B. "Pure" Λ-modules and Universal norms.

Pure Λ-modules. Let

$$\Lambda_n := \mathbf{Z}_p[\mathbf{Z}/p^n\mathbf{Z}_p] = \mathbf{Z}_p[\xi]/(1-\xi^{p^n})$$

and $\Lambda = \mathbf{Z}_p[[\mathbf{Z}_p]] = \text{proj.lim.}_{n\to\infty}\Lambda_n$. We denote by $\xi \in \Lambda$ the element that projects to the ξ's in all the Λ_n's. If M is a Λ-module, put

$$M_n := M \otimes_\Lambda /\Lambda_n = M/(1-\xi^{p^n})M.$$

If M is finitely generated as Λ-module, say that M is **pure** if M_n is a free $\mathbf{Z}_p$ module for all $n \geq 1$. If W is a $\mathbf{Z}_p$ module, let $\bar{W}$ be the torsionfree quotient of W; i.e., it is the quotient of W by its torsion submodule. If M is a finitely generated Λ-module, its **pure quotient** is the Λ-module $\tilde{M} := \text{proj.lim.}_{n\to\infty}\bar{M}_n$. Since sub-$\mathbf{Z}_p$ modules of torsionfree $\mathbf{Z}_p$ modules are again free, it follows that $\tilde{M}$ is indeed pure. A finitely generated Λ-module M is "Γ-finite" (in Iwasawa's terminology) if and only if its pure quotient is trivial. For any finitely generated Λ-module M we have a canonical exact sequence of Λ-modules

$$0 \to M^f \to M \to \tilde{M} \to 0,$$

where M^f is the maximal Γ-finite Λ-submodule of M, and $\tilde{M}$ is its pure quotient. A pure Λ-module is isomorphic modulo the class $\mathcal{C}$ of finite modules (i.e., is *pseudo-isomorphic* to the direct sum of a free Λ-module and a module on which the action of ξ is of finite order. If W is a module over an integral domain A, by the A-**rank** of M, denoted $r_A(M)$, we mean the dimension over $\mathrm{Frac}(A)$, the field of fractions of A, of the vector space $M \otimes_A \mathrm{Frac}(A)$. Any finitely generated Λ-module M has the property that

$$r_{\mathbf{Z}_p}(M_n) = r_{\mathbf{Z}_p}(\tilde{M}_n) = r_\Lambda(M)\cdot p^n + \text{constant}$$

for n sufficiently large.

By a **co-pure** Λ-module (of cofinite type) let us mean a Λ module W with the property that if W_n is the submodule of W consisting of elements which are fixed by ξ^{p^n}, then the W_n's are all free modules (of finite rank) over $\mathbf{Z}_p$ and $W = \bigcup_{n=1}^{\infty} W_n$. We have a nice $\mathbf{Z}_p$ duality theory between pure and co-pure Λ-modules as follows. If W is co-pure put $M_n := Hom_{\mathbf{Z}_p}(W_n, \mathbf{Z}_p)$; then $M := \text{proj.lim.}_{n\to\infty}M_n$ is pure and we will refer to it as the $\mathbf{Z}_p$-dual of W, denoting it $W^* := M$. If M is any Λ module of finite type, then, put $W_n := Hom_{\mathbf{Z}_p}(M_n, \mathbf{Z}_p)$; then $W := \text{ind.lim.}_{n\to\infty}W_n$ is co-pure and we will refer to it as the $\mathbf{Z}_p$-dual of M, denoting it $M^* := W$. The $\mathbf{Z}_p$-dual of M may be identified with the $\mathbf{Z}_p$-dual of its pure quotient, $\tilde{M}$, and if M is of finite type, we have a canonical identification of its pure quotient with its double

$\mathbf{Z}_p$-dual. The operation of $\mathbf{Z}_p$-duality preserves pseudo-isomorphisms, and finite direct sums.

Universal norms. For integers $0 \leq n \leq m$ consider the "norm element" $\nu_{m,n} \in \mathbf{Z}_p[\xi]/(1-\xi^{p^m})$ defined by the formula

$$\nu_{m,n} := \sum_{\alpha=0}^{p^{m-n}-1} \xi^{\alpha p^n}.$$

If W is a Λ-module, note that $\nu_{m,n}W_m \subset W_n \subset W_m$. Define the Λ-module of **universal norms** of W, denoted $UN(W)$ by setting

$$UN(W)^{(n)} := \cap_{m\to\infty}\nu_{m,n}W_m$$

for each $n \geq 0$ and putting $UN(W) := \text{proj.lim.}_{n\to\infty} UN(W)^{(n)}$. The operation UN preserves pseudo-isomorphisms, and finite direct sums.

Note that if W is a co-pure Λ-module of cofinite type, then $UN(W)$ is a pure Λ-module of finite type. Let us start now with a Λ-module M of finite type and form $UN(M^*)$, the module of universal norms of its $\mathbf{Z}_p$-dual. This operation $M \mapsto UN(M^*)$ factors through $M \to \tilde{M}$ and preserves pseudo-isomorphisms, and finite direct sums. Let us analyze this operation in two cases.

• Let $M = \Lambda$ viewed as (free, rank 1) Λ-module. So $M_n = \Lambda_n$, $M_n^* = \mathrm{Hom}_{\mathbf{Z}_p}(\Lambda_n, \mathbf{Z}_p)$, and one sees that for all $0 \leq n \leq m$, $\nu_{m,n} : M_m^* \to M_n^*$ is surjective, and therefore $UN(M^*)^{(n)} = M_n^*$. We may have an isomorphism $\iota_n : \Lambda_n \cong \mathrm{Hom}_{\mathbf{Z}_p}(\Lambda_n, \mathbf{Z}_p)$ as Λ -modules given by sending the element $\sum_{j=0}^{p^n-1}\lambda^j\xi^j$ to the $\mathbf{Z}_p$-homomorphism that takes the value λ^j on ξ^j (for $j = 0, \ldots, p^n-1$). The ι_n's are compatible with norms, in the sense that $\pi_{m,n} \cdot \iota_n = \iota_m \cdot \nu_{m,n}$, where $\pi_{m,n} : \Lambda_m \to \Lambda_n$ is the natural projection. It follows that $UN(M^*)$ is free over Λ of rank 1, that $UN(M^*)^{(n)} = UN(M^*)_n$ $(= M_n^*)$ (where the lower index n of a *lambda*-module is defined as in section 1).

• Let M be a Λ on which the action of ξ is of finite order. Here it is evident that $UN(M^*) = 0$.

Proposition. Let M be any Λ-module of finite type. Then $UN(M^*)$ is a Λ-torsionfree module of Λ-rank equal to the Λ-rank of M. Moreover the mappings

$$UN(M^*)_n \to M_n^*$$

are injective for all $n \geq 0$.

Proof. Since the passage $M \mapsto M^*$ factors through pure quotients, we may suppose that M is pure. Hence, up to pseudo-isomorphism M is a finite direct sum of a free Λ-module and one on which the action of ξ is of finite order. Since our operation $M \mapsto UN(M^*)$ preserves pseudo-isomorphisms, and finite direct sums, the analysis we have already made proves our proposition.

Remark. In particular, we have a canonical $\mathbf{Q}_p$ vector subspace of $M_1^* \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$ of dimension equal to the Λ-rank of M given by the image of $UN(M^*)_1^* \otimes_{\mathbf{Z}_p} \mathbf{Q}_p$ (call it the *universal norm subspace*). Now the recent work of Vatsal, Cornut, Bertolini (also Zhang) concerning the p-adic pro-Selmer group relative to the p-anticyclotomic $\mathbf{Z}_p$-extension over quadratic imaginary fields K an elliptic curve E over $\mathbf{Q}$ establishes the fact that (for primes p of good, ordinary reduction for E, and for quadratic imaginary fields K satisfying the appropriate splitting properties for primes dividing the conductor of E) this Selmer group is a free Λ-module of rank 1. I will also assume, unnecessarily surely, that E doesn't contain nontrivial K-rational points of order p. We have that the pro-p Selmer group, i.e., the one built from Galois cohomology of the Tate module T_pE, (for the p-anticyclotomic $\mathbf{Z}_p$-tower over K) is co-pure of co-finite type, and is the $\mathbf{Z}_p$-dual of a module of Λ-rank equal to 1. It follows, assuming finiteness of the p-primary component of $\mathrm{Sha}(E;K)$, that the universal norm subspace as described above is a $\mathbf{Q}_p$ subspace of dimension one (a *line*) in $E(K)\otimes_{\mathbf{Z}_p}\mathbf{Q}_p$. It is immediate that this line is in the null-space of the p-adic height pairing

$$E(K)\otimes_{\mathbf{Z}_p}\mathbf{Q}_p\times E(K)\otimes_{\mathbf{Z}_p}\mathbf{Q}_p\to\mathbf{Q}_p$$

attached to the p-anticyclotomic $\mathbf{Z}_p$-extension of K. What more can one say about this line? The $\mathbf{Q}$ vector space $E(K)\otimes_{\mathbf{Z}}\mathbf{Q}$ has odd dimension (denote it ρ), given the conditions regarding K that we alluded to above but didn't write down. The vector space breaks up into the "plus" and "minus" eigenspaces for the action of complex conjugation on K, whose unequal dimensions we denote $\rho^{\pm}$, so that $\rho=\rho^{+}+\rho^{-}$. It is natural to guess that the null-space of the p-adic (anti-cyclotomic) height pairing is in the eigenspace for complex conjugation which has the larger of the two ranks. Therefore, if this guess were true, it would follows from this that the universal norm line would lie in $E(K)\otimes_{\mathbf{Z}_p}\mathbf{Q}_p^{\pm}$ where the sign $\pm$ is given by whichever of $\rho^{\pm}$ is the larger. Assuming our guess, one might wonder about the "placement" of the universal norm line, defined over $\mathbf{Q}_p$, in the $\mathbf{Q}$-vector space $E(K)\otimes_{\mathbf{Z}}\mathbf{Q}^{\pm}$. One would expect it to be as transcendental as possible, barring any further ideas about how it might be constrained ...

Department of Mathematics, Harvard University, Cambridge, MA 02138 USA
E-mail address: `mazur@math.harvard.edu`