

Modular Design and the Development of Complex Artifacts: Lessons from Free/Open Source Software

Alessandro Narduzzo
Free University of Bozen-Bolzano
Bozen, Italy
narduz@unibz.it

Alessandro Rossi
University of Trento
Trento, Italy
arossi@cs.unitn.it

Abstract – Organizational and managerial theories of modularity applied to the design and production of complex artifacts are used to interpret the rise and success of development methodologies and practices in Free/Open Source software projects. Strengths and risks of the adoption of a modular approach in software project management are introduced and are related to the achievements of various Free/Open Source Software projects (among them: the GNU operating system, the Linux kernel, the HURD kernel). It is suggested that mindful implementation of the principles of modularity may improve the rate of success of many Free/Open Source software projects. Specific case studies here depicted, as well as indirect observation of common programming practices employed by Free/Open Source developers and users, suggest a possible revision towards an improved theory of modularity that may be extended also to settings different from software production.

practical application to domains different from software production.

II. MODULARITY

I. INTRODUCTION

The popularity of the GNU/Linux operating system has conveyed increasing attention to the Free/Open Source Software (F/OSS) development model, usually described as a radically different system of rules, practices, and methodologies, shared within a large and virtually distributed software developers community, alternative to proprietary and closed development techniques employed by traditional hierarchical organizations in the software industry [1, 2, 3].

A growing number of studies has analyzed the F/OSS movement, from a variety of perspectives, aimed at understanding the bases of such successful phenomenon.¹ Our study offers a complementary analysis of the design and the development of F/OSS in terms of a theory of modularity [4]. Our reconsideration of the accounts of GNU/Linux and other F/OSS projects highlights how they benefited from the typical advantages of implementing modular architectures (e.g. fast speed of development, recombination of modules, innovation through projects competition, reuse of previously developed code [3, 5, 6]), while, at the same time, many critical pitfalls typically related to managing modularity (i.e. the architectural design of modules and interfaces) were avoided.

Our account on how the principles of modular design have been originally adapted by F/OSS development allows us to move away from the stereotypical definition of modularity. As a matter of fact, while it is more and more often proposed as a fundamental paradigm for the design and production of artifacts [7], it is still regarded by some authors as a black box [8, 9]. The empirical analysis of F/OSS projects allows to shed some light on many critical issues related to modularity.

We suggest how the peculiar implementation of the principles of modularity shown by F/OSS may help in refining both existing theories of modularity, and their

Modularity has been receiving an increasing amount of attention in a variety of fields, from neuroscience and artificial intelligence to architecture, urban design and management [7]. This interdisciplinary interest is largely due to the fact that modularity is regarded as a general property of complex systems, pertaining to the degree of decomposability of a system in loosely coupled sub-parts made by tightly coupled components.

Modularity provides relevant advantages that have been neatly identified in the literature [10]. Modularity allows for product variety that is obtained by a recombination (mix and match) of components [11]. In addition, it is viewed as a base for differentiation strategies: firms may enrich their products catalog and adapt to customers' needs with limited additional costs [12]. Finally, modularity has also a great impact on production processes as it positively affects flexibility, division of labor and specialization [9].

According to Baldwin and Clark modularity in production systems is obtained by following some general rules, originally drawn from computer science and software development, concerning two different categories of information: visible and hidden information [7]. Modular systems design needs to specify only the visible rules, namely the information about: (a) the definition of the architecture, (b) interfaces specifications and (c) modules integration tests. The inner description of each module and how it works are hidden from outside: it does not need to be defined ex-ante or communicated during the process, since modules interactions exclusively follow the rules specified by the interfaces parameters.

Herbert Simon's influence the modularity literature is particularly evident. First of all, modularity is introduced within a problem solving framework and modular design is regarded as a solution to cope with uncertainty and variability. Second, as in Simon's analysis of the artificial, modularity in complex systems regards both goals and hierarchies. Third, modular solutions are based on problem decomposition; fourth, since complex systems are not quite entirely decomposable, modular design eventually needs to deal with residual interdependencies [13].

A. The "power of modularity" in software engineering

Since the early days of software engineering the issue of designing, developing, testing and releasing a large software project brought into discussion the trade-off between simplicity and speed of development.

Frederick Brooks, clearly recognized that small sharp teams performed better than large ones, but they were not

¹ See for instance the *Research Policy* special issue Vol. 32, N° 7, 2003.

sufficiently staffed to deliver large software projects under schedule pressure [14]. Efficient software engineering methodologies are meant to solve this fundamental trade-off between task partition and division of labor, on the one hand, and coordination and communication costs, on the other one. Brooks' recipe for coping with the design and the production of complex software was to vertically divide labor in order to separate high-level activities as much as possible (such as the design of a software artifact) from lower ones (such as the implementation of code). As a result, even a large software project might have been guided by a small number of architectural designers, hence reducing coordination and communication costs needed to conceive the architectural blueprint of the project. A second related element in Brook's recipe was then to assign the implementation of each part of the project to small and focused teams (the so called "surgical team").

In terms of a modern theory of modularity, the basic assumption inside Brook's seminal work is that large software projects are integral and non-decomposable systems, where interactions among parts are nontrivial and generate high communication and coordination needs. Vertical division of labor is viewed as the way to avoid as much as possible these inefficiencies by concentrating design and architectural activities on few heads. What is clearly overlooked from Brook's perspective is that interdependencies may not only be considered as given constraints, but rather they may be strongly reduced at the architectural design level, by effectively decomposing the complex system in quasi-independent subparts.

As a matter of fact, the introduction of a fully modular approach in modern software engineering methodologies has been fostered by the recognition that the degree of interdependencies may be strongly reduced if a complex software project can be decomposed in independent subparts, that is, dividing the whole project in smaller components that are loosely coupled and highly independent on each other [10, 15]. Hence, when subparts are almost independent, it is possible to divide labor minimizing the risk of coordination failures. Conceiving the design of a complex software artifact as a modular system means to apply the basic principle of "information hiding", originally developed by Parnas [16], that prescribes treating software modules as opaque entities, whose relevant information is only available to its inner programmer, while not being accessible to external programmers. Here the only information revealed is embedded in the interfaces, while the information regarding the design and how the module works is not communicated.

The above discussion seems to push modularity and information hiding as the landmark principles for combining concerns of size and division of labor with high speed of development of a software project. Nevertheless, the "dark side" of modularity, namely the pitfalls of system integration and testing of modular design, seems to be particularly substantial in the field of software artifacts.

B. Modularity as a complex design activity: managing unforeseen interdependencies in software modules

Brooks' famous essay on the difficulties of software engineering techniques in granting improvements in productivity, reliability and simplicity in developing

software programs, may support us in refining our explanations of why integrating software modules and thus producing modular software may be difficult [17]. The author speculates on the fundamental properties of software entities that may account for the difficulties in separating interdependencies and decompose large software projects: software entities differ from physical artifacts for their highly nonlinear complexity, leading to the impossibility of enumerating (not to mention understanding) all the possible states of a program. As the size of a software project increases, it becomes more and more difficult to decompose interdependencies and to design an architecture that preserves the initial conceptual integrity of the software project by a combination of loosely coupled functional software components. Moreover, software is invisible. The same intangible attributes that seem to free software entities from standard physical constraints that hardware ones have to satisfy, seem at the same time to affect human abilities of anticipating correctly component interface specifications and interdependencies. While geometric abstraction are powerful tools that may help the architectural design for assembly goods ("Contradictions become obvious, omissions can be caught." [17]), similar geometrical representations do not help much during the design phase of software structures because source of interdependencies are more subtle, not visible, and related to a series of elements ("flow of control, flow of data, patterns of dependency, time sequence, name-space relationships" [17]) whose interrelations may be only partially caught by diagrams and flow charts.

As a consequence, in software production, testing, debugging and integration phases may be much more relevant in terms of resources needed when compared to the production of physical artifacts. This is largely due to two intertwined aspects: (a) designers are boundedly rational decision makers [18] and (b) the nature of interdependencies between modules is mainly multidimensional and invisible. As a result, the act of decomposing a large software project into components is an activity that results, at its best, in a suboptimal outcome: some sources of interdependencies are well determined and taken into account in the design of components and interfaces, while others are not. In some sense, even careful decomposition of large software projects tends to be accomplished making trade-offs between sources of interdependencies, recognizing the more visible ones and disregarding the less evident or less important ones. Likewise, less careful decomposition results in even greater problems at the final stages of code integration.

III. MODULARITY IN F/OSS DEVELOPMENT

A. Imitating a previously existing design

Modular design of complex systems is a demanding job since all the modules interfaces have to be defined ex-ante. How can designers cope with such degrees of complexities? One of the lessons coming from the accounts of some well-known F/OSS projects is to take advantage of existing templates, rather than to develop a brand new project from scratch. The Free Software Foundation (FSF) GNU project and the FreeBSD

operating system project are two relevant instances of this approach to the modular design of complex artifacts, as their kinship with UNIX operating system is openly recognized.

UNIX operating system was a milestone in the computer software history and it is usually described as a highly modular, scalable and portable platform [19, 20]. The UNIX architecture is a complex and massively decomposed architecture of independent modules, characterized by high specialization of programs (“programs that do one thing and do it well”), working together by means of structures, (“pipes”), and sharing as a fundamental interface of communication text streams [21]. UNIX was the first modern operating system not developed using a hardware dependent assembly language. The kernel was written in C, ensuring portability to various hardware platforms [22, 23]. UNIX highly modular architecture had strong consequences both at the level of developers coding activities and at the level of users’ experience. Developers were able, thanks to its modular design, to carry out development of specific parts of the system in autonomy and without any need to coordinate their efforts with other sub-projects. Modularity allowed for both parallel development and contribution of new components; furthermore, the overall design of the system was significantly improved by the development of innovative modules and competition between similar projects [7]. At the end-user level, modularity invited mere users to employ mix and match strategies (recombination of different modules), allowing them to generate a wide variety of different implementations of the operating system where a large part of the modules pertaining to the user space were highly customizable and were chosen according to specific tastes or needs.

Even through this rather short and incomplete account of the early days of UNIX hackerdom, the past arguments should suggest that many of the elements pertaining to the decentralized and spontaneous nature of Linux development process are not as innovative and original as many Linux advocates often underline. They are rather mostly inherited from Linux direct ancestor, the UNIX operating system. Strangely enough, this almost self evident argument seems to be mysteriously overlooked in many popular contributions to the Linux debate.

The GNU project, started in 1984 by Richard Stallman, represented at its beginning a titanic effort to offer a free alternative to currently existing commercial and proprietary operating systems. In this respect, Stallman’s design strategy consisted in cloning an already existing project, a stable and mature architecture that had been originally conceived around fifteen years before [24].

As a matter of fact, the whole GNU project represented the attempt to recreate the pre-AT&T UNIX arcadic era, where the original architecture was preserved in essence and only some limited and marginal reworking in the design took place, in order to solve some minor technical disadvantages of UNIX (e.g. the introduction of 32-bit support). This architectural choice followed by Stallman, and later widely adopted by the hacker community, has been a conservative one. A more risky option such as undertaking a radically innovative project based on the design of a new architecture was disregarded in favor of a safe and well known alternative.

FreeBSD is another important operating system that deliberately mimicked the architecture of the UNIX operative system. Again, by adopting an existing architecture, the community spent its attention on incremental development, rather than on design discussions [25]. To conceive a new operating system characterized by a modular architecture is a challenging cognitive activity of modules and interfaces definition. First, the designer needs to conceive a system of modules, by decomposing the whole system in quasi-independent components. Second, failures in the decomposition phase results in extra costs for fine-tuning and fixing activities aimed at solving unexpected and unforeseen interdependencies.

In this respect, the FSF and the FreeBSD community were able to consciously handle what, through the lens of the theory of modularity, is a fundamental trade-off between threats at the design level and opportunities at the implementation level. As a result, the decision of establishing the GNU and the FreeBSD projects upon a stable, mature and carefully modularized architecture was the key element to benefit from the typical advantages of modularity (concurrent engineering, division of labor, decentralized development, innovation via module based evolutionary dynamics, and much more), while at the same time avoiding the classic pitfalls and drawbacks of modularity, concerning the risks of imperfect decomposition in the design of an innovative modular architecture as the backbone for the project.

B. Horizontal division of labor, task interdependencies and Brooks’ Law

The perspective of modularity seems also to offer other different interpretations contrasting many other recurring stereotypes in the debate over the revolutionary nature of F/OSS development.

One of the most criticized principles of the otherwise seminal and evocative essay *The Cathedral and the Bazaar* [1] is the one prefiguring the demise of Brooks’ Law within F/OSS development. This view is supported by a *reductio ad absurdum* argument, claiming that, if Brooks’ Law were at work, it would not be possible to observe such a thing as Linux development. Conversely, the observation of the Linux case study suggests to the author that the effects of Brooks’ Law may be overcome by other forces, such as the project leader’s capabilities in attracting, motivating and coordinating a team of skilled and talented developers, in a distributed process strongly facilitated by Internet connectivity as a shared medium of communication. This argument, that Brooks’ Law does not apply to Internet-based distributed development, has been widely criticized by many authors (see for instance Bezroukov [26] and Jones [27]).

Modularity allows us to refine and clarify these criticisms suggesting that a large number of participants in a project may be not a sufficient condition to generate dysfunctional effects such as diminishing or negative marginal return of manpower to productivity. The key aspect, in this respect, is represented by the degree of task interdependency between the various members belonging to the project. Thus, the high productivity experienced in the GNU/Linux development is interpreted as largely due to the massively modularized structure of the project,

enabling the existence of highly independent sub-projects joined by a limited number of developers, resembling in essence the theory of Brooks' surgical (small, skilled and focused) team [14], while the role of the Internet in this interpretation is of mere medium of exchange allowing distant communication.

Actually, our claim seems to be straightforward if we look at a typical sub-project within the GNU/Linux architecture. Furthermore, if we look more generally at the world of F/OSS projects, there is growing empirical evidence showing that the number of participants involved in a project is on average very small [28, 29]. Despite this, in some specific cases, such as in the development of the kernel for the GNU operating system, that has been undertaken thanks to the coordinated effort of hundreds of contributors, we need to clarify our point and to address the relationships between Brooks' Law and division of labor in the case of vertical division of labor.

C. Vertical division of labor and organization and architectural ladders

Another rather famous postulate in Raymond's *The Cathedral and the Bazaar* is the following:

"I had been preaching the UNIX gospel of small tools, rapid prototyping and evolutionary programming for years. But I also believed there was a certain critical complexity above which a more centralized, a priori approach was required. I believed that the most important software (operating systems and really large tools like the Emacs programming editor) needed to be built like cathedrals, carefully crafted by individual wizards or small bands of mages working in splendid isolation, with no beta to be released before its time. Linus Torvalds's style of development – release early and often, delegate everything you can, be open to the point of promiscuity – came as a surprise. No quiet, reverent cathedral-building here – rather, the Linux community seemed to resemble a great babbling bazaar of differing agendas and approaches (aptly symbolized by the Linux archive sites, who'd take submissions from anyone) out of which a coherent and stable system could seemingly emerge only by a succession of miracles." [1]

While finding this quote intriguing and insightful in many senses, since it clearly describes the evolutionary dynamics nature of GNU/Linux development [30], we also are convinced that it conveys many misleading interpretations of the F/OSS phenomenon as a whole. Many authors have criticized the cathedral versus bazaar metaphor. We hereby are particularly concerned with a serious and common misinterpretation of this metaphor when it comes to the topic of the architectural characteristics of GNU/Linux.

The misinterpretation of the above quote runs, slightly simplifying, as follows: GNU/Linux comes out of the blue from a chaotic mess of contributions and organizes itself as a coherent system in an apparently self-regulating way, showing a mysteriously spontaneous order. This emergent view of the genesis of GNU/Linux is misleading in that it suggests the existence of a deregulated and emergent flat architecture. In contrast, we claim that the modular architecture of GNU/Linux is characterized by being quite hierarchical, rather than flat.

Basically, it boils down to the distinct possibility of distinguishing at least two different and hierarchically ordered ladders in GNU/Linux: a higher level, the kernel space, and a lower one, the user space. As it happens, the celebrated babbling bazaar, representing the decentralized and anarchic distributed process, takes place at the user level and it is fostered by the highly modular architecture, as described formerly. Conversely, at the higher inner level of the operating system, the development process seems to be rather different: Linux inner core started to be developed as a one-person hack and only at a subsequent stage of the process contribution from other developers were introduced. Moreover, while contributions to the kernel represent an open process, the integration of code within the kernel has been a process firmly regulated by the same Torvalds, at the beginning, and later supported by a small group of "trusted lieutenants" [31, 32].

In order to preserve integrity and coherence within the most important and complex part of the system, at the kernel space ladder all initial relevant design decisions were largely taken by Torvalds and by an inner team of developers. The same holds for most of the subsequent activities of kernel development. While one has to acknowledge the role of code contribution from the bottom (the hacker community), it is also indisputable that its incorporation in the project has been fueled by a highly structured and hierarchical process of review and selection (albeit not based on formal authority but rather on competence and reputation). In other words, one of the basic features of modular product architectures, namely the isomorphic relationship between product architecture and organization traits, seems to characterize the evolution of GNU/Linux that emerged as a stable system not by a succession of miracles, but rather by exploiting modularity at the user space level, encouraging decentralization, and carefully crafting and controlling the overall consistency of the design at the kernel space, imposing a cathedral-like hierarchy in code evaluation and integration.

To summarize our point, we find the cathedral vs. bazaar distinction seriously misleading. Hence, if one really wants to compare the GNU/Linux architecture to a bazaar-like structure, he should not look at an ordinary bazaar, but rather at Kapali Carsi, Istanbul Grand Bazaar, the oldest (15th century) and largest (over 4,400 shops on 30 hectares of land) marketplace of the world. The most prominent and uncommon feature of this marketplace is that it is not uncovered and out in the open as usually bazaars are. On the contrary, it is a covered structure owning a complex architecture protecting a giant labyrinth of shops and various commercial activities. It has been observed by many that the covered architecture is a fairly regular structure, which makes the underlying bazaar even more maze-like and confusing in practice. Just as the building architecture is not affected by the underlying bazaar activities, likewise, GNU/Linux higher ladder, i.e. the kernel, is largely shielded from decentralized evolutionary dynamics happening at the user space level.

We have until now emphasized that the GNU operating system is a massive modular architecture, mostly inherited from a previous design and characterized by a hierarchical two-ladder architecture that hardly resembles the flatness of the common bazaar at all. To further refine our analysis we need to admit that, albeit largely based on the UNIX architecture, there does exist something truly innovative

and original in the GNU operating system. This pertains to its kernel. In the following we reflect on its origins, highlighting the different approaches on modularity and interdependencies decomposition followed by two different competing projects: the Linux project and the HURD project.

D. Ex-ante modularity versus evolving modularization: the development of a kernel for the GNU operating system

Along with Simon's perspective, we mentioned that the decomposition of complex problems in nearly-independent sub-problems (i.e. modules) is a complex activity itself [34]. At the beginning, designers do not know precisely how to conceptualize the modules of new artifacts; later, when a first conceptualization is reached, they still vaguely know how good is the chosen architecture, compared to the other that have not been considered.

If we underestimate the problems posed by modules identification and decomposition of new architectures, we hardly understand why modular design of complex products is so difficult and unpredictable. Another way to grasp this issue is to consider that many modular products were originally developed from interconnected solutions. While this is not a general rule, it was definitely true for Linus Torvalds's kernel: the GNU operating system is known for being a modular complex artifact and its successful development, accomplished by a distributed community of hackers, largely benefited from that. Therefore, it may be surprising that its core-component, the so-called kernel, was initially conceived as a highly integrated product and only eventually acquired a modular structure. As a developer, Linus Torvalds' major effort to the project afterward called GNU/Linux was aimed to conceive and write the kernel, that is the core part of the operative system that could use all the applications and the libraries of software that had already been developed within the GNU project.

At the time Linus Torvalds started to work on his kernel, a long debate was mounting around the advantages offered by an alternative architecture, called microkernel, designed to work in all possible and different processors [35, 36].

Compared to traditional, hardware dedicated kernels, microkernels appeared to be more complex and less efficient. They were more complex because even simple problems were treated as instances of general tasks that might have involved a higher number of specifications and instructions to interact with other parts of the kernel; therefore, they might have resulted to be less efficient as they did not take advantage of specific features of the hardware they run on. While microkernel architecture appeared to be a better solution because of its recognized technical superiority, Torvalds decided to develop his kernel in less general terms, thinking that microkernels at the beginning of the '90 were still experimental and too complex projects (at that time Microsoft was developing its new Windows NT using a microkernel structure) and they were exhibiting a much worse performance [50]. By the way, when Torvalds started to work on its kernel the Free Software community and the GNU partisans were already involved in the development of a microkernel

(called HURD), even though the task seemed to be much far away from its conclusion.

Therefore, the very first version of Linus' kernel had a monolithic structure and was also extremely hardware specific, since it was conceived for working on Intel 80386 processors only. The first effort to port Linux kernel to another processor (Motorola 68K) showed all the drawbacks of having a hardware-specific architecture, since the developers of 68K Linux had to write another hardware-specific kernel from scratch. When Torvalds started to think about porting Linux to the Alpha platform, he realized that the original design was no longer effective and in 1993 he started to rewrite the kernel code completely. He decided to keep a monolithic architecture, but he introduced some degree of modularity in the system design, in order to simplify the portability task and to incentive parallel development in some less critical parts of the system. In Torvalds words:

"With the Linux kernel it became clear very quickly that we want to have a system which is as modular as possible. The open-source development model really requires this, because otherwise you can't easily have people working in parallel. It's too painful when you have people working on the same part of the kernel and they clash.

Without modularity I would have to check every file that changed which would be a lot, to make sure nothing was changed that would effect anything else. With modularity, when someone sends me patches to do a new filesystem and I don't necessarily trust the patches per se, I can still trust the fact that if nobody's using this filesystem, it's not going to impact anything else"[35].

Therefore, the general kernel model made use of modules and it was conceived bearing in mind those elements common to all typical modular architectures (even though it was not as rigorous and general as microkernels are). Following this scheme, Torvalds could deal with them separately and confine all the hardware-specific pieces of code in modules out of the core kernel [37]. These modules could be later updated or changed by Torvalds himself and by the other Linux developers with no effect on the kernel core. Device drivers structure is a good example of the third way followed by Torvalds. One extreme solution is to put all the hardware specific into the core kernel: this is easier to do, it increased the performance, but the kernel is totally unportable. The other extreme solution, consistent with the microkernel design, urges to leave all the specific in the user space, which declines the performance and the stability of the system.

In later discussions Torvalds explained the reasons for its choice: a fully modular architecture, like the one adopted for HURD, would have posed problems to a degree of complexity that it could have compromised the accomplishment of the project. To avoid such risks and keep the degree of complexity of the project as low as possible, Torvalds decided to design a monolith and he actually wrote all the architectural specs himself, avoiding all the problems related to collective projects (e.g. division of labor, coordination, communication). On the other hand, the HURD micro-kernel, a project in direct competition with the Linux kernel, has paid for the choice of pursuing a fully modular approach from the beginning in terms of the continuous delays that have plagued its development. Nowadays, it is still under active

development and still lacks the stability and performance assured by the Linux kernel.

The validity of Torvalds choice is under our eyes and it is difficult to overestimate the consequences of this modular solution with regard to the subsequent portability and extensibility of the system through the distributed effort of the community. Nowadays Linux run on an increasing number of computers, from workstations to handheld devices and its development is assured by the effort of tens of thousands developers in the world. Torvalds and a few other people close to him control the kernel core and have the final word in the decisions related to the development of the system. Other developers, on the other hand, offer their contribution to improve and upgrade the system. We already showed how critical were the consequences of inheriting a modular UNIX-like architecture based on complementary and interconnected components. To a more hidden and critical level the development of the core of the operative system, the kernel, followed an analogous destiny. The modular structure adopted by Torvalds for its kernel happened to be successful, nevertheless it does not prevent the system from the emergence of unforeseen interdependencies within the modules that may arise with the future development of hardware and software. While HURD established itself as an attempt to develop a fully general and modular system, Linux kernel took advantage of some architectural shortcuts: as it is, the problem related to emergent interdependencies that were not expected at the beginning may become a problem for the future enduring success of Linux, even though this can be regarded as a future cost for the straightforwardness of its design. Some of these emergent interdependencies may be solved by tinkering, reworking and re-designing [38]; sometimes the adopted solutions are not adequate and the communities of developers that do not agree with the final decision may introduce alternative versions of the system. These forks may express a coordination failure when a community does not converge on a unique satisfying solution. Further, unanticipated interdependencies may end up in more serious problems than just forks, as it happens when the existing operative systems reveal itself to be inconsistent with the architecture of new processors. Torvalds himself is fully aware of this situation when he describes a future scenario of Linux's decline:

"They'll say Linux was designed for the 80386 and the new CPU's are doing the really interesting things differently. Let's drop this old Linux stuff. This is essentially what I did when creating Linux. And in the future, they'll be able to look at our code, use our interfaces, and provide binary compatibility, and if all that happens I'll be happy." [35]

It is worthwhile to point out some observations that are suggested by this story:

- even when task partitioning and division of labor issues do not really matter, the design of modular architectures from scratch may reveal to be an extremely complex task; therefore, designers may prefer integrated solutions that are easier to devise and handle;
- a modular architecture is more vulnerable to design faults, especially when the task is complex and the amount of resources are limited. In particular, an ineffective definition of modules that are not loosely-coupled enough produces an increasing amount of interdependency, rather than its opposite. As a result, individuals, rather than

groups of developers may more efficiently accomplish the early stages of new projects. Some successful F/OSS stories experienced this destiny, as they have been started as one-man projects aimed to solve specific problems and eventually evolved in structured projects involving a large number of people (e.g., Sendmail was initially developed by Eric Allman to route email for other users within UC Berkely, Perl by Larry Wall to solve some annoying problems in system administration, World Wide Web by Tim Berners-Lee as group environment for academic information sharing among high-energy physicists, and so on).

- Torvalds' kernel story enriches the perspective offered by the Conway's law about the isomorphic structure of product and process [39]. Modularity, in facts, seems to be pursued not as a dogmatic feature of the product, but it arises as a general design rule and it is boosted only when it provides some direct advantage. Therefore, the evolution of the Linux kernel towards modular design suggests that it is possible to combine together, under the same architecture, both modular components and integrated parts. Later on, the designers may introduce a higher degree of modularity by adapting the originally interconnected architecture. In other words, modularity arises more as a process of evolutionary design (modularization), rather than as an ultimate ex-ante property of an artifact.

E. Beyond the principles of modularity

Another strategy employed by the F/OSS development style in order to lower the impact of this "dark side" of modularity is represented by modifications to the principle of information hiding.

It is worth mentioning that many scholars have radically criticized the modular approach to the design of software artifacts since its introduction. As noted by Brooks:

"Harlan Mills has argued pervasively that 'programming should be a public process', that exposing all the work to everybody's gaze helps quality control, both by peer pressure to do things well and by peers actually spotting flaws and bugs". [40]

Brooks argued that information should be completely available in order for failures in the design of software to become evident and be corrected [14]. Conversely, in accordance with the principles of modularity, these processes of peer review, control and contribution to others' source code are strongly limited by information hiding constraints, since modules are not available to other developers.

Despite these criticisms, information hiding has nowadays become almost ubiquitous in software engineering. Even Brooks, in the 20th year anniversary edition of his *The Mythical Man-Month* [40], admits the following: "Parnas was right, and I was wrong on information hiding".

We claim that the fundamental innovation of F/OSS practices lies in how the basic postulate of information hiding is adapted to overcome these pitfalls, suggesting a step further in the software engineering debate on the pros and cons of modularity. While information hiding is clearly at the core of designers' activities when initially decomposing a software project in modules, the same principle is later disregarded, at the implementation level,

in day by day coding, test and integration activities. As a matter of fact, in the F/OSS community, hackers actually are overexposed to, rather than shielded from, a huge amount of code.

The free availability of the source and the absence of code ownership make programming a truly public process, since good coding solutions are shared and adapted to solve similar problems [41], and ex-post interdependency conflicts are handled by employing a wider set of fine-tuning strategies. A well-known feature of F/OSS methodologies is parallelized and distributed code debugging, where bugs are highlighted and corrected by others' "eyeballs" [1, 42]. Kuan, for instance, shows that F/OSS has a higher rate of quality improvement than closed source software [43] and similar results are obtained by Succi et al. [44]. Jorgensen reports that half of the respondents to his research survey claimed to have received a bug report from someone else within the previous month and nearly half of them credited an external contributor fixing a bug in their code. Likewise, at the code review level, similar parallel and distributed processes of peer review highlight design incoherencies introduced by others [25].

In other circumstances the "no hiding" principle allows developers to undertake much more sophisticated software engineering activities, such as redefining modules and interfaces specifications in response to the emergence of new interdependencies between separate modules. This is often the case in the introduction of radically new or substantially complex features in stable projects. For instance, the introduction of cryptography in the Freenet project affected many different modules and demanded the whole redefinition of the architecture [60]. In the worlds of developer #101:

"[...] unfortunately, any change you make in that affects not only the protocol, which is what I am working on right now, but it affects how the keys are handled (Module 4), how the client interprets the keys (Module 8), how data is verified. Basically, that little change affects pretty much everything in Freenet and, therefore, the kind of people making those changes, myself and (developer #6) mainly, have to understand everything that happens in Freenet in order to do it." The availability of other modules source code is what allowed the two developers to disentangle the complex web of interdependencies introduced by adding a public key to cryptography. Similarly, Jorgensen underlines how the free flow of information about the whole project helped FreeBSD developers to introduce a radically innovative feature to support multiprocessing (Symmetric Multiprocessing) within a mature software architecture [25].

In short, the lesson of F/OSS development is the following: since it is impossible to design ex-ante a zero-defect software architecture, it is worthwhile to embrace adaptive and flexible strategies that ease modules integration by using all the available (not anymore hidden) information.

The no-hiding policy bears one additional consequence: it does not only elicit an iterative and distributed process by which previously written code is fine-tuned and optimized by participants in the software community, but it also makes it possible for individual hackers or entire groups to write patches or variations of the original code that are not completely compatible with previous work

carried out in the same software project or with respect to other related pieces of software. While incompatibilities are most of the time unintentional and marginal and may be fixed by subsequent coding activities, sometimes these modifications are large and/or intentional and may result in forking, i.e. the introduction of an independent and partially incompatible version of the original software. As a matter of fact, within the software industry, advocates of corporate closed source software development have argued that, due to the lack of code ownership, F/OSS seems to be particularly prone to develop "multiple incompatible versions of programs, [plagued by] weakened interoperability, [and] product instability" [46]. With respect to software development activities, this may lead to duplication of efforts and may result in an inefficient allocation of scarce resources at the level of the whole F/OSS community.

Nevertheless, other studies have suggested that forking in F/OSS may be much less frequent than one might expect at a first glance, and may eventually lead to positive, rather than catastrophic, outcomes. Many F/OSS projects have a governance structure (ranging from the project leader benevolent dictatorship to the formation of complex coalitions) that prevents attempts to fork [2]. Moreover, the widespread diffusion of the GNU General Public License (GPL), seems to mitigate the incentives to fork an existing F/OSS project since, in essence, it prevents the appropriability of innovations. In fact, while anyone may fork any software project at any time, his subsequent work, due to the GPL "infectious" nature, would be available to the whole community as well. Thus, others may take advantage of the improvements of the fork. In this perspective, forking rarely happens and even when it occurs, this often translates in being beneficial to both competing projects, since the GPL allows each one to study the other and implement the most innovative features (e.g. this seems to have been the case in the rivalry between the Emacs and the XEmacs projects [47]). As a result, forking seems to take place largely in case of ultimate and irreconcilable differences in views and priorities in the development of a software project, and forks take off and succeed only if they are able to occupy different ecological niches (see for instance the existence of various GNU/Linux distributions), thus offering specialized solutions for a differentiated audience [48]. Finally, it has been noted that it is not uncommon for forks to merge back with the original project as benefits and drawbacks of "running alone" may change overtime (as in the case, for instance, of the egcs project, re-merged by the FSF with the original gcc project in 1999) [47].

Anyway this topic calls for more rigorous and analytic case studies aimed at understanding better the advantages and drawbacks in the emergence of forking within F/OSS projects.

III. DISCUSSION

In the end, modularity may be conceived as simple as it is, as long as we do not open the black box and keep track of the organizational processes behind the structure. Most quoted contributions in management studies [7, 33, 49, 50] unfold a neat and smooth theory of modularity, introduced as a cornerstone for artifact design [10].

According to this Olympic version, modularity is defined as a "particular design structure, in which

parameters and tasks are interdependent within units (modules) and independent across them” [7 p. 88]. Unfortunately, this perspective underestimates that the decomposition of complex systems generally resolves on a quasi-decomposition and not in a full decomposition, as some interdependencies may not be predicted or are left out on purpose, simply because they are regarded as marginal ones. Our reconsideration of the development of some F/OSS projects show how the modularity principles may in practice differ from what the theory prescribes. GNU/Linux and, more generally, F/OSS represent an instance of unorthodox modularity: the information hiding principle is significantly disregarded as the artifact evolves mainly through a repertoire of practices (e.g. peer coding and debugging, frank discussions, open decisions) where developers and users work apart, tinkering and patching the original modular product and, overall, violating another of the law of the Olympic modularity stating that the only available operators are represented by manipulation at the module level (splitting, substituting, augmenting, excluding, inverting, porting [7 pp. 123–146]).

In our view, reading the GNU/Linux case according to the modularity perspective provides a complementary understanding of the F/OSS phenomenon and, at the same time, offers some insights to think about the way we conceive a theory of modularity for complex systems.

With respect to the first issue, taking advantage of existing architectures like UNIX and related standards (e.g. POSIX) it has been a successful strategy as the community of developers avoided to design a modular structure from scratch. The comparison between the HURD project and Torvalds monolithic kernel shows that to develop decomposable architectures for complex products exposes the designers to the risk of unforeseen interdependencies that may ultimately endanger the whole project. Besides, as F/OSS projects are developed by distributed organizations and the community members communicate only remotely, coordination and collective decision making seem to be two fundamental issues in F/OSS development. In other words, our study of F/OSS projects through the lens of the theory of modularity outlines three main strategies that characterize the design and the development of complex systems: i) inheriting existing modular architecture, ii) evolving towards increasing degrees of modularity and iii) violating the information hiding principle. This repertoire of practices, or shortcuts as we called them in our introduction, emerge as effective and robust routines that seem to fit very well with the actors involved, i.e. distributed communities of developers, and the problem solving activity they embrace.

GNU/Linux case, on the other hand, suggests some general reflections on modularity and modularization. F/OSS developers exploit all the advantages of a modular architecture as the massive parallel activity within modules/programs witnesses; on the other hand, the modularization does not stop with the architecture design. The unforeseen interdependencies that come to the surface as the operative system evolves, revealing some inconsistencies, are met by violations of the information hiding principle. In questioning how this experience may be extendible to other contexts where modularity has already started to represent a promising approach, there

are at least two fundamental conditions that need to be clearly spelled out. First, F/OSS distinctive trait is represented by the open access to knowledge (source code and documentation) stored in the modules. In the F/OSS world, imitation and copy are encouraged and protected by a reverse form of copyright (copyleft). According to the Economics of Innovation standard models, copyleft should inhibit any investment in innovations, since anybody may take advantage of any innovation and there are no incentives for the innovators. F/OSS apparently contravenes this rule and this is way motivational analyses based on various perspectives, i.e. psychological, cultural, sociological, seem to be urgent to support an economic explanation of this phenomenon. To our viewpoint, the apparent paradox of compelling innovations in a copyleft regime is due to the second fundamental condition that characterizes the F/OSS movement, that is a deep overlap between producers and users. At least, at the beginning, most users were developers or had some skills that allowed them to perform successful adaptations. Again, most of the traditional ways to conceive innovation and product development in other domains keeps producers and users separated, even though today customers are more and more often directly involved in the definition of their own product.

As long as developers and users communities deeply overlap, copyleft regime does not inhibit innovation, but rather it ensures its open and free diffusion. On the other hand, when the communities start to be more and more different from each other, when developers are viewed as producers and users as customers, the natural system of reciprocal benefits becomes less and less salient. Therefore, looking for a possible generalization of F/OSS experience should push us towards other economic contexts where developers and users are able to establish strong relationships; in this respect, settings where customers actively participate in the development of new products [52] seem to represent a promising milieu for empirical investigation.

IV. ACKNOWLEDGMENTS

The authors would like to thank for their helpful comments the ROCK (Research on Organizations, Coordination and Knowledge) members at the University of Trento, the participants to the track session “Modularity and division of innovative labour: design, organisation and cost analysis” at EURAM’s 2nd conference on Innovative Research in Management, May 9-11, 2002, in Stockholm and the anonymous referee. The usual disclaimer applies. Financial support from MIUR under the project COFIN 02 (“Language and coordination in managerial distributed decision making”) is gratefully acknowledged.

V. REFERENCES

- [1] E. S. Raymond. *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*. Sebastopol, CA: O’Reilly & Associates, 1999.
- [2] B. Kogut and A. Metiu. Open-Source software development and distributed innovation. *Oxford Review of Economic Policy*, 17(2):248–264, 2001.
- [3] J. Feller and B. Fitzgerald. A framework analysis of

- the open source software development paradigm. In Proceedings of the twenty first international conference on Information systems, pages 58–69. Atlanta, GA: Association for Information Systems, 2000.
- [4] M. A. Schilling. Toward a general modular systems theory and its application to interfirm product modularity. *Academy of Management Review*, 25(2): 312–334, 2000.
- [5] N.W Hatch. Modular stepping stones along the firm’s technology path. Nelson and Winter Conference Aalborg, 2001.
- [6] I. Jackson. Why is software freedom useful and what does it mean? SANE’98 (18-20 November 1998), 1998.
- [7] C. Y. Baldwin and K. B. Clark. *Design Rules. Vol. I: The Power of Modularity*. Cambridge, MA: The MIT Press, 2000.
- [8] S. Brusoni and A. Prencipe. Unpacking the black box of modularity: Technologies, products and organisations. *Industrial and Corporate Change*, 10 (1):179–205, 2001.
- [9] M.G. Devetag and E. Zaninotto. The imperfect hiding: Some introductory concepts and preliminary issues on modularity. DISA Working Paper, Università a degli Studi di Trento, 2001.
- [10] R. Garud, A., Kumaraswamy, R.L. Langlois. Eds. *Managing in the modular age*. Malden, MA: Blackwell., 2003.
- [11] R. N. Langlois and P. Robertson. Networks and innovation in a modular system: Lessons from the microcomputer and stereo component industries. *Research Policy*, 21(4):297–313, 1992.
- [12] A. Camuffo. Rolling out a “world car”: Globalization, outsourcing and modularity. 2nd EURAM Conference, Stockholm, Sweden, 2002.
- [13] H. A. Simon. *The Sciences of the Artificial*, 2nd ed. Cambridge, MA: The MIT Press, 1981.
- [14] F. P. Brooks. *The Mythical Man–Month. Essays on Software Engineering*. Reading, MA: Addison Wesley, 1975.
- [15] E. von Hippel. Task partitioning: An innovation process variable. *Research Policy*, 19(5):407–418, 1990.
- [16] D. L. Parnas. On the criteria for decomposing systems into modules. *Communication of the ACM*, 15(12): 1053–1058, 1972.
- [17] F. P. Brooks. No silver bullet. In H. J. Kugler, editor, *Information Processing 1986, Proceedings of the IFIP Tenth World Computing Conference*, pages 1069–1076. Amsterdam: Elsevier Science, 1986.
- [18] H.A. Simon. *Models of Man*. New York, NY: Wiley, 1957.
- [19] D.M. Ritchie and K. Thompson. The UNIX time-sharing system. *Communications of the ACM*, 17(7): 365–375, 1974.
- [20] M. Gancarz. *The UNIX Philosophy*. Newton, MA: Digital Press, 1994.
- [21] H. P. Salus. *A Quarter Century of UNIX*. Reading, MA: Addison–Welsey, 1994.
- [22] S.C. Johnson and D.M. Ritchie. Portability of C programs and the UNIX system. *The Bell System Technical Journal*, 57(6):2021–2048, 1978.
- [23] R. Miller. UNIX – a portable operating system? *ACM Operating Systems Review*, 12(3):32–37, 1978.
- [24] D. K. Rosenberg. *Open Source. The Unauthorized White Papers*. Foster City, CA: IDG Book Worldwide, 2000.
- [25] N. Jorgensen. Putting it all in the trunk: incremental software development in the Free BSD open source project. *Information Systems Journal*, 11(4):321–336, 2001.
- [26] N. Bezroukov. A second look at the cathedral and bazaar. *First Monday*, 4(12), 1999.
- [27] P. Jones. Brooks’ law and Open Source: The more the merrier?, 2000. Retrieved Jan 2, 2003, (<http://www-106.ibm.com/developerworks/library/merrier.html>).
- [28] S. Krishnamurthy. Cave or community? an empirical examination of 100 mature Open Source projects. *First Monday*, 7(6), 2002.
- [29] A. Capiluppi, P. Lago, and M. Morisio. Characterizing the oss process: a horizontal study. 7th European Conference on Software Maintenance and Reengineering, Benevento, 2003.
- [30] K. Kuwabara. Linux: A bazaar at the edge of chaos. *First Monday*, 5(3), 2000.
- [31] E. Franck and C. Jungwirth. Reconciling investors and donors. the governance structure of open source. *Lehrstuhl für Unternehmensführung und politik Universität Zürich*, 2002.
- [32] G.N. Dafermos. Management and virtual decentralised networks: The linux project. *First Monday*, 6(11), 2001.
- [34] L. Marengo, C. Pasquali, and M. Valente. Decomposability and modularity of economic interactions. In W. Callebaut, editor, *Modularity: Understanding the Development and Evolution of Complex Natural Systems*. Cambridge, MA: The MIT Press, 2001.
- [35] L. Torvalds. The Linux edge. In C. DiBona, S. Ockman, and M. Stone, editors, *Open Sources: Voices from the Open Source Revolution*. Sebastopol, CA: O’Reilly & Associates, 1999.
- [36] C. DiBona, S. Ockman, and M. Stone. *Open Sources: Voices from the Open Source Revolution*. Sebastopol, CA: O’Reilly & Associates, 1999.
- [37] J.M. de Goyeneche and E. Apolinario Fernández de Sousa. Loadable kernel modules. *IEEE Software*, 16 (1):65–71, 1999.
- [38] M. Ratto. Re–working by the linux kernel developers. Department of Communication, University of California, San Diego, 2003.
- [39] M.E. Conway. How do committees invent. *Datamation*, 14(10):28–31, 1968.
- [40] F. P. Brooks. *The Mythical Man–Month. Essays on Software Engineering, Anniversary ed*. Reading, MA: Addison Wesley, 1995. [56] R. C. Pavlicek. *Embracing Insanity: Open Source Software Development*. Indianapolis, IN: Sams Publishing, 2000.
- [42] F. Iannacci. The linux managing model. Proceedings of the Third International Conference on Open Source, ICOS 2003, 2003.
- [43] J. Kuan. Open source software as consumer integration into production. Retrieved Jul 12, 2003, from <http://opensource.mit.edu>, 2000.
- [44] J. Succi, G. Paulson and A. Eberlein. Preliminary results from an empirical study of open source and

- commercial software products. International Conference on Software Engineering, Toronto, Ontario, Canada, 2001.
- [45] G. von Krogh, S. Spaeth, and K.R. Lakhani. Community, joining, and specialization in open source software innovation: a case study. *Research Policy*, 32(7):1217–1241, 2003.
- [46] C. Mundie. The commercial software model, 2001. Retrieved July 1, 2003, (<http://www.microsoft.com/presspass/exec/craig/05-03sharesource.asp>).
- [47] R. Moen. Fear of forking essay, 2003. Retrieved July 1, 2003, from <http://linuxmafia.com/Erick/essays/forking.html>.
- [48] R. van Wendel de Joode, H. de Bruijn, and M. van Eeten. Protecting the virtual commons. Self-organizing communities and innovative intellectual property regimes, 2003. NWO/ITeR series.
- [49] C. Y. Baldwin and K. B. Clark. Managing in the age of modularity. *Harvard Business Review*, 75(5):84–93, 1997.
- [50] K. Ulrich. The role of product architecture in the manufacturing firm. *Research Policy*, 24:419–440, 1995.
- [51] R. Sanchez and J. T. Mahoney. Modularity, flexibility, and knowledge management in product and organizational design. *Strategic Management Journal*, 17(winter special issue):63–76, 1996.
- [52] E. von Hippel. Economics of Product Development by Users: The Impact of "Sticky" Local Information. *Management Science*, 44(5):629–644, 1998.